STL PARTNERS

Executive Briefing

GEN AI: WHERE SHOULD TELCOS START?

Gen AI has great potential usefulness to telcos, provided its capabilities and limitations are well understood, and projects implemented judiciously. This report directs telcos to the most promising use cases.

Charlotte Patrick, Associate Analyst | charlotte.patrick@stlpartners.com | September 2023

Executive Summary

After the hype of Generative AI (Gen AI)

The appearance of ChatGPT in early 2023 created a surge of interest in Gen AI and speculation that it will be deployed across telcos – pushing forward automation and significantly increasing the levels of intelligence deployed within an organisation. Interest has since declined as those that have tried ChatGPT have realised that, although a genuinely revolutionary technology, Gen AI still has some way to go if it is to be a reliable part of automation or the day-to-day working life of staff in a telco.

This research tackles the complex question of what Gen AI projects telcos and telco vendors should invest in to optimise short- to medium-term returns. This decisioning is impacted by a mix of factors, including the maturity of Gen AI models plus the maturity of other AI models, which it will work in conjunction with. Also, the maturity and data availability of the processes in which Gen AI will be introduced. Other factors are commercial: Where is Gen AI likely to create a significantly better customer experience or more revenue opportunities?

The most compelling uses for Gen AI in telcos

Use cases that are expected to deliver short-term benefits to the telco include:

- **Text and image generation for teams creating content for customers**, such as marketing and sales. These tools are already available from specialist vendors and provide, at minimum, a starting point for content creation and a range of viable options from which the content creator can choose.

- **The use of large language models (LLMs) to generate snippets of code within developer tools**. This will be a welcome development for telcos, which often struggle to attract sufficient developer talent. Quoted reductions in the time spent coding vary, but it opens future options for telcos to develop more in house.

- **The use of models for pattern and anomaly detection**. This is a telco-wide requirement that benefits many areas, particularly the network where data sets are large and must be monitored for real-time issues. Gen AI models bring improvements to anomaly detection.

- **Generating synthetic data (new data sets built from existing data) and augmenting data sets that are currently incomplete**, thanks to Gen AI's ability to create novel content. Use cases vary, but a typical example would be the ability to create better-quality training data for predictive models where the training data has previously been missing a minority class.

- **Improving telco knowledge bases and catalogues**. Models will be able to use their summarisation capabilities to create bespoke deliverables from these stores and can improve existing functions such as frequently asked questions (FAQs), which will automatically update as needed.

- **The use of LLMs to improve the quality of interactions with chatbots, interactive voice responses (IVRs) and digital assistants**. This offers a range of new capabilities. For example, interactions can be multilingual, more "human," more engaging and answering more complex questions. These benefits will take time to develop and deploy, but the area should be a short-term focus for telcos as it will be a long road to potentially significant savings in customer services.

How to scope a Gen AI project

These technologies are costly. A recent Mckinsey article[1] gives an estimated one-time cost of between US$0.5 million and US$2 million for an off-the-shelf coding assistant, US$2 million to US$10 million for a model which is fine-tuned with sector-specific information and US$5 million to US$200 million for the training of a foundational model (costs varying dependent on the complexity of the task).

This level of costs will not provide an immediate return on investment if the model output quality is low and integration costs are high.

Telco decision-makers will need to spend a significant and ongoing amount of time in-market understanding how to enable proper risk assessments as new models are rapidly delivered to provide brand new capabilities and improve existing ones.

Many models will likely be trialled by humans who find they receive only limited value in the short term. However, the building of familiarity will be necessary, allowing more rapid deployment when models mature and are ready for integration into "high-stakes" environments where they can create customer experience issues.

Gen AI's ability to generate code snippets in core telco areas, such as network function deployment, will create interesting strategic questions. In what timeframe will telcos be able to become more self-sufficient in areas currently dominated by specialist vendors?

4 core recommendations

1. **Develop Gen AI knowledge levels across the company** so that everyone is clear about what it can and cannot accomplish.

2. **Use the workforce's familiarity with the concept of Gen AI to encourage the trialling of new tools**. In particular, the creation of draft versions of the content or first iterations of larger tasks that can provide a speedy start to new projects.

---

[1] https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/technologys-generational-moment-with-generative-ai-a-cio-and-cto-guide#/

3. **Prepare for a range of deployment scenarios.** Telcos will deploy Gen AI using one of five approaches, each with different parties involved in purchasing and managing solutions. This range of deployment scenarios is not unusual in a telco but given the nascent nature of Gen AI as a technology, telco IT teams need to be ready:

   a) Specialist applications, particularly in areas such as content creation, will likely not need significant integration and be purchased as SaaS. However, centralised oversight will be needed to manage risks such as data privacy.

   b) Gen AI providing minor augmentations to existing models on vendor platforms (e.g., generation of content within a campaign management tool) will come as upgrades, and vendors should be asked about short-term roadmaps to ensure benefits can be taken in the short term.

   c) Models will improve the data and processes which support other models and automations (e.g., generation of synthetic training data). These may come from specialist vendors in, say, data management but have ROI impact across different parts of the organisation – making them a focus for the Chief Data Officer.

   d) A set of core models will be considered necessary to own – requiring a new platform team to provision and manage approved models. This team will need a range of Gen AI skills as well as broader skill sets such as integration and risk management.

   e) A small number of Gen AI projects will be complex but potentially high-value and will require cross-team focus, potentially for relatively long periods as the models grow in sophistication (e.g., chatbot and IVR customer experience).

4. **Calculate the exact needs for human intelligence in new Gen AI deployments.** OpenAI hired 6,000 individuals to label an appropriate subset of data for its ChatGPT application. Obviously, telcos are unlikely to want to create their own foundational models, but the exact requirements for human intelligence in new Gen AI deployments should be carefully considered.

Next steps

Growing use of artificial intelligence, including Gen AI, has significant implications for telco organisations. Upcoming research from STL looking at what skills telco most need to acquire will help build understanding of how to best harness these technologies. STL is also planning a report looking at how to build trust and acceptance of AI in telco organisations.

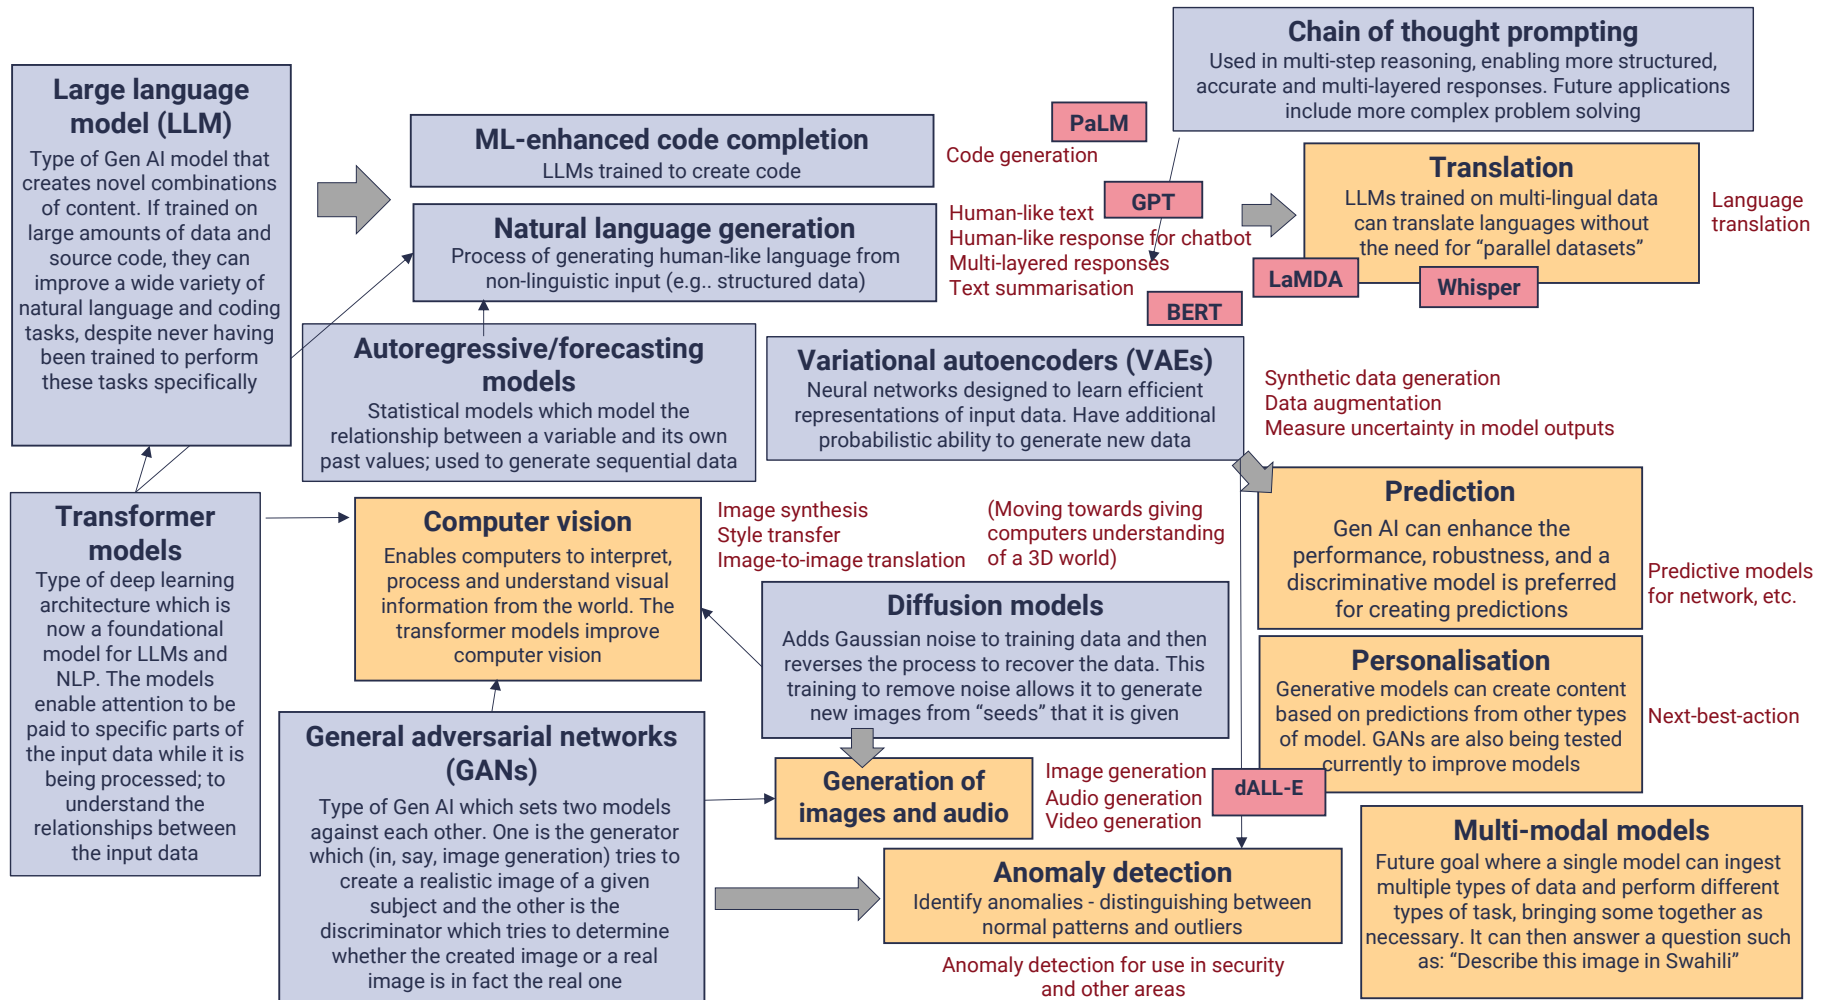# Table of Contents

Table of Figures

Introduction

ChatGPT defines Generative AI as:

> *"a class of artificial intelligence models and algorithms that have the capability to generate new, original content. Unlike traditional AI models that operate based on predefined rules or patterns, generative AI models can produce novel data that resembles the patterns observed in the training data. Generative AI models can generate content in various domains, such as natural language, images, music, and more".*

Figure 1 describes the range of models, concepts and uses that are seen in discussions around Gen AI.

- Blue boxes describe the main models, architectures and concepts that underpin various Gen AI capabilities (e.g., large language models).

- Orange boxes describe the general capabilities of these models (e.g., natural language generation).

- Grey arrows show the main models used to create capabilities in an orange box, and smaller black arrows show where other models can also be used (e.g., diffusion models provide image generation capabilities).

- Red text gives some of the uses made of the capabilities shown in the yellow boxes (e.g., generation of novel text).

- Red boxes highlight some of the popular foundational models for these uses (e.g., ChatGPT).

Figure 1: Concepts in Gen AI

Large language model (LLM)
Type of Gen AI model that creates novel combinations of content. If trained on large amounts of data and source code, they can improve a wide variety of natural language and coding tasks, despite never having been trained to perform these tasks specifically

ML-enhanced code completion
LLMs trained to create code

Code generation

PaLM

Chain of thought prompting
Used in multi-step reasoning, enabling more structured, accurate and multi-layered responses. Future applications include more complex problem solving

GPT

Natural language generation
Process of generating human-like language from non-linguistic input (e.g.. structured data)

Human-like text
Human-like response for chatbot
Multi-layered responses
Text summarisation

LaMDA

BERT

Translation
LLMs trained on multi-lingual data can translate languages without the need for "parallel datasets"

Whisper

Language translation

Transformer models
Type of deep learning architecture which is now a foundational model for LLMs and NLP. The models enable attention to be paid to specific parts of the input data while it is being processed; to understand the relationships between the input data

Autoregressive/forecasting models
Statistical models which model the relationship between a variable and its own past values; used to generate sequential data

Variational autoencoders (VAEs)
Neural networks designed to learn efficient representations of input data. Have additional probabilistic ability to generate new data

Synthetic data generation
Data augmentation
Measure uncertainty in model outputs

Computer vision
Enables computers to interpret, process and understand visual information from the world. The transformer models improve computer vision

Image synthesis
Style transfer
Image-to-image translation

(Moving towards giving computers understanding of a 3D world)

Prediction
Gen AI can enhance the performance, robustness, and a discriminative model is preferred for creating predictions

Predictive models for network, etc.

Diffusion models
Adds Gaussian noise to training data and then reverses the process to recover the data. This training to remove noise allows it to generate new images from "seeds" that it is given

Personalisation
Generative models can create content based on predictions from other types of model. GANs are also being tested currently to improve models

Next-best-action

General adversarial networks (GANs)
Type of Gen AI which sets two models against each other. One is the generator which (in, say, image generation) tries to create a realistic image of a given subject and the other is the discriminator which tries to determine whether the created image or a real image is in fact the real one

Generation of images and audio

Image generation
Audio generation
Video generation

dALL-E

Multi-modal models
Future goal where a single model can ingest multiple types of data and perform different types of task, bringing some together as necessary. It can then answer a question such as: "Describe this image in Swahili"

Anomaly detection
Identify anomalies - distinguishing between normal patterns and outliers

Anomaly detection for use in security and other areas

Source: Charlotte Patrick Consult

Definitions of terms in the graphic:

- **Generative models** create something new based on examples they are given.

- **Foundational models** introduce a significant breakthrough, a new architecture or a novel approach that paves the way for subsequent advancements in the field.

- **Parallel dataset** is a data set which provides exact translations of all words in one language to the other.

- **Discriminative model** is a type of machine learning or statistical model that classifies input data points into different categories or classes.

- **GPT** (Generative Pre-trained Transformer) is a foundational model which can generate text responses.

- **LaMDA** is a Google project to provide a language model designed to allow more free-flowing conversations.

- **dALLe** is an OpenAI system that creates realistic images and art from natural language.

- **Whisper** is an automatic speech recognition system with improved ability to understand accents, technical language and background noise.

- **BERT** (Bidirectional Encoder Representations from Transformers) is a transformer-based language model for text generation.

- **PaLM** is a Google model that works on advanced reasoning tasks including code, mathematical problems, classification and question answering.

What does Gen AI bring to the telco?

Given the level of hype around Gen AI and its potential to outstrip human-created content, it is useful to begin by summarising its general capabilities and limitations:

Strengths

- Generation of novel content (text, image, audio, video and 3D).

- Support for a range of other content-related tasks (e.g., language translation).

- Ability to generate a set of first thoughts/rough plans/ideas for humans to work from.

- Support for the curation, management and consumption of information (e.g., ability to summarise).

- Improvement of natural language processing.

- Detection of anomalies.

- Creation of synthetic data and augmentation of data sets to improve/complete them.

- Future abilities to generate code and supporting capabilities, such as being able to explain what the code does and debug it.

- Ability to support process creation (e.g., documentation of new processes).

Weaknesses

- Gen AI models make mistakes and are prone to issues such as drift.

- Like all models, they are only as good as the data they consume.

- Natural language inference has some ability to make decisions and predict the next-best action, but it is unlikely to be the core model type used in future.

- Offers support to other AI/ML, such as predictive models, but cannot replicate their capabilities.

- Can generate snippets of code but is a distance from writing software.

- Can add capabilities to processes and automation (e.g., Gen AI can generate new application programming interface [API] calls) but is not a core automation technology.

- Can only solve simple mathematical problems.

Gen AI use cases in a telco

Having understood the potential, and also the limitations, of Gen AI, through looking at vendor websites, academic papers and industry commentary, potential use cases that can benefit telcos have been identified and catalogued. The analysis in this report is based on a database of 95 potentially beneficial telco use cases for Gen AI. These use cases can be categorised under seven groupings - seven ways that telcos can benefit from Gen AI. In this section we explain these seven use case types, how Gen AI will create benefits, and where they will be of most use in the telco organisation.

1. Content creation

There are many start-ups with products using Gen AI to create novel text, images, video, audio and 3D models. Although, when trying these solutions, the first iterations often have limited usefulness as they attempt to mimic sophisticated human activities with variable success.

Text-based content generation has the broadest set of applications, with products for marketing, sales, contact centre and digital channels. Moving to image creation, it is possible to articulate a number of potential uses for a telco, including personalised marketing and network design. Video creation is similar where applications are niche, and the capabilities still in their infancy.

Several content-related capabilities sit alongside these content-generating products. For example:

- The ability to summarise documents, allowing quick review by the user.

- Improvements in language translation.

- General adversarial networks (GANs) can transfer the style of one design onto another, merging different aesthetics and creating innovative combinations. Diverse use cases include the cloning of AI-generated voice with different human emotions.

- In the future a single multi-modal Gen AI model will generate a mix of content types (e.g., text, images or audio) or even a combination. For example, a single model could create images with descriptive text captions or presentations that include both text and images.

2. Human-machine interactions

Natural Language Understanding (NLU), Natural Language Processing (NLP), Natural Language Generation (NLG) and Natural Language Inference (NLI) are related techniques deployed in chatbots, IVRs and digital assistants. Figure 2 outlines the types of language generation models.

Figure 2: Types of language generation models
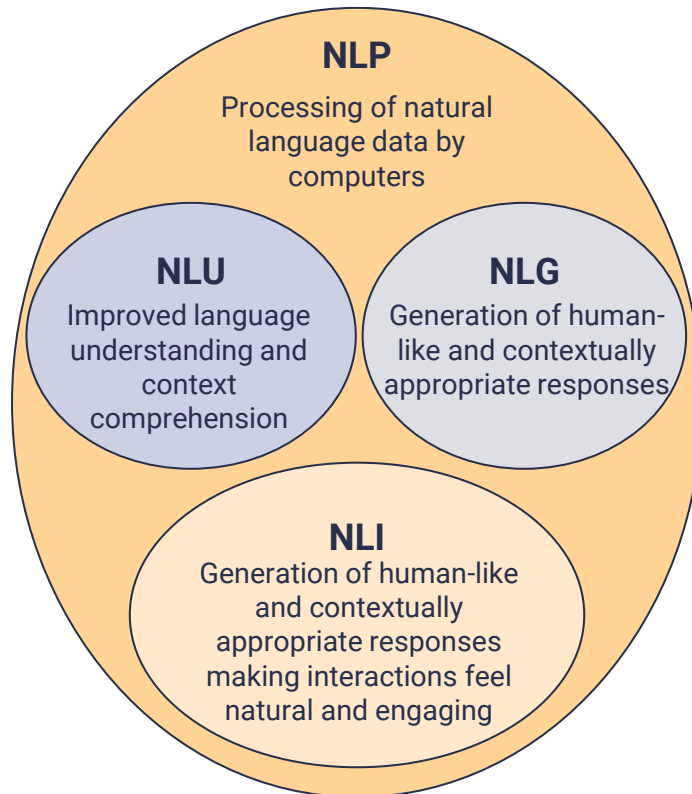
Source: Charlotte Patrick Consult

The use of Gen AI in these language processing models will create significant new value for telcos, with GPTs improving the understanding and interpretation of user input by providing better language understanding and generating contextually appropriate text.

Like many use cases for Gen AI, it will also work with various other models to provide a complete solution. For example, NLG already includes the use of models for sentiment analysis and these will benefit from the deployment of Gen AI in several ways:

• LLM will analyse the data context, considering domain-specific knowledge, to effectively differentiate between positive, negative and neutral sentiments

• GAN will help clean and preprocess unstructured data by removing irrelevant information

• GAN will also generate synthetic text data which can be used to augment existing training datasets or create new ones

NLI (also known as Recognizing Textual Entailment) is part of the solution for machine translation (the automated process of translating text/spoken language from one natural language to another), question-answering and text summation applications. It has been discussed as a significant future driver of more intelligent machines. However, it is more likely just one part of the solution for future sophisticated decisioning/prediction activities leading to human-like and detailed responses.

3. Human-human interactions

In addition to improving interactions between humans and machines, Gen AI has implications for enabling better human interactions. Examples are mainly from the contact centre environment and include the improvement of digital assistants, allowing agents to access good-quality responses to customer questions in a shorter time. Also, improving processes underneath human interactions, for example, using models to accomplish automatic call summarisation and topic extraction that can update the CRM with instant insight about a customer call.

4. Knowledge management

Telco knowledge management solutions (KMS) create, curate, manage and distribute knowledge, making it accessible to different teams, including network support, customer services, sales and field services. If (and this is a significant if) telcos can build the accurate, complete and up-to-date data sets on which GPTs need to be trained, they may significantly change how a KMS can access and consume information. Impact areas include:

- **Knowledge generation**

    – automatic creation of knowledge articles from existing data sources, such as product documentation, customer support tickets and employee training materials

- **Knowledge management**

    – identification of gaps and patterns in knowledgebase, providing new insight or highlighting potential improvements

    – automatic tagging and categorisation of documents, articles and data

    – analysis of large datasets to identify emerging trends and patterns (for example, helping customer service teams anticipate and respond to changing customer needs)

- **Assistance**

    – improvement of digital assistants that provide access to information in a KMS

    – new functionality to allow a human to ask complex questions (e.g., "What went wrong last week?")

    – The automatic creation of FAQs and their corresponding answers – updated in real-time

5. Process improvements

Code generation

Gen AI can generate snippets of code and make them available via a digital assistant, reducing the time developers need to look up boilerplate code. One example of benefits seen comes from Google: its ML-Enhanced Code Completion uses LLMs trained on source code (instead of natural language text) to improve developer efficiency. It reports that its 500-million parameter language model (used by a cohort of 10,000 software developers) provides 2.6% of the code, reducing coding iteration time for these developers by 6%.

This is the beginning of a shift towards coding being undertaken by machines, leaving developers to be more strategic. But in the mid-term, it is unlikely that Gen AI will be able to generate a complete program automatically, as it will still require human input and the usual testing cycles; plus, a complete program is complex, and it will be a significant challenge to create a complete program from a simple set of text prompts.

In the meantime, Gen AI adds new capabilities to code development, such as bug identification and the ability to explain code. There are also code-development use cases where relatively simple code is already being generated in real-time, such as API calls and policy management changes.

There is high interest from telcos in this capability, particularly where:

- small changes to code are needed, and it saves money to make them internally.

- telcos wish to develop their own solutions but lack developer talent.

- telcos want to, in future, become less dependent on vendors (e.g., management of network functions).

Improved process management

This is a catch-all description for potential Gen AI use cases that help to improve processes. Several suggested use cases have been seen:

- In the short term, the accurate documentation of standards and procedures when processes are created or changed.

- In future, Gen AI could generate small changes and develop new ideas for process improvements. For example, the ProcessGPT model (currently under development) aims to suggest the best next step in the current process pipeline.

- GANs could also be used to transfer good practice from one process to another.

Validation and testing

Lastly, use cases around validation and testing are included within this category. The strength of Gen AI in anomaly detection and the creation of synthetic data provides opportunities in this area. Real-world testing is expensive and time-consuming; risks develop when test data sets are not kept up to date, leaving the network or other equipment under test vulnerable. Gen AI can generate test data samples more quickly and efficiently than current solutions and GANs can be used for adversarial testing to test the network, and other models deployed in the network, against attack.

6. Data management

This category attempts to capture all use cases related to data gathering, generation and management.

Synthetic data generation

GANs and variational autoencoders (VAEs) can create synthetic data that resembles or augments a real data set. Synthetic data does not use (and potentially expose) sensitive data of companies and individuals, and it can act as a substitute for training data where a good quality data set is hard to obtain or is deficient in certain types of data (e.g. outliers) – improving model accuracy. However, it is prone to bias and privacy issues (if there is a model overfit while generating the synthetic data), and it may also not adequately represent real-world conditions.

As well as generating datasets, Gen AI offers a range of additional capabilities around the generation of synthetic data; these include:

- generation of synthetic samples for a minority data class, improving a model's ability to predict rare events accurately

- the augmentation of datasets by creating additional samples with slight variations

- "data synthesis" describes the creation of new composite datasets to create larger and more diverse datasets

- synthetic data can model a normal distribution so that anomalies can be detected when the real data significantly deviates

Data accessibility

Another general category of data-related use cases for Gen AI is around "accessibility" - using its image and text generation capabilities to improve a human's experience of using data sets. Opportunities include:

- the creation of new types of visualisations and summaries of the data, making it easier for data analysts and scientists to explore and understand complex datasets

- the generation of descriptions and summaries of aggregated data, making it easier to work with

Data quality

The main vendors involved in data management typically discuss a couple of Gen AI use cases, including adding natural language capabilities to their existing tools; but do not yet describe a comprehensive list of uses across the breadth of data management capabilities. This may be because Gen AI is not the best tool for the, mostly, more rules-based models. However, the use of anomaly detection to learn the normal patterns within data sets and identify outliers that deviate from the learned distribution should be useful in data quality exercises – which remains a significant issue for telcos.

7. Intelligence improvements

The final category includes all uses of Gen AI to improve the functionality of models and algorithms used across the telco:

Anomaly detection

Gen AI models such as VAEs bring new anomaly detection capabilities to the telco. These models learn to capture the regular patterns and structures in a data set, making them sensitive to anomalies that deviate from these patterns. There are many uses across the telco, including detecting bad actors in security or detecting network issues for predictive maintenance.

Prediction

One significant benefit of rules-based algorithms, rather than the use of Gen AI, in predictive modelling is the increased chance of accurate predictions – especially important in areas where models are beginning to replace human decisioning. Today, Gen AI's use of unsupervised modelling techniques makes it a suboptimal tool – as evidenced by its current use as an augmentation to human activity rather than a replacement. However, it can be used to push forward the creation of accurate predictive models by enhancing them in certain areas:

- The creation of synthetic data set or improvement of an existing data set (e.g., handling a class imbalance within the data set where it adds in data from missing minority classes).

- Use in the prevention of overfitting in prediction algorithms.

- VAEs can provide measures of the uncertainty in a prediction by learning the uncertainty in the data distribution.

Personalisation

Like predictive models, Gen AI is part of the solution for improving personalisation. Its ability to generate new content can be used to create personalised messaging by analysing human-created marketing messages crafted for a particular demographic and generating new messages with the same intentions. In the academic literature, GANs are also being tested for their ability to tackle problems seen in current personalisation models including data augmentation and the issue of "data noise" (irrelevant, inaccurate or inconsistent information present in the data).

Other model improvements

Gen AI offers a range of improvements to other models:

- **Proposal and generation of code**. For example, Gen AI can contribute to neural network architectures by suggesting innovative variations to its layers, activation functions and optimisation algorithms

- **Creating multiple simulated scenarios** based on existing data to enable the exploration of different possibilities and outcomes.

Figure 3 maps these use types onto the telco organisation to identify where they should be considered.

Figure 3: Gen AI use cases in a telco

	Network management	IT	Operations	Channels	Marketing and sales	Other
1. Content creation		Design and optimisation (e.g. network)			Content creation – text, images, video, audio, 3D	
					Support for content creation – suggestion of alterative ideas, editing and proofreading, style transfer, language translation	
2. Human-machine				Chat bots, IVR, digital assistant response improvement – improved response to questions, personality, human intent recognition, sentiment analysis, flow and response variation		Accessibility
3. Human-human				Improved human interactions – sentiment analysis, translation, customer profiling		
				Training – simulations, personalise learning, assessment, post-interaction analysis	Process improvements – quality management, transcriptions, call monitoring	
4. Knowledge management		Knowledge stores (including product manuals) – creation of knowledge articles using summarisation				
		Improved access to knowledge – bringing improvements to digital assistants, automated FAQs				
		Knowledge management – validation, gap analysis				
5. Process improvements	Security and privacy – synthetic data, testing	Code creation – creating snippets				
		Improving process management – generate simple API calls and process documentation, suggest small changes				
		Validation and testing – test data creation, anomaly detection, adversarial testing				
6. Data management		Synthetic data generation and data augmentation (improving data sets with additional data)				
		Data management – data quality improvements using anomaly detection, possible improvements in compression and other tasks				
		Data accessibility – data exploration, summarising, visualisation				
7. Intelligence improvements		Anomaly detection – learning normal behaviour, scoring, handling novel anomalies				
		Speed up delivery of predictive models – prevent overfitting, estimates uncertainty, data augmentation		Improvements to personalisation – content generation, data augmentation, noise reduction	Computer vision improvement	

Most impactful Gen AI solutions for the telco

Source: Charlotte Patrick Consult

Where is the value in Gen AI?

Important types of Gen AI

This section tries to estimate the value different Gen AI uses can bring to telcos, to help them begin prioritising. It starts with a piece of modelling that considers the value of each category in Figure 3. The bubbles in Figure 4 represent each of these categories, and the graphic is generated from a database of 95 use cases for Gen AI, which were placed in one of the categories and then scored against two measures:

1. **Likely availability.** The y-axis measures the maturity of the use case. A score of 3 expects the use case to be available and of reasonable quality for a telco to use in the next year or two. Moving down to 1, which assumes it will take many years to create a usable product. Factors considered include whether the use case will be needed by multiple verticals – which is assumed to speed up deployment.

2. **Usefulness to telco.** The x-axis scores the likely impact on the telco. Use cases gained maximum scores if Gen AI (often used with other models) was likely to reduce costs or increase revenues once it matured significantly. Lower scores were given if there was no significant impact expected or the majority of the benefit would be provided by other technologies. For example, Gen AI will improve sentiment analysis, but it can be offered today without its use.

The bubble size represents the number of use cases in each of the seven categories in Figure 4.

Figure 4: Importance of Gen AI types to telcos
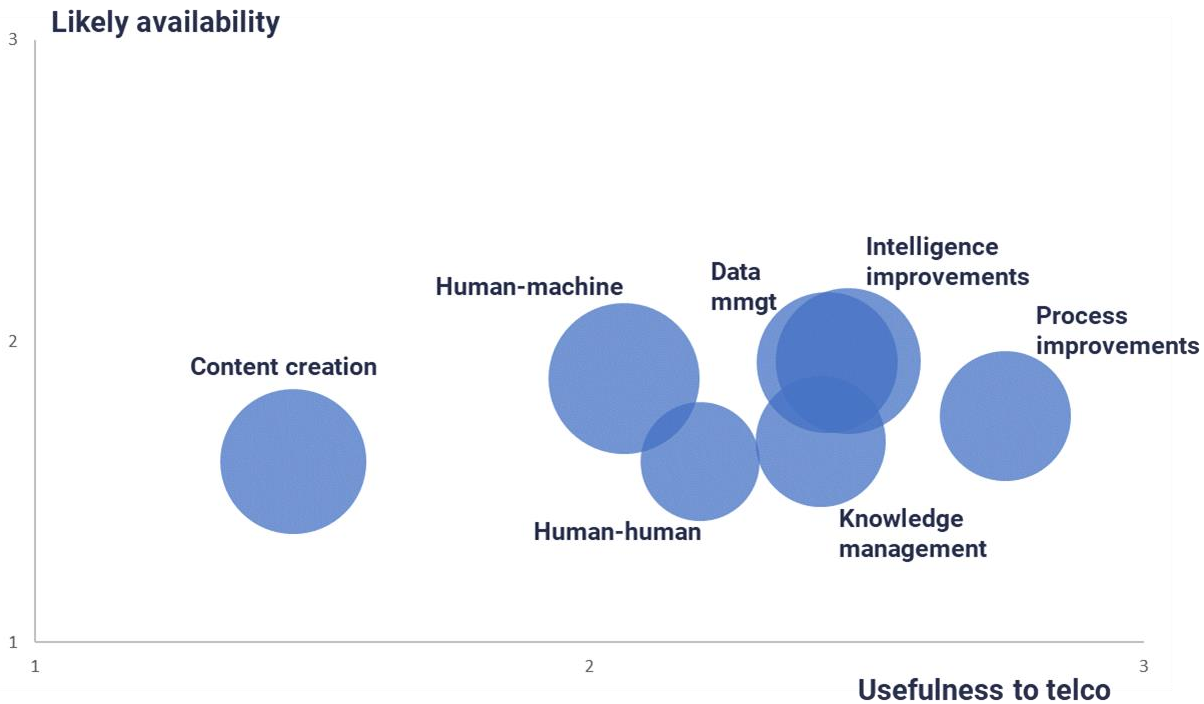
Source: Charlotte Patrick Consult

We can see that the size of the bubbles does not vary very significantly. The position of the bubbles shows greater differentiation, especially with regard the potential usefulness to telcos. Below we discuss the position of the different bubbles, starting with Process Improvements, the use case type found to have the biggest potential usefulness to telcos.

Process improvements

Code development will be increasingly important across the telco, as will some of the other process management activities included within this category (e.g., the ability of Gen AI to improve the documentation of processes as they are introduced or improved). Gen AI's benefits in testing and validation have also pushed this category to the right. However, the maturity of most use cases is not considered high – with some needing to be explicitly created for telco use.

Intelligence improvements

On the x-axis, this category benefits from the many possible uses across the telco for anomaly detection and improvements in prediction and personalisation. It is thought that most items in the category will be delivered reasonably rapidly contributing to a higher value on the y-axis, although some are yet to be developed and proven effective outside of the lab.

Data management

Use cases pushing this category higher on the x-axis include the ability for Gen AI to generate synthetic data and augment existing data sets. Data augmentation has some intriguing use cases, such as filling missing data points or helping in transfer learning (where knowledge from one domain can be used

to improve data in another). However, it is assumed that many of these theoretical uses are still mainly in the lab.

Knowledge management

Although only used in certain parts of the telco, knowledge management is essential for telcos; it underpins their ability to automate, the quality of their customer services and the user experience of new self-service capabilities in digital channels. Gen AI offers a range of capabilities that look useful to a telco (e.g. data summation, validation of accuracy, and the ability to fill in gaps in the information). However, it is hard to assess quite how long these will take to develop into tools that will cope with telco-specific data, especially as complexity grows with the addition of 5G and IoT services. These unknowns have pushed this category down the y-axis.

Human-machine

Capabilities in this category, such as improving chatbots and IVR with better NLU and NLG, provide the opportunity to improve customer experience significantly. The bubble is, however, not far along the x-axis because of a range of additional capabilities that Gen AI can also help with that seem less useful (the ability to create personality, improve conversational flow and vary response). However, these solutions' maturity is considered higher than in other bubbles.

Human-human

This bubble is slightly further along the x-axis (compared to the Human-Machine category) due to the generally applicable set of capabilities around building more useable digital assistants and a group of other Gen AI capabilities such as conversation transcription and Gen AI's use for making improvements possible in staff training (e.g., post-interaction analysis). It is not known precisely how mature these tools are, but it is assumed that they will require considerable work before significantly benefiting telco staff.

Content creation

Given the hype around ChatGPT, it might be assumed that the bubble would be higher up the y-axis than others; however, as discussed previously, even text creation is still in its infancy as a helpful tool and will require telco-specific training and improvement in the models before it can be customer-facing without human-decision. The bubble is also pulled down the y-axis by the immaturity of other content types, even though image generation is seeing significant improvements in model capabilities. Its place to the left on the x-axis is due to the varied usefulness to a telco of these content creation tools – with solid use cases in marketing and other content consumption teams but less use for image, video, audio and 3D content creation – and associated use cases such as rapid prototyping.

Based on this analysis, telcos wanting to prioritise Gen AI focus areas may find the most reward by concentrating on use cases that fall under process improvements, intelligence improvements and knowledge management.

Impact of Gen AI in the telco organisation

This section considers where each of the categories of use cases in Figure 3 could provide value to telco teams. Figure 5 shows the seven categories in the columns and various telco teams and activities in the rows.

Figure 5: Benefits of Gen AI to telcos

Telco team	1. Content creation	2. Human-machine	3. Human-human	4. Knowledge management	5. Process improvements	6. Data management	7. Intelligence improvements
Ecosystem management order handling • Product lifecycle management • Revenue management • Partner management	Creation of FAQs, summarisation of information and other helpful content around billing and product insight	Improvement in chat bot functionality for customer self-service	Improvement in digital assistant functionality to aid agents dealing with customers	Improvements to underlying product catalogues	Code snippets Automated API creation Generation of automation documentation	Data quality improvements in the product catalogue	Anomaly detection and improvements in prediction algorithms
Service management • Fulfilment/service design • Service orchestration • Assurance • Trouble ticketing			Improvement in digital assistant functionality Improved training				
Resource management • Network planning • Resource management • Incident management • Network deployment • Security and privacy				Improvements to network inventory system		Data quality improvement for knowledgebases	
Contact centre	Creation of any customer-facing content sent from contact centre	Improvement to IVR	Improvement in digital assistant functionality to aid agents dealing with customers Improved training	Improvement agent knowledgebases			

Telco team	1. Content creation	2. Human-machine	3. Human-human	4. Knowledge management	5. Process improvements	6. Data management	7. Intelligence improvements
Digital channels	Creation of more personalised content	Improvement in chatbots		Improvement to customer-facing knowledgebases	Code snippets Automated API creation Generation of automation documentation	Data quality improvement for knowledgebases	Improvements to personalisations
Sales	Creation of personalised content and novel copy/images for sales and marketing documentation		Improvements in chatbot and digital assistant functionality used in sales process	Improvements in knowledgebases, product catalogues, etc.		Data quality improvement for sales and marketing data	
Marketing							Improvement in recommender and campaign management systems

Source: Charlotte Patrick Consult

Conclusion

Use case prioritisation

STL has identified seven groupings of Gen AI use cases that can be beneficial for telcos, which are thoroughly described and analysed in the body of this report. These groupings are:

- Process improvements

- Intelligence improvements

- Data management

- Knowledge management

- Human-human interaction support

- Human-machine interaction support

- Content creation

All of these use case types have the potential to ring benefits to telcos, but based on our analysis prioritising use cases in the category of process improvements, intelligence improvements, and data management should lead to the best benefit in the short to medium term, based on both their potential usefulness to a telco, and the maturity of the specific Gen AI technology supporting these use cases.

How to proceed

As the hype surrounding Generative AI begins to subside, questions of its fitness for purpose and a range of other hurdles appear. Telcos should focus on working around these barriers in the short term.

- **Overestimation of capabilities**. Gen AI models have been described as able to "hallucinate confidently" – generating false or irrelevant information. For example, Meta's Galactica, a model trained on 48 million science articles claiming to summarise academic papers, solve mathematical problems, and write scientific code was taken down after less than three days of being online as the scientific community found it was producing incorrect results after misconstruing facts. This is perhaps an exaggerated example, but telcos must be realistic about Gen AI's capabilities, and test them in relevant contexts as potential implementations move forwards.

- **Unexpected outcomes**. It can be hard to control the model's behaviour and if performance is unstable, it can generate unexpected outcomes. In addition, they also lack the necessary "topic-centric" grounding and can quickly go off-topic.

- **Action drift**. The inclusion of Gen AI within large, complex workflows to trigger actions in automation will create a new level of risk, defined as "action drift". Here the underlying

relationship between the human request and the actions triggered by the system can change over time (e.g., moving from making a recommendation to taking an unwanted action).

- **General AI governance Concerns**. Reports suggest that vendors of current Gen AI solutions are typically operating under a self-governing model through limited-release strategies, monitored use of models, and controlled access to APIs. However, some companies such as Stability AI, believe that these models should be openly released to democratize access and create the greatest possible impact on society and the economy. This looks set to be a general issue across industries increasing the risk of bad actors using Gen AI for their own purposes and may require telcos to deploy defensive algorithms into the future.

It is important to remember that the costs of implementing Gem AI can be significant. Cost of acquiring the technology might range from $0.5mn to $200mn, and implementation is not necessarily straight forward. Telcos will need to be sure they have sufficient skills and resource available to support new Gen AI implementations. Return on investment will be hampered if the model output quality is low and integration costs are high.

Telco decision-makers will need to spend a significant and ongoing amount of time in-market understanding how to enable proper risk assessments as new models are rapidly delivered to provide brand new capabilities and improve existing ones.

Core recommendations

- Developing Gen AI knowledge across the whole company will have wide ranging benefits – if everyone is clear on what it can a cannot accomplish, and encouraged to think about or trial potential tools, progress should be quicker.

- Telcos will deploy Gen AI using one of five approaches, each with different parties involved in purchasing and managing solutions. This range of deployment scenarios is not unusual in a telco but given the nascent nature of Gen AI as a technology, telco IT teams need to be ready:

  – Specialist applications, particularly in areas such as content creation, will likely not need significant integration and be purchased as SaaS. However, centralised oversight will be needed to manage risks such as data privacy.

  – Gen AI providing minor augmentations to existing models on vendor platforms (e.g., generation of content within a campaign management tool) will come as upgrades, and vendors should be asked about short-term roadmaps to ensure benefits can be taken in the short term.

  – Models will improve the data and processes which support other models and automations (e.g., generation of synthetic training data). These may come from specialist vendors in, say, data management but have ROI impact across different parts of the organisation – making them a focus for the Chief Data Officer.

– A set of core models will be considered necessary to own – requiring a new platform team to provision and manage approved models. This team will need a range of Gen AI skills as well as broader skill sets such as integration and risk management.

– A small number of Gen AI projects will be complex but potentially high-value and will require cross-team focus, potentially for relatively long periods as the models grow in sophistication (e.g., chatbot and IVR customer experience).

3. **Prepare for a range of deployment scenarios.** Telcos will deploy Gen AI using one of five approaches, each with different parties involved in purchasing and managing solutions. This range of deployment scenarios is not unusual in a telco but given the nascent nature of Gen AI as a technology, telco IT teams need to be ready:

   f) Specialist applications, particularly in areas such as content creation, will likely not need significant integration and be purchased as SaaS. However, centralised oversight will be needed to manage risks such as data privacy.

   g) Gen AI providing minor augmentations to existing models on vendor platforms (e.g., generation of content within a campaign management tool) will come as upgrades, and vendors should be asked about short-term roadmaps to ensure benefits can be taken in the short term.

   h) Models will improve the data and processes which support other models and automations (e.g., generation of synthetic training data). These may come from specialist vendors in, say, data management but have ROI impact across different parts of the organisation – making them a focus for the Chief Data Officer.

   i) A set of core models will be considered necessary to own – requiring a new platform team to provision and manage approved models. This team will need a range of Gen AI skills as well as broader skill sets such as integration and risk management.

   j) A small number of Gen AI projects will be complex but potentially high-value and will require cross-team focus, potentially for relatively long periods as the models grow in sophistication (e.g., chatbot and IVR customer experience).

4. **Calculate the exact needs for human intelligence in new Gen AI deployments.** OpenAI hired 6,000 individuals to label an appropriate subset of data for its ChatGPT application. Obviously, telcos are unlikely to want to create their own foundational models, but the exact requirements for human intelligence in new Gen AI deployments should be carefully considered.

Index

STL PARTNERS

Research

Consulting

Events