The AI Paradigm:
A Special Report

Written and Designed by Humans

w: lab45 | a *lab45* Publication

2022 will be remembered as an important year. We saw a form of intelligence emerge that begins to show early signs of similarity to us. In the next decade, this form of intelligence will improve and become integral to our way of life. Fueling the growth and adoption of this intelligence are both, VC and Corporate Investments. VC investments have increased 10x in the past several years and major corporations are investing in the billions of dollars.

At lab45, we bring you this special report that will answer many of the key questions that you may have about Artificial Intelligence (AI). We discuss the science behind the technology that has led to generative AI and explain key concepts such as foundation models. We break down the building blocks that will help you understand the technology. We proceed to map the emerging business landscape to the building blocks so that you can develop strategies.

We then embark on a deep dive into the enterprise benefits and uses of this technology. We outline the different gains that businesses may expect in the coming years and discuss ways for business and IT leaders to implement AI in their enterprises.

We also carve out a section for government leaders, regarding the ethics and regulations in this field. We take a short-, medium- and long-term approach towards regulations. For citizens, we outline a range of new careers that will emerge.

Throughout our report, we bring to life many of our ideas through case studies. Case studies are examples of products, applications or companies in this field.

While there is a lot of speculation about the future of AI, we conclude with a positive note. As a human race, we have the ability to shape our collective future through our actions and our intent.

Rather than race ahead blindly towards monetary gains, we encourage a more human-centric approach towards business & innovation. We believe that we can't go wrong if we begin each AI project with ethics at the forefront— with the following question, top of mind— "How can my initiative leverage AI to bring a positive outcome for me, my society and my planet?".

We hope that this report informs you, inspires you and prepares you for the coming decade.

Contents

Section **01**

Breakthroughs
Leading to GenAI

AI models such as GPT-4 will impact enterprises, governments and consumers alike. They exhibit impressive abilities like acing the AP Bio Exam. At the same time, they can also exhibit inconsistencies in their logic. In this section we unpack the technology to make sense of this.

AI models will be in everyone's future, be it enterprises, governments or consumers. An advanced AI model (such as GPT3.5) is smart enough to ace the AP Bio Exam with a score of 59/60. Paradoxically, it can also fail to understand that a safety pin cannot be used as a dog leash.

Understanding the technology of AI models will help you understand this paradox and ensure that your company is applying it in the appropriate manner. GenAI and other AI models used today emerged from a long process of innovation and evolution. Understanding the evolutionary path of the different AI models that led to the development of GenAI will help any user or organization to best understand the technology and, most importantly, best apply the technology.

Key Insights

→ AI research in the last three decades has been founded on the concept of prediction engines called Models.

→ Foundation AI Models are AI Models trained on broad data sets that can then be applied to many different situations.

→ GenAI represents a paradigm shift from information retrieval to information composition, using models trained on vast datasets. These models create outputs such as text, audio, or video, tailored to address user prompts based on their semantic understanding.[1]

01

Deep learning has progressed in the last decade from only being able to classify information to being able to generate information.

Prediction-Based Modeling: The Science Behind the Models

AI research over the last three decades has been founded in the concept of prediction engines called machine learning. The goal of a prediction engine is to predict the right outcome when given an input. The engine is trained on a large dataset using a training algorithm. Another name for the prediction engine is "model."

Model development is an iterative process. In each iteration, the prediction engine aims at minimizing the error between the model outcome and the predicted outcome. After each iteration, the model is updated.

After enough iterations, the learning process updates, and the model "LEARNS" to predict the correct outcome. The model is then used for specific applications. One common example is sentence correction on a mobile phone. The mobile machine learning model, learning from its training on the English language, predicts the best outcome for a given sentence, when it detects an error.

Classical machine learning aspects, such as decision trees, often required explicit engineering to ensure the models would predict a particular outcome. Classical machine learning eventually evolved into deep learning.

The past two decades, deep learning has been considered one of the most advanced prediction engine models. Deep learning, based on a technology called neural networks, reduces the need for explicit feature engineering. Neural networks are a method of data processing modeled after the human brain. They allow the features to be learned implicitly from data examples, similar to the brain.

Advances in Deep Learning

Deep learning became more powerful than conventional models due to massive amounts of data collected and stored by various applications. Many virtual assistants and chatbots today are based on deep learning algorithms. Information is fed in the form of text, images, data, audio and video.
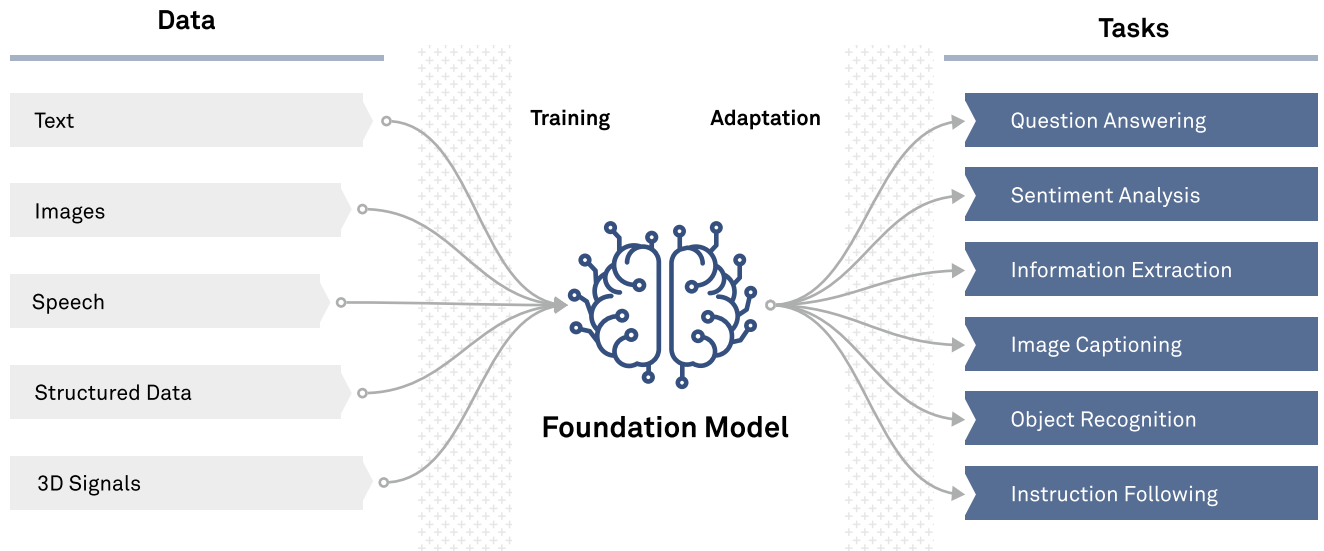
Deep learning progressed in the last decade from only being able to classify information to being able to generate information.

Steps in this evolution include:

1. **The availability of massive computational power (GPUs)**

2. **The availability of massive datasets.**

3. **Two breakthrough models that were developed in deep learning, named GAN in 2014 and Transformer in 2017**

The GAN Model

Introduced by Ian Goodfellow in 2014, the Generative Adversarial Network (GAN) model consists of two neural networks, a generator and a discriminator, which engage in an adversarial training process. GANs have become popular for applications like image synthesis and artistic style transfer, thanks to significant developments such as Wasserstein GANs and StyleGANs. The continued evolution of GANs will likely expand their applications and further advance AI and deep learning.

Data

Text

Images

Speech

Structured Data

3D Signals

Training

Adaptation

Foundation Model

Tasks

Question Answering

Sentiment Analysis

Information Extraction

Image Captioning

Object Recognition

Instruction Following

Source Credit : Stanford Research Institute for AI studies

The Transformer Model

The Transformer AI model, introduced in 2017, revolutionized natural language processing with self-attention mechanisms, enabling efficient parallelization and scaling to larger datasets. Serving as the foundation for state-of-the-art NLP models like BERT, GPT, and T5, it has been adapted for tasks such as machine translation and text generation. The transformer architecture's ongoing development and widespread success continue to shape AI research and language-based applications.

The GAN and Transformer models have been used by innovative startups such as OpenAI to train AI models on very large datasets. One of the key features that emerged during this training was in-context learning.

Foundation Models are pre-trained AI models that can learn from one situation and apply the learnings in another situation making them foundational to many tasks. Once trained, these models can be used for different applications. In the case of OpenAI, the foundation model was GPT4. GPT was trained on data that comprised of a large portion of the internet.

Foundation AI Models

Foundation AI Models will be in everyone's future. Foundation Models (seen in the graphic above) are AI models that are pre-trained on vast data sets making them valuable across many general purpose tasks. Companies spend significant investments training such models. Foundation Models, like BERT, GPT, and CLIP can then be fine-tuned towards specific domains and tasks, potentially shaping the future of enterprises. Any enterprise, government, or other organization can adapt these models and train them to meet their domain-specific needs. When adapting these models to specific domains, organizations can save time and resources. However, there's a risk that their limitations and biases will become ingrained in society, and our understanding of their inner workings is still in its infancy.

CASE STUDY:

NLLB-200

Meta has released an open-source AI model capable of translating 202 different languages.

The model is called NLLB-200 and is named after the company's No Language Left Behind initiative. Meta says it will improve the quality of translations across its technologies by an average of 44% with that number jumping to 70% for some African and Indian languages, as shown by its Bilingual Evaluation Understudy (BLEU) benchmark scores.

The GenAI Model

GenAI models are a class of deep learning architectures that create new data samples based on patterns learned from existing data. They encompass various techniques, including those of Generative Adversarial Networks (GANs) and Transformer models. These models have enabled applications like image synthesis, artistic style transfer, text generation, and machine translation.

> As GenAI continues to evolve, it has the potential to revolutionize numerous domains by generating realistic, contextually relevant content.

CASE STUDY:

Software Engineering Applications

Software engineering has emerged as one of the most practical and promising use cases for GenAI. With models like Codex by OpenAI or Ghostwriter by Replit, engineers can breeze through the mundane elements of coding, while focusing on more significant work.

While all code should be reviewed and curated, studies have shown that engineers can generate 30% of their code with AI, and that number is likely to grow as models become more accurate. We could even see in the near future, GenAI being used to automate engineering beyond software, including chip design and product design.

CASE STUDY:

Healthcare Applications

GenAI is showing promise to augment doctors with productive tools. It is still too risky to use GenAI to diagnose patients or for patients to go directly to these tools for medical care.

Yet GenAI could be used to help doctors write notes after patient visits, or to create realistic copies of patient data with anonymized private identifying information. For many hospitals, GenAI is being seen as a potentially powerful tool to lessen the documentation burden for their medical professionals, and models that specialize in these tasks could be very impactful for medical professionals across the world.

"The development of AI is as fundamental as the creation of the microprocessor, the personal computer, the Internet, and the mobile phone."

— Bill Gates

Section **02**

Building blocks of
GenAI

In this section, we cover the building blocks that an organization would have to consider to build Generative AI. Depending on the organization's AI maturity and business goals, they can choose to build all these blocks or work with other organizations that have platforms supporting these blocks. We encourage you to think of these as lego blocks in your build process.

We define GenAI as consisting of several layers and identify the blocks in each layer with case studies to illustrate them. We address key considerations like output interpretability, explainability, and user experience to establish trustworthy systems. These layers and building blocks are crucial for training, building, and using effective and reliable GenAI models.

Key Insights

Technology Architecture of GenAI

→ Two factors heavily influence your AI model : Data and Model training techniques. The data quality and model training forms the foundation of all machine intelligence. Consequently, robust datasets and effective training techniques are essential for achieving high-quality AI-generated results.

→ AI technology consists of a number of layers, each of which is optimized by particular companies or organizations. Startups, companies, and others contribute to each layer, where they are best suited.

02

lab45

The Technology Building Blocks of GenAI

A.

Consumption Engine	Application Layer	Text to Image, Text to Video, Image to Audio, Text to Audio. Text to 3D ChatGPT | DALL E | Riffusion | Vall E		
		API Engine	Intelligent Personal Assistant (IPA) - Hate/ Bias/ Personality/ Empathy/ Persuaion	
	Presentation Layer	Composition	Rendering	Veracity of Composition

B.

Composition Engine	Manipulator	User Intent	User Context	User Profile
	Orchestrator	Dynamic Workflows Mode(s) orchestration	Knowledge Representation Content sharing across models	Semantic Understanding Mode context / Content understanding

C.

Construction Engine	AI Model	Hyper Local AI Models (Profiling & Personalization)	Specific AI Models (Contextual Awareness & Domain Understanding)
		General AI Model (General Awareness) | Reasoning | Understanding | Generation) GANs | VAE | Transformers (LLM - GPT3 | BERT) | Diffusion Models (DALL E2, VALL E)	
	Architectures & Models	CNN, VAE, GAN, LLM, Transformers, CLIP	
	Techniques & Components	Encoder-Decoder, Attention Mechanisms, Diffusion, GAN	
	Foundational	Generative, Discriminative	

D.

Data Engine	Representation Layer	Knowledge Graph	Feature Set	Textual Embeddings
	Ingestion Layer (Reference Data)	Pre-Existing Data (Audio | Video)	Sensor Data	Text Data (Half Written Code)
Training	Training	Supervised	Reward (Annotation)	Reinforcement

Ethical Design

Explainability

Causality

Trust

▲

In the above graphic, we show the layers A, B, C and D, and the building blocks of GenAI in each of these layers. Think of these as lego blocks that you can build to meet your enterprise needs.

A.

Consumption Engine

The consumption engine empowers any business (small or large) to build intelligent applications that offer personalized outputs in a user-friendly format for its customers. With the ability to generate text, images, code, videos, audio, and more from simple natural language prompts, the consumption engine represents a significant opportunity for developers to create innovative applications that transform and change the way people interact with information and brands.


CASE STUDY:

ChatGPT

The ChatGPT app has single-handedly propelled AI into the public consciousness.

It's a versatile chatbot that makes GenAI models useful for the masses with its high-quality text construction. It's being used to summarize information and as a creative ideation tool that can mimic creative styles. It has launched as a free product for hundreds of millions of users. As an app, it plugs into multiple other APIs to present usually relevant information within a consumer-friendly user interface.

B.

Composition Engine

The composition engine, with its manipulator layer, allows GenAI models to understand user intent and context, delivering personalized responses. Beyond a GenAI application, the composition engine can manage multiple GenAI applications at once, by virtue of the user's original intent, without the need for multiple prompts.

GenAI offers creators and artists efficient tools for diverse content production.

Innovators and platform providers can develop custom orchestration tools to manage GenAI workflows, enabling seamless connections between text prompts, images, videos, and voice, and fostering a new era of creativity.

CASE STUDY:

AutoGPT

AutoGPT has the ability to manage multiple AI workflows at once. For example, it can plan a weekend using a search AI, a chatbot AI, and a payment AI, without manual intervention.

AutoGPT and other composition engines can create, manage, complete, and reprioritize tasks to achieve a given user intent.

Composition engines can not only manage other AIs to complete tasks— they can also be used on a more meta level, to build and deploy other Ais. See the below case study for an example.

C.

Construction Engine

The construction engine is where the real magic of GenAI happens. It is where models are developed and refined using advanced architectures like transformers. Transformers utilize encoder-decoder and attention mechanisms. Technology providers, open-source enthusiasts, and entrepreneurs can all build the necessary foundation and fine-tuned models. These models can range from the generic, like GPT-4 and BERT, to fine- tuned models that address specific industries and use cases.

Additionally, hyper-personalized models can be created to offer contextually appropriate responses.

With the help of the construction engine, businesses can leverage the power of GenAI to solve complex problems and interact with the technology in a more effective way.

CASE STUDY:

DiffDock

MIT researchers have developed DiffDock, a groundbreaking model that has the potential to revolutionize drug development.

DiffDock is capable of identifying protein components that were previously unknown and generating original 3D coordinates that can aid protein manipulation. This innovation is just one way that GenAI can be used to accelerate biological research. In the future, GenAI could become very useful across all medicine in generating and evaluating possible drug candidates, by pinpointing the most promising options and concentrating their efforts on them.

D.

Data Engine & Model Training

Quality data is crucial for developing the mind and internal universe of a GenAI model. The data engine layer is vital for processing that data and using the data to create advanced GenAI models that can tackle complex problems and transform industries.

Enterprises must focus on quality data and data training techniques. Technology providers can offer support through data hygiene services, knowledge graph synthesis, feature extraction, and textual embeddings.

Additionally, providers can contribute manpower skilled in supervised learning, reward-based learning, and reinforcement learning with human feedback for effective model training.

CASE STUDY:

DeepMind & DeepNash

DeepMind has long been training deep learning models to compete against humans, to showcase the depth of complex decision-making and optimization that can be achieved with AI.

The goal in training models is to develop AI that can solve real-world problems with unpredictable scenarios, such as in autonomous driving or traffic optimization. In 2022, DeepMind shared research of a new model, called DeepNash, which had an 84% win rate against human players in the game of Stratego. DeepNash shows the ability of neural models to train data and infer decisions, with an eye for long-term goals.

Section **03**

How is the Industry **Organizing Itself to Deliver Value Around GenAI?**

We see a rapid rise in the investments in GenAI leading to an explosion of players in this space. In this section, we map the players to the building blocks of GenAI. Since we are unable to list all the players, we have chosen a few players as examples in each block.

In this section, we illustrate how companies are organizing themselves in the landscape to deliver value using GenAI. This is an early assessment and we expect this landscape to change significantly in the next 6 months to a year as new value is discovered and the market hype stabilizes.

Key Insights

→ Foundation models have been at the forefront of the excitement of Gen-AI. We should expect to see a rise in narrower, fine-tuned models, trained on a foundation model such as GPT 4, to provide generative content for industry or function specific use cases.

→ Data is the most valuable resource in the race to develop AI apps and models. Enterprises should seek to understand how they can develop their proprietary data to be valuable and with minimal bias.

03

The Emerging Enterprise AI Model Landscape

When we map the theoretical construct of the previous section to the emerging business landscape in 2023, we find that the market has organized itself into three areas:

1. **Applications (consumption engines)**

2. **AI models (construction engine)**

3. **Tools for model development (model training and data engine)**

Some companies such as Microsoft play in all three areas forming both verticals and horizontals. Most are preoccupied with one of these areas alone.

Enterprises that are at the forefront of innovation, may be able to unlock immense value by developing the tools of model development. These include infrastructure for data hosting and provisioning, software and hardware for AI processing, development engines for customized model training, and tooling for enterprises to improve their data ETL processes.

Most organizations do not have the resources to build tools for solutions in model development or data operations. These organizations are deriving value by building GenAI at the top of the stack. They are

1. Building their own fine-tuned generative models, trained on their proprietary data.

2. Building apps on top of models from other companies.

3. Integrating generative capabilities within their apps.

Whichever be the approach your organization takes we recommend balancing speed with caution.

Companies that participate in building foundation models are generally very large and are also cloud service providers such as Amazon, Microsoft, Google, etc.

Companies/Products in the End Applications Space

Consumption Engine	Lifesciences Biomatter, MDClone	Images Mid Journey, Stable Diffusion	Text Jasper, Grammarly, Dramatron	Search Vectara, Perplexity.ai, mem, you.com
	Chatbots ChatGPT, Google Bard, Character.ai	Audio ChatGPT, Google Bard, Character.ai	Speech Assembly.ai	

Companies/Products in the Model Space

Construction Engine (Models)	Non-Language Foundational Models DALL-E-2 (images), Stable Diffusion, OpenAI's Whisper	Foundational Language Models Whisper, Riffusion, VALL E	Fined-tuned Models Nvidia NeMo, Aleph Alpha, Github CoPIlot, Adobe Firefly, runwayML, GPT-J

Companies/Products in the Tools Space

Model Training (Machine Leaning Ops)	Model Training & Deployment Vertex.ai, Datature, Kubernetes, Flyte	Model Ops Amazon SageMaker, Comet, Mlflow, Weights & Biases	Machine Learning Processing Cerebras CS-2, Google TPU v4, Nvidia A100, Intel Xeon
Data Engine (Data Ops)	Data Pipeline Mgmt Apache Airflow, Apache Spark, Prefect, Dagster, Flyte, scale.ai	Data Processing Snorkel, Encord, squared.ai, ango.ai, Feast, timbr.ai	Data Hosting Amazon S3, Google Cloud, Azure, Parquet, Snowflake, Databricks, BigQuery, SQL

Company names are added purely for illustrative purposes and are examples.

▲

In the graphic above, we show the emerging landscape of applications, foundational models, and fine-tuned models, We also show an overview of the data and ML operations landscape, to help the reader get an understanding of the type of products and companies in the space. In our writeup, we provide tactical recommendations for how enterprises can build solutions in this space.

Applications (Consumption Engine)

Enterprises are racing to capture the value that GenAI consumption engines will create.

Much of the explosion of interest in GenAI is in the market for consumer apps. Due to the rapid progression of neural networks, there is a continuous stream of new generative capabilities in text, images, video, code, and audio. Text and images are the two areas that are showing the most popular near-term use cases.

GenAI is already disrupting application segments in chatbots, search, and software programming. Image generation and text generation are emerging as new categories of consumer apps.

CASE STUDY:

Chatbots

ChatGPT Plus is attracting customers from individuals and enterprises who want access to the latest model developments.

Unlike traditional chatbots, generative chatbots have semantic capabilities and can generate an output that sounds human-like, even with minimal factual accuracy. Improving accuracy is the ultimate goal. The rise of generative search has led to a range of tailored chatbots relevant to specific industries, use cases, or styles of generative output, such as Character.AI's chatbots that mimic Albert Einstein's style.

CASE STUDY:

Content Creation

GenAI is transforming content creation by enabling marketers to generate image and text content quickly and at scale.

Jasper.ai and Grammarly are examples of companies offering such services. Gen-AI is also being used to generate software code, with tools such as Github Co-Pilot and Tabnine increasing programmer efficiency and enabling them to focus on higher-level work.

CASE STUDY:

Search

Generative search is disrupting the search engine industry, utilizing language model neural networks to compose outputs that match user inputs.

This provides a seamless and efficient experience for individuals seeking quick access to information. However, these neural networks are few-shot learners, meaning they may not be accurate in generating factual outputs, particularly when training data is limited or not relevant to the query. Despite this limitation, generative search and chatbots remain popular and offer directional guidance for users.

Models (Construction Engine)

Foundation Models

Beyond the application layer, there is an equally-important market in easing and scaling AI development for other builders. The complexity of the model and the problem, along with data, talent, and computing costs, can cost upwards of millions of dollars, depending on the complexity of the model and the problem to solve.

While the cost to develop models should significantly decrease, hosting and running inference on these models could still cost many thousands of dollars a week.

The construction engine market can be broken into two categories: foundational model builders and fine-tuned model builders.

CASE STUDY:

BERT by Google

BERT, developed by Google, is a groundbreaking foundation language model that revolutionized natural language processing tasks.

Trained on vast, generalized datasets, BERT utilizes a bidirectional approach to understand context in language, allowing for improved performance in various text-based tasks. As a highly influential model in the AI landscape, BERT has paved the way for numerous applications and advancements in language understanding and generation.

Fine-Tuned or Narrow Models

Foundation language models are at the forefront of much of the hype surrounding GenAI.

Organizations that do not have the resources to build foundation models may instead develop smaller, fine-tuned models.

There are two approaches to fine-tuned models:

1. Developers can access open-source and shared datasets and generative models, upon which they can fine- tune to meet their organizational needs.

2. Many foundational models have open APIs, allowing anyone to build fine-tuned models for specific industries or AI tasks.

We expect to see a rise in the development of narrower models, for industry or enterprise specific-use cases. We are starting to see such models emerge, in industries as varied as finance and fashion. Enterprises can charge users for access to these fine-tuned models, or sell access to its API, for other organizations to fine-tune the models even further, for even more granular use cases.

For example: A retail organization can build a fine-tuned generative model for developing sneaker product mockups, built on Adobe Firefly, and train the model on their history of sneaker designs. If this model becomes the most effective tool to rapidly build sneaker designs, they can lease the model to other sneaker companies, and broaden the model to further clothing categories.

Enterprises can even offer their model to potential customers or organizations for free, leading the organization to have a software business unit for other retail enterprises, while much of the heavy lifting of developing a text-to-image neural network was developed by a foundation model provider.

CASE STUDY:

BloombergGPT

BloombergGPT is being developed by Bloomberg, which blends financial data with general-purpose datasets and is trained on their current datasets and documents

This method is expected to produce higher accuracy and performance in generating finance-related text, as well as help develop an understanding of financial products or companies.

Section 04

How Can Enterprises **Harness the Power of GenAI?**

Enterprises need a well-defined AI strategy to stay competitive in today's business landscape.

AI has the potential to revolutionize industries and functions. While the C-suite within an enterprise will always have a competing list of innovation priorities, AI should be on the top of the list. Failure to adopt a strategic approach to AI will result in missed opportunities, reduced market share, and the risk of falling behind in innovation.

In this section we outline both a framework to capture gains from Generative AI and also a framework to integrate Generative AI in the enterprise.

Key Insights

→ Enterprise gains will move from productivity gains to revenue gains to innovation gains.

→ Due to the potential for inherent bias and errors in existing tools, adopting the Copilot model for AI is recommended for enterprises in the short term.

04

In the short term, enterprises should evaluate the productivity gains through the augmentation of Generative AI in enterprise functions. In the medium term, enterprises could consider building products and services using inhouse fine-tuned models or external foundational models and in the long-term enterprises need to consider the innovation gains that can be achieved with GenAI. We outline all three scenarios and possibilities in each scenario.

Productivity Gains

Short Term

Revenue Gains

Mid Term

Innovation Gains

Long Term

In the graphic above, we show the roadmap of gains an enterprise may witness with GenerativeAI.

Productivity Gains (Near term)

GenAI will play a significant role as an assistant in enterprise functions , assisting humans with functional tasks such as content creation, content organization and content presentation. However, it is unlikely that GenAI will create perfect artifacts. Humans will still need to curate and execute a final product.

As GenAI streamlines and accelerates the creative process, it can help organizations respond more quickly to changing market trends and consumer demands, enabling them to stay ahead of the competition and achieve unprecedented levels of success. With GenAI tools, designers can generate multiple product mock-ups quickly and efficiently, giving them more freedom to experiment with different design options. AI-powered content creation tools can assist in crafting personalized and localized content that resonates with individual audiences.

Enterprises must be cautious when deploying AI in their creative organizations. They must establish guardrails for common GenAI errors and design human-machine workflows that optimize and empower their creative workforce.

GenAI can assist enterprise professionals in the following functions—

Design & Product Development

GenAI has the potential to transform product design and development by reducing human effort required for design exploration and production. This technology can be used in creative departments of any enterprise and could become a must-have tool for generating visual mock-ups of product ideas and rendering 3D prototypes of 2D objects. As Gen-AI apps become more accessible, they are likely to alter all design processes.

It can generate visual mock-ups of product ideas, rendering 3D prototypes of 2D objects, and altering the design in real-time based on customer feedback. GenAI can also create hyper-targeted and personalized product designs, tuned to individual customer preferences, resulting in higher customer satisfaction and engagement.

The human touch will still be fundamental in product development, with human designers executing the final product design, while GenAI assists in the exploration and ideation process.

CASE STUDY:

AiDLab

In 2022, a group of Hong Kong designers gathered to produce a GenAI fashion show.

They used tools from several tech companies, including Fashable and Cala, and tapped into the power of GenAI for numerous design variations.

CASE STUDY:

Adidas

In 2017, Adidas partnered with Carbon, a 3D printing company, to use GenAI algorithms and 3D printing to generate a new shoe and type of sole.

They used GenAI to analyze data on running mechanics, cushioning, and support requirements. The AI generated a range of potential designs for the shoe, optimizing for performance, comfort, and material efficiency.

Marketing

GenAI will enable hyper-targeted and personalized marketing with the rise of immediately generated bespoke content for specific consumers based on their preferences. The marketing industry is expected to see a proliferation of off-the-shelf Gen-AI tools. Many companies will integrate these tools into their products and services. Ad agencies, marketing firms, and other creative service agencies can utilize GenAI for creating ads, marketing displays, magazines, online articles, television ads, social media campaigns, etc.

Software Engineering

GenAI shows potential in creating computer code, with tools like Microsoft's Copilot already used by developers to accelerate code generation and optimization. This could even enable non-IT teams to create low-code apps, and AI-augmented testing tools can improve testing efforts. However, AI-augmented approaches need to be integrated into overall DevOps and MLOps strategies and cloud architecture. The human touch will still be necessary to understand full functionality and develop secure code.

CASE STUDY:

Amazon CodeWhisperer

CodeWhisperer is a free code assistant developed by Amazon.

Trained on billions of lines of code and powered by machine learning, CodeWhisperer continually examines the code of its users, and presents the user with corrections that are syntactically correct. Recommendations are based on coding style and variable names, personalized for the coder.

CASE STUDY:

Google Ads

Google is planning to incorporate GenAI into its ad business.

Advertisers will supply Google with creative content— images, text, etc— and Google AI will then mix this content and reformulate it for particular audiences and demographics— along with sales targets.

GenAI will enable hyper-targeted and personalized marketing with the rise of immediately generated bespoke content for specific consumers based on their preferences.

Recruiting

AI can be used for recruiting to automate and streamline various tasks, such as resume screening, candidate matching, and interview scheduling.

> **AI algorithms can analyze resumes and applications to identify the most qualified candidates, reducing the time and effort required by recruiters.**

Additionally, AI-powered chatbots can interact with candidates to answer their questions and provide guidance throughout the recruitment process.

However, it's also been shown that AI can exacerbate biases in the hiring process. AWS used an AI recruiting tool, and later had to discontinue use of the tool after it became evident that it was biased against women.

CASE STUDY:

Screenloop

Remote work and video conferencing have become embedded into every organization, and are especially relevant in talent recruitment, where over 60% of HR professionals using video.

Screenloop is a startup that seeks to develop platforms for hiring intelligence, and are using speech transcription for recruitment insights. Through transcription AI models, they can summarize key sections of interviews, and develop algorithms that score how well a user interviewed, reducing hiring bias. With transcription intelligence, Screenloop is able to develop interview intelligence, and training tools for their customers. Screenloop's customers now realize up to 90% less time on manual hiring and interview tasks.

Drug Development

AI can enhance drug development at every stage, from target discovery to clinical trials. By undercutting costs and reducing timelines for getting an Investigational New Drug (IND) application, AI has the potential to revolutionize drug development. In the clinical development phase, AI-guided drug development can also be used for testing efficacy and tolerable safety profiles.

Studies suggest timelines for IND applications could reduce by as much as 50%. Additionally, AI can generate synthetic data, which is especially useful in healthcare where sample datasets are small and can help to mask sensitive patient data, transforming data custody practices and regulations while lowering compliance costs.

CASE STUDY:

Insilico Medicine

Insilico Medicine used its AI platform, GENTRL, to design a new drug candidate in just 46 days, showcasing the potential of AI to accelerate drug development.

The platform employs deep learning and reinforcement learning to design novel molecular structures with desired properties, dramatically reducing the time and cost associated with traditional drug discovery methods. The achievement was published in the journal Nature Biotechnology, highlighting the significance and potential impact of AI in pharmaceutical research and development.

Revenue Gains through Products/ Services (short-med term)

In the short to medium term, we expect an explosion in adoption of products and services built on fine tuned and foundation models. We see enterprises capture value in two ways : enterprises that have access to data will build fine-tuned AI models that can be monetized as a service and enterprises that don't have access to data will build apps/products using such AI models.

Legal and ethical guard rails must be put in place before AI tools can be used without human supervision.

Regardless— GenAI will drive revenue improvements and we are already seeing evidence of this with less-advanced pre-GenAI models, as the case studies below highlight.

CASE STUDY:

Uber

Uber employed AI to optimize its dynamic pricing strategy, which in turn improved revenue.

Uber's AI algorithms factored in variables such as demand, location, and time to determine the best prices for rides. This dynamic pricing strategy contributed to Uber's significant revenue growth, with 2018 Q4 revenue reaching $3 billion, a 24% increase from the previous year.

CASE STUDY:

Starbucks

Starbucks used AI-driven predictive analytics to boost revenue by personalizing offers for its customers.

Their AI platform, Deep Brew, analyzed customer data to identify patterns and preferences, tailoring recommendations and promotions to individual tastes.

CASE STUDY:

H&M

H&M, a major fashion retailer, implemented AI to optimize inventory management and improve customer experiences, leading to increased revenue.

By analyzing sales data and customer preferences, AI algorithms helped H&M tailor their inventory to local trends and demands. This resulted in a 7% rise in Q3 2018 sales.

Innovation Gains (long-term)

In the long run, AI's pattern recognition capabilities will enhance human innovation in areas like drug discovery, material science, and product design. For example, Gen-AI will enable creative breakthroughs in medicine, with generative design accelerating the discovery of new drugs and materials. By 2025, over 30% of these discoveries may come from GenAI, though understanding molecule integrity and the risks remains crucial.

In the field of material science, GenAI will transform industries by suggesting materials with specific properties through inverse design. "Inverse design reverses the conventional design process, allowing new materials and compounds to be "reverse-engineered" simply by inputting a set of desired properties and characteristics and then using an optimization algorithm to generate a predicted solution", says MIT's research enterprise in Singapore. This approach benefits sectors like aerospace, defense, and chemical.

CASE STUDY:

AlphaFold

Google's DeepMind developed AlphaFold, an AI system that predicts protein structures with remarkable accuracy.

This groundbreaking innovation has the potential to revolutionize fields like drug discovery and biology, enabling researchers to better understand diseases and design novel treatments. AlphaFold's success was recognized in the 2020 Critical Assessment of protein Structure Prediction (CASP).

AI's potential also lies in optimizing part design across various industries, such as using generative design in automotive to create lighter parts for better fuel efficiency.

The practical integration of AI within enterprises (short to medium term)

GenAI Implementation Options

Every enterprise leader will be expected to have a roadmap for their adoption of GenAI or other types of AI. We use ChatGPT below and OpenAI's foundation model, GPTx, to illustrate several approaches to integrating GenAI over the short- and medium-term horizons.

Enterprises will need to upskill employees and transform their capabilities, as automated technologies continue to replace time-consuming tasks that aren't using creative brainpower.

The result could be incredibly creative capabilities democratized across the enterprise.

For example: an animated film company like Disney could use image generation tools to generate mockups across every element of its business, rather than only relying on its animation team. The Disneyland marketing team could instantly generate architectural mockups of a new ride without having to consult the experience design team.

GenAI is similar to photography. At the advent of camera phones, many worried that professional photographers may soon be out of a job. Mobile phones didn't make photographers obsolete but instead democratized the photo-taking process. There are more professional photographers now than 20 years ago, and over 1.7 trillion photos taken annually.

As a disruptive technology, AI requires proactive and holistic design for positive societal impact.

Implementation
(using OpenAI's models as an example)

Short-term		Medium-term
(Approach 1) Deploy ChatGPT Plus3 SaaS in selective functions within the enterprise	**(Approach 2)** Deploy ChatGPT API in some pilot programs to determine its ROI	**(Approach 3)** License GPT3.5, the base model of ChatGPT, and train it using your own domain specific intelligence.

Enterprise	Enterprise	Enterprise
	Enterprise Specific Interface	
	API	Enterprise Specific Interface
ChatGPT Plus Interface	ChatGPT Plus Interface	Domain Specific Training Layers
GPT3.5 (or newer) Open AI Models	GPT3.5 (or newer) Open AI Models	GPT3.5 (or newer) Clone Model

Pros	Pros	Pros
• Assess ROI and use cases for GenAI quickly. • Use information to build near to medium term strategy.	• Reduce risk of poor consumer experience and litigation from algorithm errors. • Have time to set up governing models and guard rails.	• Gain a competitive edge by being the first in the industry to combine base AI model with domain specific intelligence.

Cons	Cons	Cons
• Poor consumer experience from incorrect output. • Litigation risk from algorithm errors.	• Delayed organization- wide learning of GenAI.	• Resource and time-intensive program.

VS   VS

Section **05**

The New AI Economy:
New Careers, New Lifestyles

GenAI may represent a new frontier for innovation, progress, and human possibility.

AI will democratize human creative expression across music, art, literature, etc. With AI, anyone can be a creator by inputing a few prompts. Harnessing the power of algorithms and neural networks, human creativity will see a new paradigm.

As AI integrates into our everyday life, an entirely new set of careers will emerge from AI model trainers to AI ethicists. We outline some of these careers in this section.

Key Insights

→   Gen-AI allows anyone to create without physical limitations, stretching the boundaries of human imagination and creativity. Similar to how the internet reduced the distribution costs of information to zero, Gen-AI can push the costs of creation toward zero.

05

lab45

Unleash Your **Creativity**

Until now, the barriers to creative expression have been high— requiring hard-won skill in order for creative expression to be possible. Now, thanks to GenAI, this is no longer the case. This technology enables anyone to create content simply from a prompt or an idea, allowing unprecedented exploration of creativity and imagination. In the future, to be creative, all that will be required is the human imagination and the power of GenAI.

A Brave New World?

As GenAI advances over the coming years, both what we do for a living and how we do what we do are set to transform.

> **We need to take big bold steps into the careers of the future, with curiosity and confidence.**

Will the legal and medical professions undergo massive transformations? Will graphic designers be irrelevant? Will call center jobs become redundant?

The big question is whether GenAI will make working professionals more productive, redundant, or both. This has prompted notable public figures like Elon Musk, Yuval Noah Harari, and Andrew Yang to call for a "halt" on "giant AI experiments." It also raises a challenging inquiry: "Should we automate all jobs, including those that bring fulfillment?"

Fortunately, we expect new jobs, that may also be fulfilling, to emerge. These roles include prompt engineering, AI interaction designer, AI model trainer, and AI artist. We expect these to present exciting opportunities.

"Some people call this artificial intelligence, but the reality is this technology will enhance us. So instead of artificial intelligence, I think we'll augment our intelligence."

— **Ginni Rometty, CEO of IBM**

The Creation of New Jobs

Prompt Engineering: the new programming paradigm

Prompt engineering is an emerging field that has the potential to revolutionize various industries and applications, particularly in the realm of new-age digital creativity. It involves producing natural language prompts for GenAI models to generate a desired output. Currently the skill is considered fairly specialized, but as GenAI technology advances and becomes more widespread, prompt engineering will likely evolve into a more ubiquitous skill set— similar to typing and use of spreadsheets. The increasing use of AI in several sectors raises many questions about the future of work.

AI Model Trainers

The emergence of GenAI creates a pressing need for skilled AI trainers who design and train models with fairness, accountability, and ethical considerations. They will help ensure that the model is beneficial for the organization, the user, and the society. They identify and address potential biases and privacy concerns. They monitor and retrain models for regulatory compliance.

AI model trainers must have a deep understanding of the data used to train AI models, and ensure that the data is diverse and representative of the population. As AI technology continues to advance, the demand for skilled AI trainers is expected to grow exponentially

Will the legal and medical professions undergo massive transformations? Will graphic designers be irrelevant? Will call center jobs become redundant?

AI Artists

AI Artists use GenAI-powered tools to create innovative art. They will blend technology and the arts to push the limits of human creativity. They are pioneers in this emerging field, shaping the future of art and design by driving its evolution and pushing boundaries.

AI Ethicists

AI ethicists develop ethical guidelines for GenAI, ensuring transparency, accountability, and fairness. They will be the guardians of the society— ensuring that AI does not harm society at large. They promote inclusivity and diversity to avoid societal inequalities and protect individual rights.

AI Interaction Designers

 AI interaction designers will be focused on integrating GenAI tools into human workflows to create more relatable and engaging technology interfaces— they will pave the way for human creatives and GenAI to work together. They design applications for users that are convenient and effective— furthering democratizing AI for all.

Section **06**

GenAI— **Safety, Legal, and Ethical Implications**

Enterprises have shaped human lifestyle and will continue to do so. As enterprises race to build GenAI products and services, they need to balance speed with caution and ensure that they are also making investments towards AI safety.

For all the potential of AI-generated content, there are also a myriad of complex challenges and potential risks. If AI capabilities are exploited for malicious use or created without any thoughtful human-machine design or guardrails, it may lead to an erosion of trust in society and government at large. Ethical design and regulation needs to be at the forefront of GenAI.

Further, regulating technologies for social good is often more complex than developing the technologies themselves.

Key Insights

→ GenAI carries many embedded risks and potentials for incredible disruption

→ Ethical design must be a priority with respect to GenAI

06

We have the ability to harness this technology for incredible good.

Governments differ in their concerns and priorities. This makes regulation of technology, particularly digital technology, often fragmented, difficult to apply, and challenging to uphold in court.

Guidance and best practices for responsible regulation would ideally begin with the enterprises that build the tech. Governments also have an important role to play in collaborating with these technology providers in crafting the regulation that guides responsible usage of GenAI. Ideally, the two will work together to mitigate complex risks while also enabling GenAI to better the lives of people that do or do not use AI.

In the table below, we examine the emerging and complex ways in which these disruptive technologies can bring risk and ethical considerations – whether it is safety around inherent limitations, or regarding the technology's impact on society, enterprises, citizens, or the planet. We will also examine some recent government action to provide sensible and responsible regulation around these technologies.

▼

The Short Term—the fundamental limitations of AI technology and its effects on citizens

Class of Risk	Explanation of risk	What can be done?
False information	Machine outputs are sensible enough to sound like real insights, but are often factually incorrect	New AI fields that focus on output quality and include explainable AI (e.g., chain of thought), interpretable AI (e.g., sourcing output), ethical AI (e.g., reducing societal risks), and human-centered AI (e.g., reducing consumer risks).
Acceleration of Misinformation	Acceleration and spread of misinformation at an accelerated rate— sometimes intentional, and sometimes not	Regulation can combat misinformation, with emerging legislation in US states like Virginia, California, and Texas criminalizing deep fakes. The EU may mandate watermarking of AI-generated artifacts. China's Cyberspace Administration prohibits the use of gen-AI for activities that harm public interest and mandates auditing of AI-generated content and user prompts by service providers.
Data Biases	Biases that exist in the real world are often reflected in AI. Ensuring gender and racial equality and representation in the data is essential	To combat bias in AI, enterprises need to adopt best practices that include creating model datasets with a balanced mix of equality and accuracy. A centralized bias authority, either from a government or non-profit organization, would be ideal. Model developers can distinguish themselves by measuring the bias in their data and how it affects the quality of the gen-AI output.
Lack of data consent	Creative classes having their data trained for models, without any compensation or their consent	AI developers may face more lawsuits for using artists' data without consent, particularly with foundation models. To prevent this, model developers can make sure that all data in their datasets are proprietary and legal, as Adobe is doing with Firefly.
Data security	An employee could give away highly sensitive information through prompts	Employees need to treat Chat GPT as a person and undergo retraining. Confidentiality agreements are necessary to prevent the disclosure of trade secrets to OpenAI or other model owners.

Short-Medium Term – The societal concerns of Gen-AI

Class of Risk	Explanation of Risk	What can be done?
AI to create inequality and sow divisiveness	AI may become a polarizing force in society. Some may benefit disproportionately, and some may watch their capabilities be rendered useless. The result will be social polarization.	Governments must lead and regulate the development and usage of AI, collaborating with enterprises to meet the needs of citizens. A universal basic income may be part of the solution.
Information convergence	AI-generated content is often limited to a single perspective, resulting in an abundance of similar outputs that stifle creativity and innovation.	AI should be developed with an awareness of its limitations and the need for interpretability and explainability. This includes understanding the data that is not captured in its outputs. Meta AI is a new field in computer science that aims to address these issues.
Creative classes of work become automated	As gen-AI becomes more prevalent in the enterprise, there is a risk of automation replacing large groups of creative workers, rendering their skills obsolete.	Enterprises should wisely exercise caution in adopting Gen-AI and allow regulatory guidance to develop. They must also establish ethical principles to ensure responsible and beneficial uses of AI for their employees.

Medium-Long Term – Ordered effects of AI on humans and the planet

Class of Risk	Explanation of Risk	What can be done?
Technology Over-reliance	Technologies that replace humans in creative tasks may lead to overreliance and addiction to the technology, rather than relying on their own creative capacities, undermining the goal of improving human potential.	Education on AI and technology, including its capabilities and limitations, will ideally be implemented from a young age. Best practices in education should be established to ensure effective learning. Ideally, humans will be educated about the technology well-before the begin using it in daily life.
High energy consumption of model	High energy consumption, especially during model training and inference, results in an ever-increasing energy footprint.	AI must be designed to be green, with an eye for sustainability at every level of the tech step. Presently it, consumes an enormous amount of energy. For now, one way model developers can make an immediate improvement is to publish energy consumption information and aim to reduce their energy usage through sustainable best practices.

Section **07**

How Can We **Build a Better World with AI?**

Lab45 envisions a positive AI-driven future, urging businesses to consider how they can harness AI for the benefit of humanity.

As companies develop foundational, fine-tuned, and even personalized AI models, balancing human welfare and profit is crucial. We have the ability to harness this technology for incredible good.

There are many areas where AI can contribute to the collective good. The discovery of new medicines, the discovery of new biodegradable materials, enabling remote work, decentralizing currency, solving the climate crisis, accelerating our exploration of space, eradicating poverty, educating the masses, exploration of self, and securing our food and water resources— the list of potential good is endless.

Likewise, as discussed in the previous section, AI portents numerous risks for society at large.

The widespread popularization of GenAI makes 2023 a pivotal year. Business leaders, governments, and other organizations must carefully implement and/or handle these emerging technologies.

Key Insights

→ AI will revolutionize society, and bring about changes for the economy, governments, and regular citizens.

→ We're optimistic. The future looks bright.

07

The internet democratized access to knowledge, AI will democratize access to intelligence

Emergence of the Creator Economy

The internet democratized communication and led to a surge in user-generated content. GenAI is exponentially accelerating this democratization— allowing all of us to have access to tools that can help us create compelling content. We expect a new boom in content creation to explode onto the scene, with AI- powered content across various domains. We may see a deluge of AI-powered search, video games, stories, art, and marketing content.

AI will drive both content curation and composition. The challenge for the creator economy resides in curating and composing both the inputs and the outputs of the AI. Quality source data must be curated for the AI on the input side, while its outputs will need to be curated further and ultimately composed into an artifact.

GenAI could transform both the creator and curator economies. It will spawn both new types of content and new tools for personalized, synthetic content composition. It has the ability both to interpret data, and to infer from it. These capabilities are already transforming search with GenAI model's instant and contextualized responses. GenAI's impact on the internet may be as significant as the transition to the mobile device, reshaping how we interact with the digital world.

AI : as a Democratizing Force

Creativity is no longer solely a human domain. It is becoming something that can be synthesized, amplified, manipulated, and commodified. Society must proactively define human roles in the future of work, when AI capabilities may surpass human abilities. This may sound uncertain and potentially dystopian. Despite this, we must consider how AI can benefit individuals and the entire society at large.

We will always have the need to create, to learn, and to discover. Further, the need for human talent, creativity, and curation will always exist.

AI can potentially democratize creativity, empowering all of us in an intelligent and connected future.

Role of Governments and NGOs

Governments and NGOs must provide funding and policies to ensure AI reduces inequality instead of increasing it. As Bill Gates emphasizes, market forces alone won't create AI solutions for the poorest. The lack of regulation can lead to goals for the AI that are misaligned with human values. It may instead perpetuate inequality and allowing companies to define ethics and safety subjectively— with regard only for the profit motive. We believe in a world where we all thrive and AI should be used to bridge the inequality gap.

AI's benefits are immense, but costs and societal disruption must be considered, especially as the poor often bear the brunt of it. Careful management is needed to ensure AI's benefits reach all socio-economic classes and improve life for everyone.

Conclusion

Our curiosity to understand the world around us and our desire to innovate will continue to propel us forward. Progress cannot be stopped, and Artificial Intelligence represents progress. It is here to stay and will continue to evolve from its current form of large language models into more sophisticated reasoning and thinking machines.

Though the evolution of AI's thinking and reasoning capabilities is inevitable, its ability to develop awareness or consciousness remains uncertain. Further, as humans, we are yet to display a sincere understanding of awareness or consciousness ourselves— so our ability to create or recognize a conscious AI is uncertain.

For now, AI's ability to compute vast permutations & combinations, access large memory banks, and process information rapidly, could benefit humanity.

With AI, a tremendous gift has been handed to all of us. Working together with this incredible new tool, humans and AI will be able to create a new and better world; one of less drudgery, greater efficiency, and unbridled creativity.

The question to ask yourself as a leader, a business owner, an employee, or a citizen, is:

> **"How do I leverage this marvelous new technology to make lives better for me and better for the world around me?"**

Appendix

AI Model Training: **A How-To Explainer**

How to Train Your AI Model

Models are generally trained on vast amounts of data, which can be sourced from various unstructured formats, cleaned, and prepared for model training. Numerous processed datasets, each suitable for specific problem sets, are available for training different models.

When preparing a dataset for training, data needs to be divided into training data and a test set, which is used to evaluate the model training. When neural nets ₂ are trained, the network nodes are iteratively adjusted and weighted, so the model's output fits the intended predictive output. The importance of each node in the network is learned through training.

Datasets are divided into training data and a test set for evaluation. Neural network training involves iteratively adjusting nodes to achieve the desired output while learning each node's importance.

The data representation layer is crucial for efficient learning. Depending on the training method and model requirements, data can be represented as a flat list, knowledge graph, feature set, or vector embedding.

	Learning Description	Types of Problems	AI Techniques	Business Use Cases
Supervised Learning	Labelled data to map inputs to outputs. Useful when the target variable is already known, such as identifying images. Training is externally supervised.	Regression and classification	Linear regression, logistic regression, support vector machines, K Nearest neighbors, decision trees, random forest	Sales forecast, stock price analysis, identifying cancer cells in x-rays, emotion detection
Unsupervised Learning	Models discovering patterns in huge amounts of unlabeled data with no fixed output, trained to generate new sequences of data.	Clustering and association	K-means clustering, heirarchical clustering, db scan, principal component analysis	Customer segmentation based on behavior (clustering customers into a group), customer churn analysis, language generation
Reinforcement Learning	Trains machines to maximize rewards according to criteria of the situation. Uses an agent and an environment to produce actions and rewards; there is no predefined target variable.	Reward-based	Q-learning, monte carlo, Sarsa, Deep Q Network	Building games, task-training robots, fine-tuning language models

Generative Modeling Techniques

New generative techniques are responsible for expanding the use-cases, improving the performance, and shrinking the costs to generate new data. Different techniques use fundamentally different components to achieve their results and are best suited for varying use-cases. We will review some of the leading generative modeling techniques, which are driving the expansion of the gen-AI app ecosystem and consumer adoption. [3]

Variational autoencoders (VAEs)

Popularized encoder-decoder architecture for data generation using past data. VAEs excel in data compression, video-to-text captioning, and image generation.

Transformers:

Utilize positional encoding and self-attention mechanisms, enabling understanding and tracking of sequential data relationships. Transformers excel in NLP tasks and computer vision problems.

Generative Adversarial Networks (GANs)

Utilize positional encoding and self-attention mechanisms, enabling understanding and tracking of sequential data relationships. Transformers excel in NLP tasks and computer vision problems.

Diffusion

Gradually corrupts data with noise and trains a model to reverse the process. Diffusion models create high-fidelity images, are easier to train than GANs, and are highly effective.

These techniques can be combined for specific problems, with neural nets addressing classification, regression, and clustering. Depending on data types, different neural net structures like RNNs, LSTMs, or CNNs can be implemented.

Connecting the Input with the Output

Algorithms must build a semantic understanding of prompts to generate data for intended problem sets. Intent and Classification models extract entities and intent from prompts to achieve this.

The manipulator layer consists of algorithms identifying and classifying user desires from prompts. These models interact with user input and learn context from query histories. As models improve, personalized AI systems will adjust outputs based on user biases and preferences.

After understanding the inference, the neural network generates information matching the prompt, orchestrating it semantically. Workflows can run models sequentially, organizing outputs into unique modes tailored to the algorithm, aiming to solve user input probabilistically.

Works Cited

[i] GPT-4 Is Here. How Can Doctors Use Generative AI Now?, https://www.medpagetoday.com/special- reports/exclusives/103616

[ii] How Transformers Work, Medium https://towardsdatascience.com/transformers-141e32e69591

[iii] Speeding up drug discovery with diffusion generative models, MIT https://news.mit.edu/2023/speeding-drug-discovery-with-diffusion-generative-models-diffdock-0331

[iv] Generative AI Unlocking the Future of Fashion, McKinsey, https://www.mckinsey.com/industries/retail/our- insights/generative-ai-unlocking-the-future-of-fashion#:~:text=Generative%20AI%20has%20the%20potential%20to%20affect%20the%20entire%20fashion,communicati ons%2C%20and%20speed%20up%20processes.

[iv] Google to use Generative AI in its Ad business, Reuters https://www.reuters.com/technology/google-deploy-generative- ai-create-sophisticated-ad-campaigns-ft-2023-04-20/

[vi] Amazon scraps secret AI recruiting tool that showed bias against women, Reuters, https://www.reuters.com/article/us- amazon-com-jobs-automation-insight-idUSKCN1MK08G

[vii] Artificial intelligence in drug discovery and development, National Library of Medicine, https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7577280/

[viii] An AI system identified a potential new drug in just 46 days, MIT Tech Review https://www.technologyreview.com/2019/09/03/133175/an-ai-system-identified-a-potential-new-drug-in-just-46-days/

Further Reading

Azumo. A Quick Guide to Generative Models and GPT-3. Marina Bottacchi, Dec 2022

BBC. The Weird and Wonderful Art Created when AI and Humans Unite. Alexander Reben, Nov 2022

Bessemer Venture Partners. Entering the Era of Intelligent Search. Bhavik Nagda, Dec 2022

E2E Networks. Feedforward and Feedback Neural Networks: What's the difference? Niharika Srivastava, Sept 2022

Fast Company. How Generative AI will Impact the Future of Work, According to a CEO Who Has Spent 2,000 Hours Studying it. Erol Toker, Nov 2022

Forbes. AI: The Next Candidate For A Market Boom, Bubble, Bust. Clem Chambers, Nov 2022

Forbes. Twenty-Five Eye-Opening 2023 Predictions About Generative AI And ChatGPT Including A Splash Of AI Ethics And AI Law Tossed In. Lance Eliot, Dec 2022

Gartner. Top Trends in Generative AI 2022: Generative AI. Anthony Mullen, Nicole Greene, Brent Stewart, Marc Halpern, Soyeb Barot, Oct 2021

Gartner. Innovation Insight for Generative AI. Brian Burke, Arun Chandrasekaran, Erick Brethenoux, Anthony Mullen, Alexander Linden, Svetlana Sicular, Oct 2021

Gartner. Building a Digital Future: Emergent AI Trends Arun Chandrasekaran, Brian Burke, Erick Brethenoux, Apr 2022

Gartner. Go Beyond Machine Learning and Leverage Other AI Approaches. Ben Yan, Erick Brethenoux, Mike Fang, Pieter den Hamer, Apr 2022

Gartner. A CTO's Guide to Top Artificial Intelligence Engineering Practices. Arun Chandrasekaran, Farhan Choudhary, Erick Brethenoux, Oct 2021

Gartner. Emerging Tech Impact Radar: Artificial Intelligence. Annette Jump, Eric Goodness, Svetlana Sicular, Jim Hare, Alys Woodward,Alan Priestley, Erick Brethenoux, Jonathan Davenport, Eric Hunter, Danielle Casey,Ehtisham Zaidi, Oct 2022

The Indian Express. Here's how much it costs to run OpenAI's ChatGPT chatbot per day. Tech Desk, Dec 2022

LA Times. Beware a world where artists are replaced by robots. It's starting now. Molly Crabapple, Dec 2022

Machine Learning Mastery. A Gentle Introduction to Long Short-Term Memory Networks by the Experts Jason Brownlee, May 2017

McKinsey. Generative AI: Unlocking the Future of Fashion Holger Harreis, Theodora Koullias, Roger Roberts, Kimberly Te., March 2023

Medium. The AI Powering ChatGPT. Jesus Rodriguez, Dec 2022

MedPage Today. GPT-4 Is Here. How Can Doctors Use Generative AI Now? Michael DePeau Wilson, March 2023

ML6. Using generative AI for image manipulation: discrete absorbing diffusion models explained. Bert Christaens, Aug 2022

MIT Technology Review. An AI system identified a potential new drug in just 46 days Charlotte Gee, Sept 2019.

MIT Technology Review. Biotech labs are using AI inspired by DALL-E to invent new drugs. Will Douglas Heaven, Dec 2022

MIT Technology Review. Speeding up drug discovery with diffusion generative models. Alex Ouyang, March 2023

National Library of Medicine. Artificial intelligence in drug discovery and development. Debleena Paul, Gaurav Sanap, Snehal Shenoy, Dnyaneshwar Kalyane, Kiran Kalia, and Rakesh K. Tekade, Oct 2020

NBC News. An AI chatbot went viral. Some say it's better than Google, others worry it's problematic. Kalhan Rosenblatt, Dec 2022

NFX. Generative Tech Market Map and 5-Layer Tech Stack James Currier, Dec 2022

Reuters. Amazon scraps secret AI recruiting tool that showed bias against women Jeffrey Dastin, Oct 2018.

Reuters. Google to use Generative AI in its AD Business, Financial Times reports. Reuters staff, April 2023.

Skift. Generative AI Is the Travel Industry's Future, Get Used to It Rafat Ali, Dec 2022

Tom's Guide. Forget AI Art — AI-Powered Cyberattacks are What We Should Be Worried About. Anthony Spadafora, Dec 2022

Towards Data Science. Software 3.0 — How Prompting Will Change the Rules of the Game. Alberto Romero, Jun 2021

Towards Data Science. How Transformers Work Giuliano Giacaglia, March 2019.

VentureBeat. How Deep Learning Will Ignite the Metaverse in 2023 and Beyond. Victor Dey, Dec 2022

Vice. Artists Are Revolting Against AI Art on ArtStation Chloe Xiang, Dec 2022

Washington Post. Creative AI Is Generating Some Messy Problems. Parmy Olson, Dec 2022

World Economic Forum. Generative AI: a game-changer that society and industry need to be ready for Benjamin Larsen, Jayant Narayan, Jan 2023

Lead Authors @lab45

Nagendra Singh
Technology Futurist

Nagi is a trend jockey, with a track record of relentlessly spotting technology trends and analyzing them to accurately predict their impact on industries. He works on Artificial and Ambient Intelligence, Metaverse, Robotics and other emerging areas.

in

Tommy Mehl
Emerging Technology Strategist

Tommy is focused on designing technology systems from a first principles perspective. His passion lies in guiding emerging technologies to best serve enterprises, consumers, and society.

in

Siddhant Raizada
General Manager, Lab45

Passionate about building products that solve real world problems

in

Arvind Ravishunkar
GM, Emerging Tech.Strategy

Arvind specializes in identifying important trends among the noise of today's world. He works best in situations of ambiguity and diversity. He has an MBA from UC Berkeley and MSc in EE.

in

Contributing Authors @lab45

Arindam Chatterjee
Principal Data Scientist

Arindam works on core AI R&D at lab45. He is very passionate about research and emerging trends in AI. He has a Masters in Computer Science from IIT Bombay and is currently pursuing his Ph.D.

in

Noha El-Zehiry
Director of AI Research

Noha has over two decades of experience in advanced machine learning, AI, computer vision and image processing. She is a scientist and a patented inventor. Noha has a B.SC in EE, MA in applied mathematics and PhD in computer science.

in

Robert Walker Cohen
Senior Writer & Content Strategist

Robert is a business, finance, and technology writer and content strategist. He holds a BA in Philosophy from Columbia University.

in

Lab45 is a visionary space developing ground-breaking solutions to foster and accelerate ideation throughout Wipro.

At Lab45, engineers, research analysts, and scientists come together to infuse creative ways of incubating solutions for customers that will transform the future. It is a space filled with ambition at the vanguard of far-reaching research across cutting-edge technologies.

Established with the Silicon Valley culture of free-flowing creativity, Lab45's goal is to make bold ideas a reality and to invent the future of enterprise. So come, collaborate, and see what happens when ideas are left unbound.

Feedback		Click to Know More

It will take less than a minute!

Ambitions Realized.