

Operationalizing Generative AI on Vertex AI using MLOps

Authors: Anant Nawalgaria,
Gabriela Hernandez Larios, Elia Secchi,
Mike Styer, Christos Aniftos
and Onofrio Petragallo



Acknowledgements

Reviewers and Contributors

Nenshad Bardoliwalla

Warren Barkley

Mikhail Chrestkha

Chase Lyall

Lakshmanan Sethu

Erwan Menard

Curators and Editors

Antonio Gulli

Anant Nawalgaria

Grace Mollison

Technical Writer

Joey Haymaker

Designer

Michael Lanning





Table of contents

Introduction	5
What are DevOps and MLOps?	6
Lifecycle of a gen AI system	7
Discover	9
Develop and experiment	10
The foundational model paradigm	11
The core component of LLM Systems: A prompted model component	13
Chain & Augment	16
Tuning & training	20
Data Practices	23
Evaluate	27
Deploy	30
Deployment of gen AI systems	31
Version control	31
Continuous integration of gen AI systems	32
Continuous delivery of gen AI systems	33
Deployment of foundation models	34
Infrastructure validation	34

Compression and optimization	35
Deployment, packaging, and serving checklist	36
Logging and monitoring	37
Govern	41
The role of an AI platform for gen AI operations	42
Key components of Vertex AI for gen AI	43
Discover: Vertex Model Garden	44
Prototype: Vertex AI Studio & Notebooks	47
Customize: Vertex AI training & tuning	48
Train	49
Tune	49
Orchestrate	51
Chain & Augment: Vertex AI Grounding, Extensions, and RAG building blocks	52
Evaluate: Vertex AI Experiments, Tensorboard, & evaluation pipelines	55
Experiment	56
Evaluation	57
Predict: Vertex AI endpoints & monitoring	57
Govern: Vertex AI Feature Store, Model Registry, and Dataplex	59
Conclusion	61
Endnotes	63



Emergence of foundation models and generative AI (gen AI) has introduced a new era for building AI systems.

Introduction

The emergence of foundation models and generative AI (gen AI) has introduced a new era for building AI systems. Selecting the right model from a diverse range of architectures and sizes, curating data, engineering optimal prompts, tuning models for specific tasks, grounding model outputs in real-world data, optimizing hardware - these are just a few of the novel challenges that large models introduce.

This whitepaper delves into the fundamental tenets of MLOps and the necessary adaptations required for the domain of gen AI and Foundation Models. We also examine the diverse range of Vertex AI products, specifically tailored to address the unique demands of foundation models and gen AI-based applications. Through this exploration we uncover how Vertex AI, with its solid foundations of AI infrastructure and MLOps tools, expands its capabilities to provide a comprehensive MLOps platform for gen AI.

What are DevOps and MLOps?

DevOps is a software engineering methodology that aims to bridge the gap between development (Dev) and operations (Ops). It promotes collaboration, automation, and continuous improvement to streamline the software development lifecycle, introducing practices such as continuous integration and continuous delivery.

MLOps builds upon DevOps principles to address the unique challenges of operationalizing Machine Learning systems rapidly and reliably. In particular, MLOps tackles the experimental nature of ML through practices like:

- Data validation: Ensuring the quality and integrity of training data.
- Model evaluation: Rigorously assessing model performance with appropriate metrics.
- Model monitoring: Tracking model behavior in production to detect and mitigate drift.
- Tracking & reproducibility: Maintaining meticulous records for experiment tracking and result reproduction.

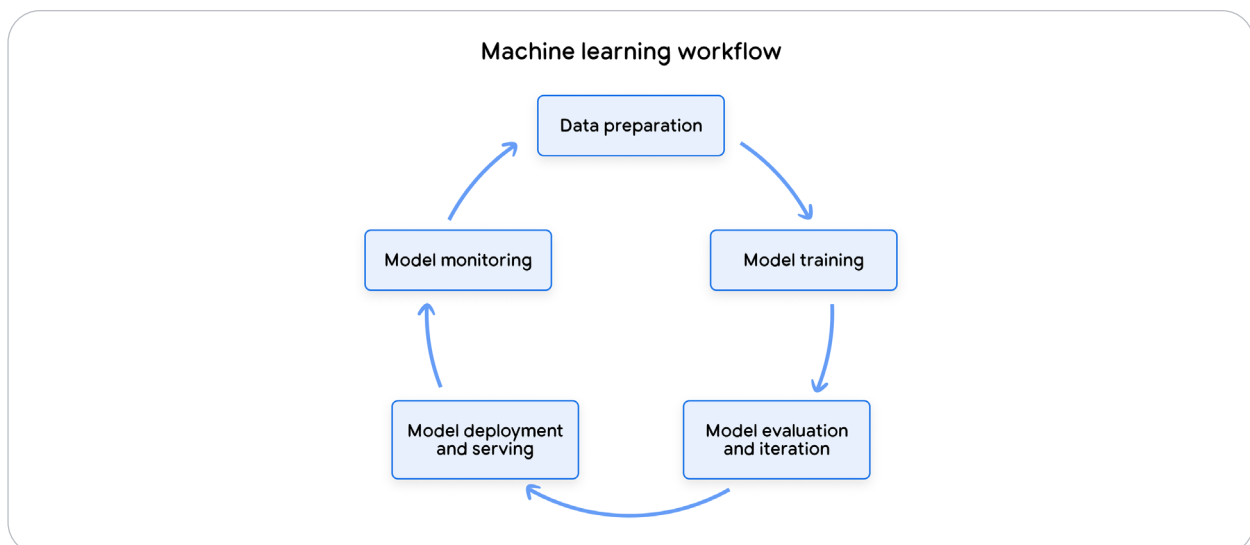


Figure 1. Machine learning workflow

Lifecycle of a gen AI system

Imagine deploying your first chatbot after months of dedicated work, and it's now interacting with users and answering questions. Behind this seemingly simple interaction lies the complex and fascinating life cycle of a gen AI System, which can be broken down into five key moments.

First in the **discovery phase**, developers and AI engineers must navigate the expanding landscape of available models to identify the most suitable one for their specific gen AI application. They must consider each model's strengths, weaknesses, and costs to make an informed decision.

Next, **development and experimentation** become paramount, with prompt engineering playing a crucial role in crafting and refining input prompts to elicit desired outputs based on an understanding of the model's intricacies. Few-shot learning, where examples are provided, can further guide model behavior, while additional customization may involve parameter-efficient fine-tuning (PEFT). Most gen AI systems also involve model chaining, which refers to orchestrating calls to multiple models in a specific sequence to create a workflow.

Data engineering practices have a critical role across all development stages, with factual grounding (ensuring the model's outputs are based on accurate, up-to-date information) and recent data from internal and enterprise systems being essential for reliable outputs. Tuning data is often needed to adapt models to specific tasks, styles, or to rectify persistent errors.

Deployment needs to manage many new artifacts in the deployment process, including prompt templates, chain definitions, embedding models, retrieval data stores, and fine-tuned model adapters among others. These artifacts each have unique governance requirements, necessitating careful management throughout development and deployment. Gen AI system deployment also needs to account for the technical capabilities of the target infrastructure, ensuring that system hardware requirements are fulfilled.

Continuous monitoring in production ensures improved application performance and maintains safety standards through responsible AI techniques, such as ensuring fairness, transparency, and accountability in the model's outputs.

Continuous Improvement as a concept is still key for Gen AI-based applications, though with a twist. For most Gen AI applications, instead of training models from scratch, we're taking foundation models (FMs) and then adapting them to our specific use case. This means constantly tweaking these FMs through prompting techniques, swapping them out for newer versions, or even combining multiple models for enhanced performance, cost efficiency, or reduced latency. Traditional **continuous training** still holds relevance for scenarios when recurrent fine-tuning or incorporating human feedback loops are still needed.

Naturally, this lifecycle assumes that the foundational model powering the gen AI system is already operationalized. It's important to recognize that not all organizations will be directly involved in this part of the process. In particular, the operationalization of foundational models is a specialized set of tasks that is typically only relevant for a select few companies with the necessary resources and expertise.

Because of that, this whitepaper will focus on practices required to operationalize gen AI applications using and adapting existing foundation models, referring to other whitepapers in the book should you want to deepdive into how foundational models are operationalized.

This includes active areas of research such as model pre-training, alignment (ensuring the model's outputs align with the desired goals and values), evaluation or serving.

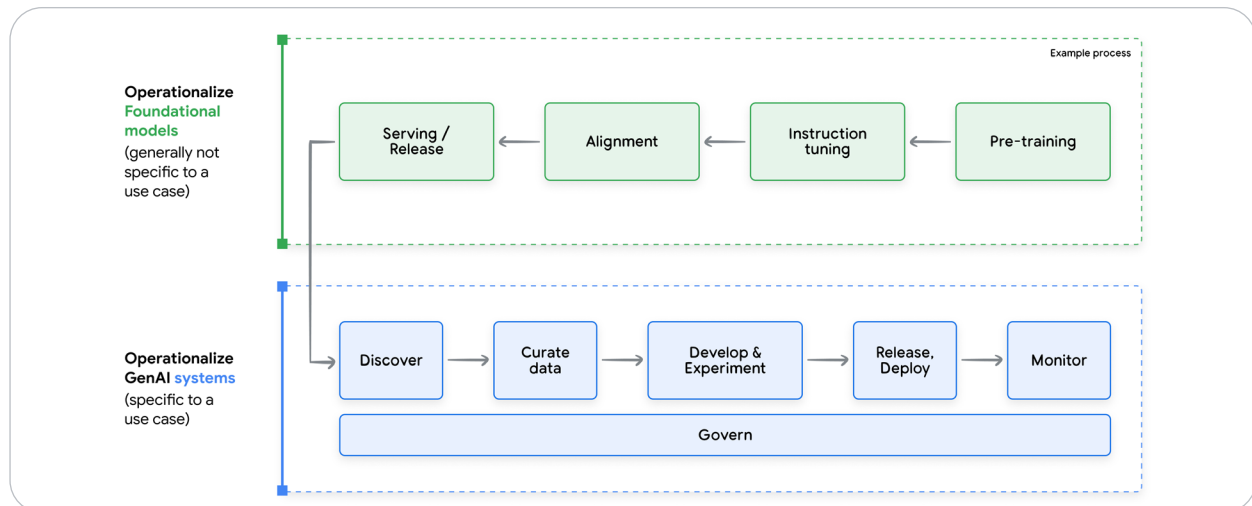


Figure 2. Lifecycle of a Foundational Model & gen AI system and relative operationalization practices

Discover

As mentioned before, building foundational models from scratch is resource-intensive. Training costs and data requirements are substantial, pushing most practitioners towards adapting existing foundation models through techniques like fine-tuning and prompt engineering. This shift highlights a crucial need: efficiently discovering the optimal foundation model for a given use case.

These two characteristics of the gen AI landscape make model discovery an essential MLOps practice:

1. **An abundance of models:** The past year has witnessed an explosion of open-source and proprietary foundation models. Navigating this complex landscape, each with varying architectures, sizes, training datasets, and licenses, requires a systematic approach to identify suitable candidates for further evaluation.

2. **No one-size-fits-all solution:** Each use case presents unique requirements, demanding a nuanced analysis of available models across multiple dimensions.

Here are some factors to consider when exploring models:

1. **Quality:** Early assessments can involve running test prompts or analyzing public benchmarks and metrics to gauge output quality.
2. **Latency & throughput:** These factors directly impact user experience. A chatbot demands lower latency than batch-processed summarization tasks.
3. **Development & maintenance time:** Consider the time investment for both initial development and ongoing maintenance. Managed models often require less effort than self-deployed open-source alternatives.
4. **Usage cost:** Factor in infrastructure and consumption costs associated with using the chosen model.
5. **Compliance:** Assess the model's ability to adhere to relevant regulations and licensing terms.

Because the activity of discovery has become so important for gen AI systems, many model discoverability platforms were created to support this need. An example of that is Vertex Model Garden,¹ which is explored later in this whitepaper.

Develop and experiment

The process of development and experimentation remains iterative and orchestrated while building gen AI applications. Each experimental iteration involves a tripartite interplay between data refinement, foundation model(s) selection and adaptation, and

rigorous evaluation. Evaluation provides crucial feedback, guiding subsequent iterations in a continuous feedback loop. Subpar performance might call for gathering more data, augmenting data, or further curating the data. Similarly, the adaptation of the foundation model itself might need tweaking - optimizing prompts, applying fine-tuning techniques, or even swapping it out for a different one altogether. This iterative refinement cycle, driven by evaluation insights, is just as critical for optimizing gen AI applications as it's always been for traditional machine learning.

The foundational model paradigm

Foundation models differ from predictive models most importantly because they are multi-purpose models. Instead of being trained for a single purpose, on data specific to that task, foundation models are trained on broad datasets, and therefore can be applied to many different use cases. This distinction brings with it several more important differences between foundation models and predictive models.

Foundation models also exhibit what are known as 'emergent properties',² capabilities that emerge in response to specific input without additional training. Predictive models are only able to perform the single function they were trained for; a traditional French-English translation model, for instance, cannot also solve math problems.

Foundation models are also highly sensitive to changes in their input. The output of the model and the task it performs are strongly affected, indeed determined, by the input to the model. A foundation model can be made to perform translation, generation, or classification tasks simply by changing the input. Even insignificant changes to the input can affect its ability to correctly perform that task.

These new properties of foundation models have created a corresponding paradigm shift in the practices required to develop and operationalize Gen AI systems. While models in the predictive AI context are self-sufficient and task-specific, gen AI models are multi-purpose and need an additional element beyond the user input to function as part of a gen AI Application: a prompt, and more specifically, a prompt template, defined as a set of instructions and examples along with placeholders to accommodate user input. A prompt template, along with dynamic data such as user input, can be combined to create a complete prompt, the text that is passed as input to the foundation model.

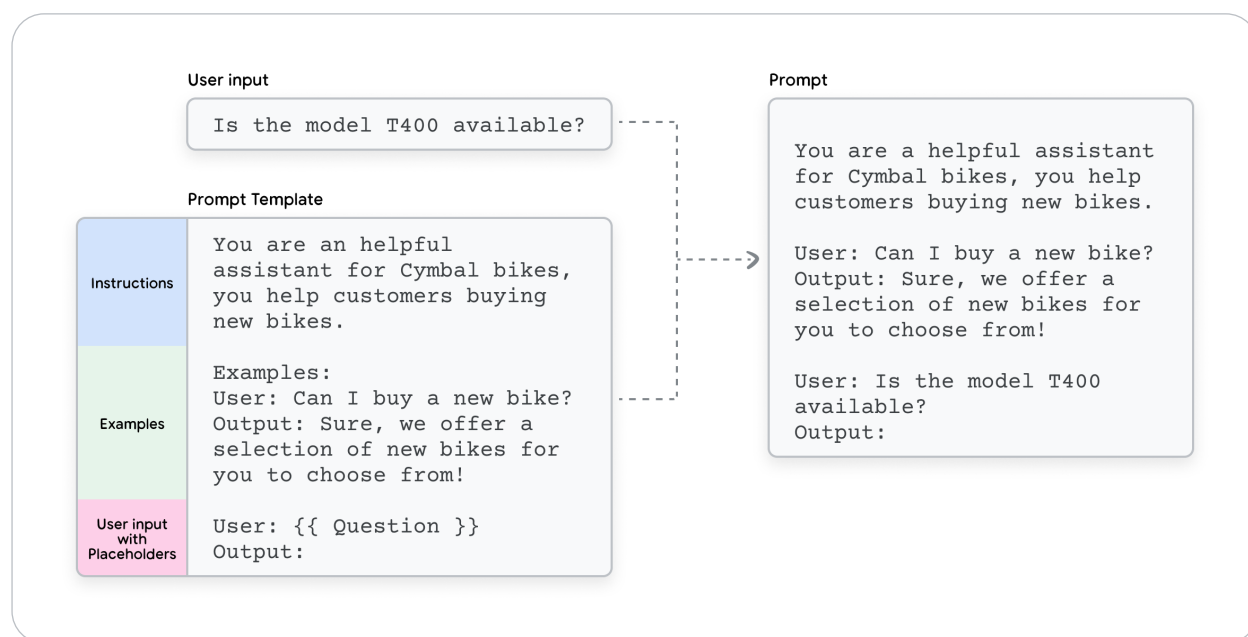


Figure 3. How Prompt Template and User input can be combined to create a prompt

The core component of LLM Systems: A prompted model component

The presence of the prompt element is a distinguishing feature of gen AI applications. Neither the model nor the prompt is sufficient for the generation of content; gen AI needs the combination of both. We refer to the combination as a ‘prompted model component’. This is the smallest independent component sufficient to create an LLM application. The prompt does not need to be very complicated. It can be a simple instruction, such as “translate the following sentence from English to French”, followed by the sentence to be translated. Without that preliminary instruction, though, a foundation model would not perform the desired translation task. So a prompt, even just a basic instruction, is necessary along with the input to get the foundation model to do the task required by the application.

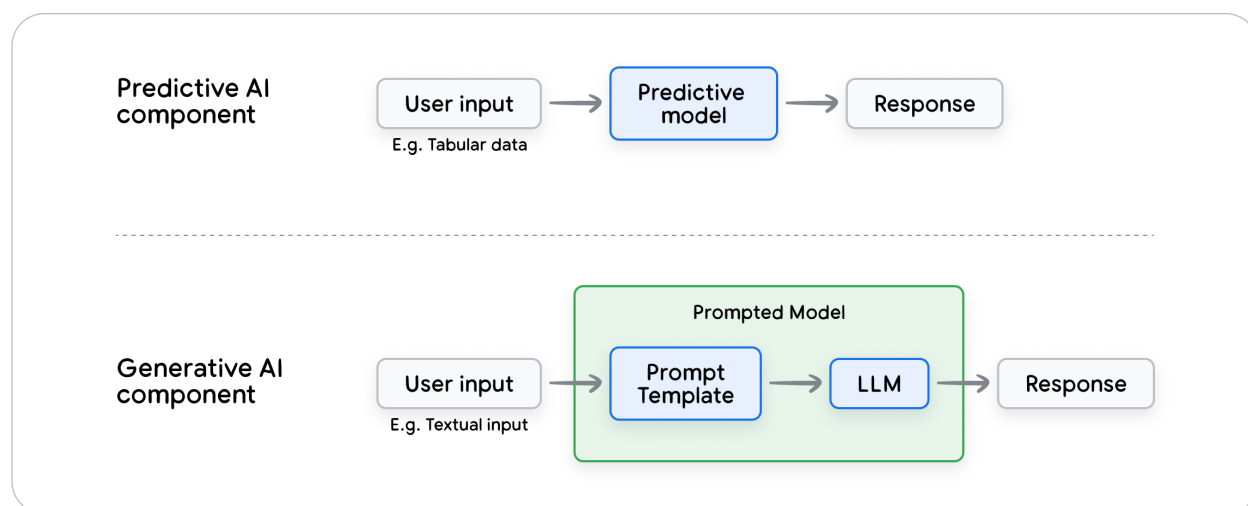


Figure 4. Predictive AI unit compared with the gen AI unit

This introduces an important distinction when it comes to MLOps practices for gen AI. In the development of a gen AI System, experimentation and iteration need to be done in the context of a *prompted model component*, the combination of a model and a prompt. The Gen

AI experimentation cycle typically begins with testing variations of the prompt – changing the wording of the instructions, providing additional context, or including relevant examples, etc., and evaluating the impact of those changes. This practice is commonly referred to as prompt engineering.

Prompt engineering involves two iterative steps:

1. **Prompting:** Crafting and refining prompts to elicit desired behaviors from a foundational model for a specific use case.
2. **Evaluation:** Assessing the model's outputs, ideally programmatically, to gauge its understanding and success in fulfilling the prompt's instructions.

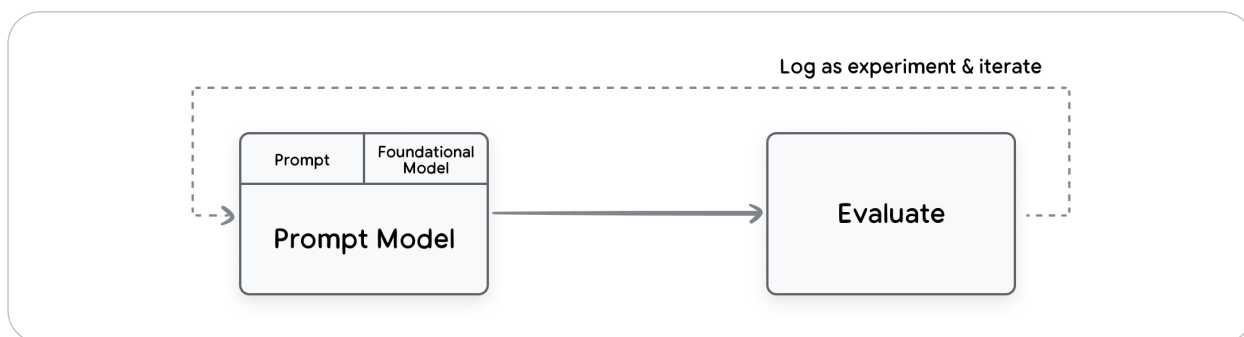


Figure 5. The activity of prompt engineering

Results of an evaluation can be optionally registered as part of an experiment, to allow for result tracking. Since the prompt itself is a core element of the prompt engineering process, it becomes a first class citizen within the artifacts part of the experiment.

However, we need to identify which type of artifacts they are. In the good old days of Predictive AI, we had clear lines - data was one thing, pipelines and code another. But with the “Prompt” paradigm in gen AI, those lines get blurry. Think about it: prompts can include anything from context, instructions, examples, guardrails to actual internal or external data pulled from somewhere else. So, are prompts data? Are they code?

To address these questions, a hybrid approach is needed, recognizing that a prompt has different components and requires different management strategies. Let's break it down:

Prompt as Data: Some parts of the prompt will act just like data. Elements like few-shot examples, knowledge bases, and user queries are essentially data points. For these components, we need data-centric MLOps practices such as data validation, drift detection, and lifecycle management.

Prompt as Code: Other components such as context, prompt templates, guardrails are more code-like. They define the structure and rules of the prompt itself. Here, we need code-centric practices such as approval processes, code versioning, and testing.

As a result, when applying MLOps practices to gen AI, it becomes important to have in place processes that give developers easy storage, retrieval, tracking, and modification of prompts. This allows for fast iteration and principled experimentation. Often one version of a prompt will work well with a specific version of the model and less well with a different version. In tracking the results of an experiment, both the prompt and its components version, and the model version must be recorded and stored along with metrics and output data produced by the prompted model.

The fact that development and experimentation in gen AI requires working with the prompt and the model together introduces changes in some of the common MLOps practices, compared to the predictive AI case in which experimentation is done by changing the model alone. Specifically, several of the MLOps practices need to be expanded to consider the prompted model component together as a unit. This includes practices like **evaluation, experiment tracking, model adaptation and deployment, and artifact management**, which will be discussed below in this whitepaper.

Chain & Augment

Gen AI models, particularly large language models (LLMs), face inherent challenges in maintaining recency and avoiding hallucinations. Encoding new information into LLMs requires expensive and data-intensive pre-training, posing a significant hurdle. Additionally, LLMs might be unable to solve complex challenges, especially when step-by-step reasoning is required. Depending on the use case, leveraging only one prompted model to perform a particular generation might not be sufficient. To solve this issue, leveraging a divide and conquer approach, several prompted models can be connected together, along with calls to external APIs and logic expressed as code. A sequence of prompted model components connected together in this way is commonly known as a chain.

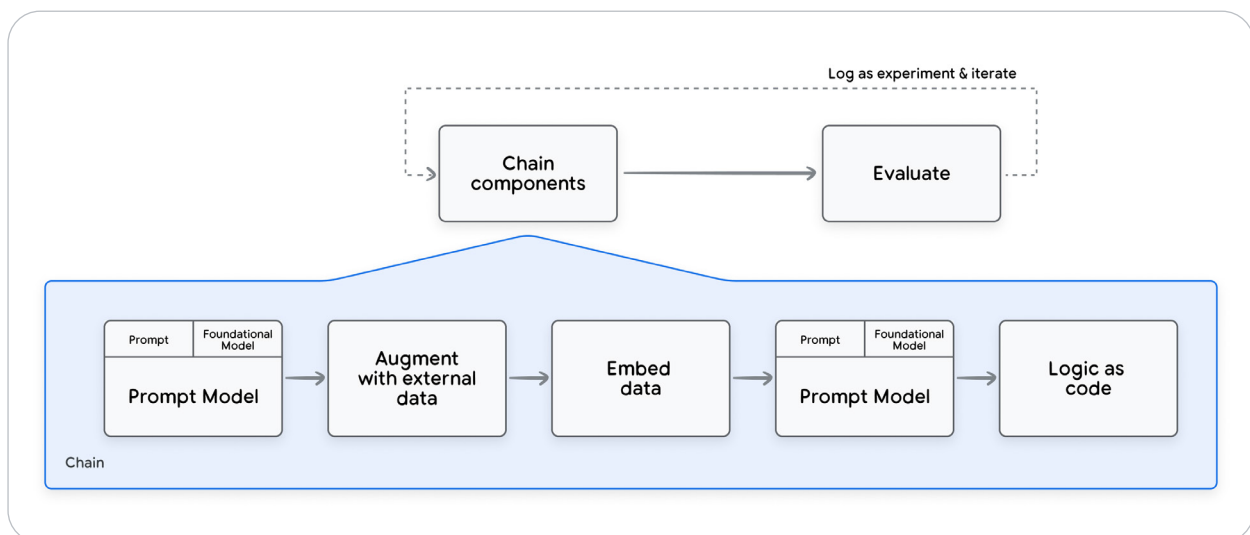


Figure 6. Components of a chain and relative development process

Two common chain-based patterns that have emerged to mitigate recency and hallucinations are retrieval augmented generation (RAG)³ and Agents.

- RAG addresses these challenges by augmenting pre-trained models with “knowledge” retrieved from databases, bypassing the need for pre-training. This enables grounding and reduces hallucinations by incorporating up-to-date factual information directly into the generation process.
- Agents, popularized by the ReAct prompting technique,⁴ leverage LLMs as mediators interacting with various tools, including RAG systems, internal or external APIs, custom extensions, or even with other agents. This enables complex queries and real-time actions by dynamically selecting and utilizing relevant information sources. The LLM, acting as an agent, interprets the user’s query, decides which tool to utilize, and how to formulate the response based on the retrieved information.

RAG and Agents approaches can be combined to create multi-agent systems connected to large information networks, enabling sophisticated query handling and real-time decision-making.

The orchestration of different models, logic and APIs is not a novelty of gen AI Applications. For example, recommendation engines have long combined collaborative filtering models, content-based models, and business rules to generate personalized product recommendations for users. Similarly, in fraud detection, machine learning models are integrated with rule-based systems and external data sources to identify suspicious activities.

What makes these chains of gen AI components different, is that, we can't a priori characterize or cover the **distribution** of component inputs, which makes the individual components much harder to evaluate and maintain in isolation.

This results in a paradigm shift in how AI applications are being developed for gen AI.

Unlike Predictive AI where it is often possible to iterate on the separate models and components in isolation to then chain in the AI application, in gen AI it's often easier to develop a chain in integration, performing experimentation on the chain end-to-end, iterating over chaining strategies, prompts, the underlying foundational models and other APIs in a coordinated manner to achieve a specific goal. No feature engineering, data collection, or further model training cycles is often needed; just changes to the wording of the prompt template.

The shift towards MLOps for gen AI, in contrast to predictive AI, brings forth a new set of demands. Let's break down these key differences:

1. **Evaluation:** Because of their tight coupling, chains need end-to-end evaluation, not just on a per-component basis, to gauge their overall performance and the quality of their output. In terms of evaluation techniques and metrics, evaluating chains is not dissimilar to evaluating prompted models. Please refer to the below segment on evaluation for more details on these approaches.
2. **Versioning:** A chain needs to be managed as a complete artifact in its entirety. The chain configuration should be tracked with its own revision history for analysis, reproducibility, and understanding the impact of changes on output. Logging should also include the inputs, outputs, and intermediate states of the chain, and any chain configurations used during each execution.

3. **Continuous Monitoring:** Establishing proactive monitoring systems is vital for detecting performance degradation, data drift, or unexpected behavior in the chain. This ensures early identification of potential issues to maintain the quality of the generated output. The activity of monitoring Chains is discussed in detail in the section 'Logging and Monitoring'.
4. **Introspection:** The ability to inspect the internal data flows of a chain (inputs and outputs from each component) as well as the inputs and outputs of the entire chain is paramount. By providing visibility into the data flowing through the chain and the resulting content, developers can pinpoint the sources of errors, biases, or undesirable behavior.

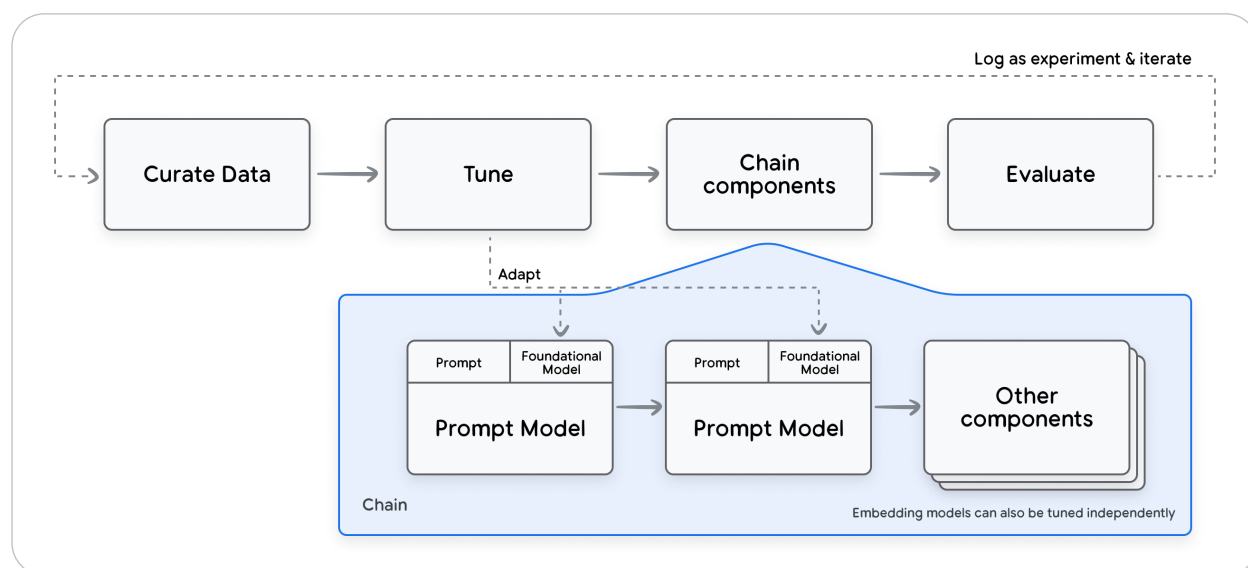


Figure 7. Putting together chains, prompted models and model tuning

There are several products in Vertex AI that can support the need for chaining and augmentation, including Grounding as a service,⁵ Extensions,⁶ and Vector Search,⁷ Agent Builder.⁸ We discuss the products in the section “Role of a AI Platform”. Langchain⁹ is also integrated with the Vertex SDK,¹⁰ and can be used alongside the core Vertex products to define and configure gen AI chained applications.

Tuning & training

When developing a gen AI use case and a specific task that involves LLMs, it can be difficult, especially for complex tasks, to rely on only prompt engineering and chaining to solve it. To improve task performance practitioners often also need to fine-tune the model directly. Fine-tuning lets you actively change the layers or a subset of layers of the LLM to optimize the capability of the model to perform a certain task. Two of the most common ways of tuning a model are:

1. Supervised fine-tuning: This is where we train the model in a supervised manner, teaching it to predict the right output sequence for a given input.
2. Reinforcement Learning from Human Feedback (RLHF): In this approach, we first train a reward model to predict what humans would prefer as a response. Then, we use this reward model to nudge the LLM in the right direction during the tuning process. Like having a panel of human judges guiding the model's learning.

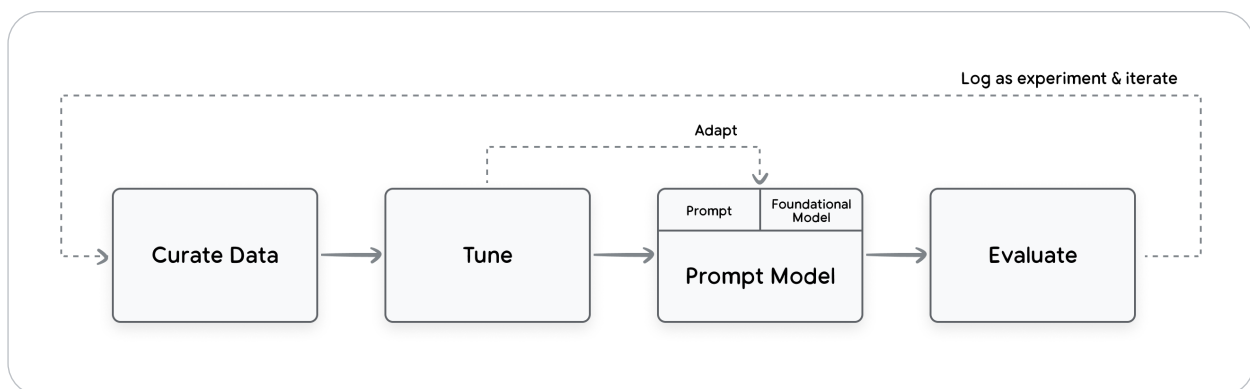


Figure 8. Putting together chains, prompted models and model tuning

When viewed through the MLOps lens, fine-tuning shares similar requirements with model training:

1. The capability to track artifacts being part of the tuning job. This includes for example the input data or the parameters being used to tune the model.
2. The capability to measure the impact of the tuning. This translates into the capability to perform evaluation of the tuned model for the specific tasks it was trained on and to compare results with previously tuned models or frozen models for the same task.

Platforms like Vertex AI¹¹ (and the Google Cloud platform more broadly) provide a robust suite of services designed to address these MLOps requirements: Vertex Model Registry,¹² for instance, provides a centralized storage location for all the artifacts created during the tuning job, and Vertex Pipelines¹³ streamlines the development and management of these tuning jobs. Dataplex,¹⁴ meanwhile, provides an organization-wide data fabric for data lineage and governance and integrates well with both Vertex AI and BigQuery.¹⁵ What's more, these products provide the same governance capability for both predictive and gen AI applications, meaning customers do not need separate products or configurations to manage generative versus AI development.

Continuous Training & Tuning

In machine learning operations (MLOps), continuous training is the practice of repeatedly retraining machine learning models in a production environment. This is done to ensure that the model remains up-to-date and performs well as real-world data patterns change over time. For gen AI models, continuous tuning of the models is often more practical than retraining from scratch due to the high data and computational costs involved.

The approach to continuous tuning depends on the specific use case and goals. For relatively static tasks like text summarization, the continuous tuning requirements may be lower. But for dynamic applications like chatbots that need constant human alignment, more frequent tuning using techniques like RLHF based on human feedback is necessary.

To determine the right continuous tuning strategy, AI practitioners must carefully evaluate the nature of their use case and how the input data evolves over time. Cost is also a major consideration, as the compute infrastructure greatly impacts the speed and expense of tuning. We discuss in detail monitoring of GenAI systems in the Logging and Monitoring section of this whitepaper.

Graphics processing units (GPUs) and tensor processing units (TPUs) are key hardware for fine-tuning. GPUs, known for their parallel processing power, are highly effective in handling the computationally intensive workloads and often associated with training and running complex machine learning models. TPUs, on the other hand, are specifically designed by Google for accelerating machine learning tasks. TPUs excel in handling large matrix operations common in deep learning neural networks.

To manage costs, techniques like model quantization can be applied. This represents model weights and activations using lower-precision 8-bit integers rather than 32-bit floats, which reduces computational and memory requirements.

We discuss in detail the support for tuning in Vertex AI in the Customize: Vertex AI Training & Tuning section.

Data Practices

Traditionally, ML model behavior was dictated solely by its training data. While this still holds true for foundation models – trained on massive, multilingual, multimodal datasets – gen AI applications built on top of them introduce a new twist: model behavior is now determined by how you adapt the model using different types of input data (Figure. 9).

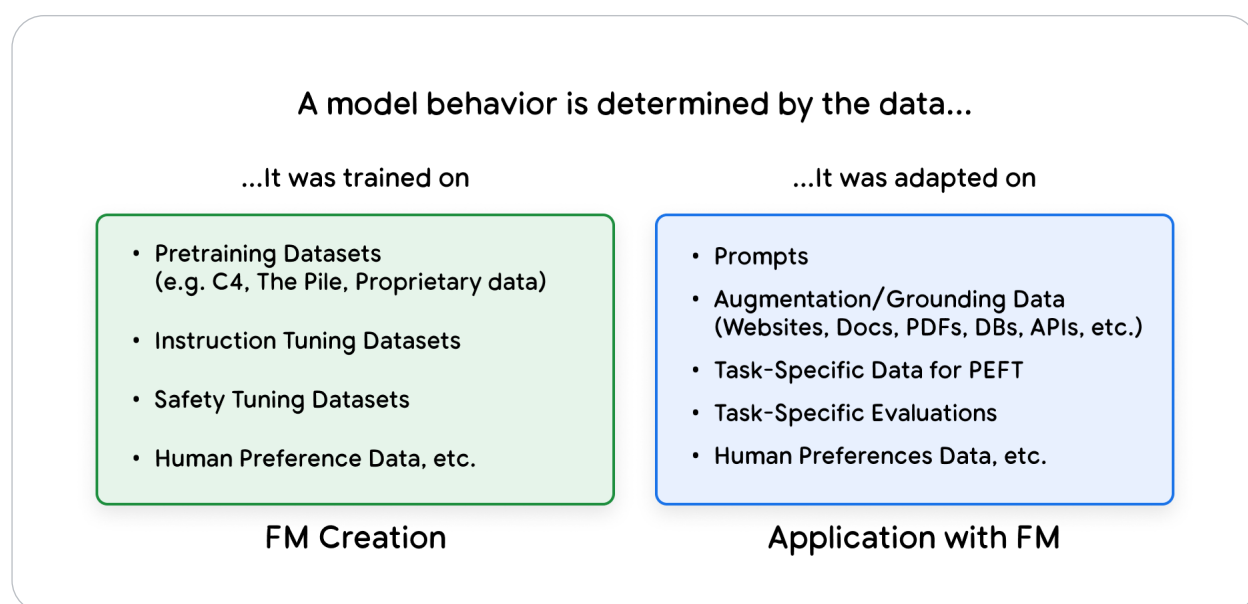


Figure 9. Examples of data spectrum for foundation models – creation (left) vs. adaptation (right)

The key difference between traditional predictive ML and gen AI lies in where you start. In predictive ML, the data is paramount. You spend a lot of time on data engineering, and if you don't have the right data, you cannot build an application. Gen AI takes a unique approach to this matter. You start with a foundation model, some instructions and maybe a few example inputs (in-context learning). You can prototype and launch an application with surprisingly little data.

This ease of prototyping, however, comes with a challenge. Traditional predictive AI relies on apriori well-defined dataset(s). In gen AI, a single application can leverage various data types, from completely different data sources, all working together (Figure 10). Let's explore some of these data types:

- **Conditioning prompts:** These are essentially instructions given to the Foundation Model (FM) to guide its output, setting boundaries of what it can generate.
- **Few-shot examples:** A way to show the model what you want to achieve through input-output pairs. This helps the model grasp the specific task(s) at hand, and in many cases, it boosts performances.
- **Grounding/augmentation data:** Data coming from either external APIs (like Google Search) or internal APIs and data sources. This data permits the FM to produce answers for a specific context, keeping responses current, relevant without retraining the entire FM. This type of data also supports reducing hallucinations.
- **Task-specific datasets:** These are used to fine-tune an existing FM for a particular task, improving its performance in that specific area.
- **Human preference datasets:** These capture feedback on generated outputs, helping refine the model's ability to produce outputs that align with human preferences.
- **Full pre training corpora:** These are massive datasets used to initially train foundation models. While application builders may not have access to them nor the tokenizers, the information encoded in the model itself will influence the application's output and performance.

This is not an exhaustive list. The variety of data used in gen AI applications is constantly growing and evolving.

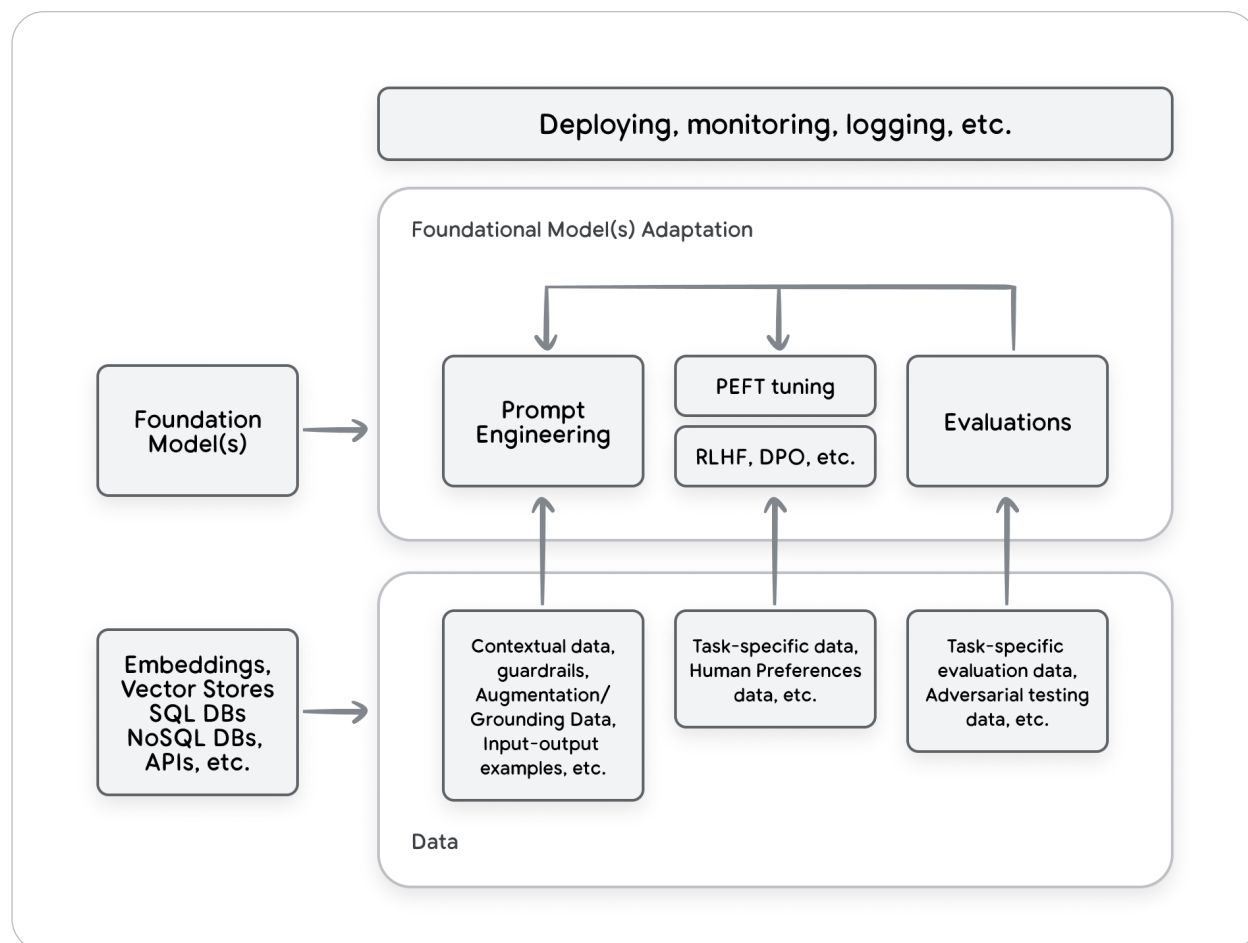


Figure 10. Example of high-level data and adaptations landscape for developing gen AI applications using existing foundation models

This diverse range of data adds another complexity layer in terms of data organization, tracking and lifecycle management. Take a RAG-based application as an example: it might involve rewriting user queries, dynamically gathering relevant examples using a curated set of examples, querying a vector database, and combining it all with a prompt template. This involves managing multiple data types: user queries, vector databases with curated few-shot examples and company information, and prompt templates.

Each data type needs careful organization and maintenance. For example, the vector database requires processing data into embeddings, optimizing chunking strategies, and ensuring only relevant information is available. The prompt template itself needs versioning and tracking, the user queries need rewriting, etc. This is where traditional MLOps and DevOps best practices come into play, with a twist. We need to ensure reproducibility, adaptability, governance, and continuous improvement using all the data required in an application as a whole but also individually. Think of it this way: in predictive AI, the focus was on well-defined data pipelines for extraction, transformation, and loading. In gen AI, it's about building pipelines to manage, evolve, adapt and integrate different data types in a versionable, trackable, and reproducible way.

As mentioned earlier, fine-tuning foundation models (FMs) can boost gen AI app performance, but it needs data. You can get this data by launching your app and gathering real-world data, generating synthetic data, or a mix of both. Using large models to generate synthetic data is becoming popular because it speeds things up, but it's still good to have a human check the results for quality assurance. Here are few ways to leverage large models for data engineering purposes:

1. **Synthetic data generation:** This process involves creating artificial data that closely resembles real-world data in terms of its characteristics and statistical properties, often being done with a large and capable model. This synthetic data serves as additional training data for gen AI, enabling it to learn patterns and relationships even when labeled real-world data is scarce.
2. **Synthetic data correction:** This technique focuses on identifying and correcting errors and inconsistencies within existing labeled datasets. By leveraging the power of larger models, gen AI can flag potential labeling mistakes and propose corrections, improving the quality and reliability of the training data.

3. **Synthetic data augmentation:** This approach goes beyond simply generating new data. It involves intelligently manipulating existing data to create diverse variations while preserving essential features and relationships. Thus, gen AI can encounter a broader range of scenarios during training, leading to improved generalization and ability to generate nuanced and relevant outputs.

Evaluating gen AI, unlike predictive AI, is tricky. You don't usually know the training data distribution of the foundational models. Building a custom evaluation dataset reflecting your use case is essential. This dataset should cover essential, average, and edge cases. Similar to fine-tuning data, you can leverage powerful language models to generate, curate, and augment data for building robust evaluation datasets.

Evaluate

Even if only prompt engineering is performed, as any experimental process, it does require evaluation in order to iterate and improve. This makes the evaluation process a core activity of the development of any gen AI systems.

In the context of gen AI systems, evaluation might have different degrees of automation: from entirely driven by humans to entirely automated by a process.

In the early days of a project, when you're still prototyping, evaluation is often a manual process. Developers eyeball the model's outputs, getting a qualitative sense of how it's performing. But as the project matures and the number of test cases balloons, manual evaluation becomes a bottleneck. That's when automation becomes key.

Automating evaluation has two big benefits. First, it lets you move faster. Instead of spending time manually checking each test case, you can let the machines do the heavy lifting. This means more iterations, more experiments, and ultimately, a better product. Second, automation makes evaluation more reliable. It takes human subjectivity out of the equation, ensuring that results are reproducible.

But automating evaluation for gen AI comes with its own set of challenges.

For one, both the inputs (prompts) and outputs can be incredibly complex. A single prompt might include multiple instructions and constraints that the model needs to juggle. And the outputs themselves are often high-dimensional - think a generated image or a block of text. Capturing the quality of these outputs in a simple metric is tough.

There are some established metrics, like BLEU for translations and ROUGE for summaries, but they don't always tell the full story. That's where custom evaluation methods come in. One approach is to use another foundational model as a judge. For example, you could prompt a large language model to score the quality of generated texts across various dimensions. This is the idea behind techniques like AutoSxS.¹⁶

Another challenge is the subjective nature of many evaluation metrics for gen AI. What makes one output 'better' than another can often be a matter of opinion. The key here is to make sure your automated evaluation aligns with human judgment. You want your metrics to be a reliable proxy for what people would think. And to ensure comparability between experiments, it's crucial to lock down your evaluation approach and metrics early in the development process.

Lack of ground truth data is another common hurdle, especially in the early stages of a project. One workaround is to generate synthetic data to serve as a temporary ground truth, which can be refined over time with human feedback.

Finally, comprehensive evaluation is essential for safeguarding gen AI applications against adversarial attacks. Malicious actors can craft prompts to try to extract sensitive information or manipulate the model's outputs. Evaluation sets need to specifically address these attack vectors, through techniques like prompt fuzzing (feeding the model random variations on prompts) and testing for information leakage.

Automating the evaluation process ensures speed, scalability and reproducibility
An automation of the evaluation process can be considered a proxy for the human judgment
Depending on the use case, the evaluation process will require a high degree of customization.
To ensure comparability it is essential to stabilize the evaluation approach, metrics, and ground truth data as early as possible in the development phase.
It is possible to generate synthetic ground truth data to accommodate for the lack of real ground truth data.
It is important to include test cases of adversarial prompting as part of the evaluation set to test the reliability of the system itself for these attacks.

Table 1. Key suggestions to approach evaluation of gen AI systems

Deploy

It should be clear by this point that production gen AI applications are complex systems with many interacting components. Some of the common components discussed include multiple prompts, models, adapter layers and external data sources. In deploying a gen AI system to production, all these components need to be managed and coordinated with the previous stages of gen AI system development. Given the novelty of these systems, best practices for deployment and management are still evolving, but we can discuss observations and recommendations for these components and indicate how to address the major concerns.

Deploying gen AI solutions necessarily involves multiple steps. For example, a single application might utilize several large language models (LLMs) alongside a database, all fed by a dynamic data pipeline. Each of these components potentially requires its own deployment process.

For clarity, we distinguish between two main types of deployment:

1. **Deployment of gen AI systems:** This focuses on operationalizing a complete system tailored for a specific use case. It encompasses deploying all the necessary elements - the application, chosen LLMs, database, data pipelines, and any other relevant components - to create a functioning end-user solution.
2. **Deployment of foundational models:** This applies to open-weight models, where the model weights are publicly available on platforms like Vertex Model Garden or Hugging Face, or privately trained models. Deployment in this scenario centers around making the foundational model itself accessible to users. Given their multipurpose nature, these deployments often aim to support various potential use cases.

Deployment of gen AI systems

Deployment of gen AI systems is broadly similar to deployment of any other complex software system. Most of the system components – databases, Python applications, etc. – are also found in other non-gen AI applications. As a result, our general recommendation is to manage these components using standard software engineering practices such as version control¹⁷ and Continuous Integration / Continuous Delivery (CI/CD).¹⁸

Version control

Gen AI experimentation is an iterative process involving repeated cycles of development, evaluation, and modification. To ensure a structured and manageable approach, it's crucial to implement strict versioning for all modifiable components. These components include:

- **Prompt templates:** Unless leveraging specific prompt management solutions, version them through standard version control tools like Git.
- **Chain definitions:** The code defining the chain (including API integrations, database calls, functions, etc.) should also be versioned using tools like Git. This provides a clear history and enables easy rollback if needed.
- **External datasets:** In retrieval augmented generation (RAG) systems, external datasets play a key role. It's important to track these changes and versions of these datasets for reproducibility. You can do that by leveraging existing data analytics solutions such as BigQuery, AlloyDB, Vertex Feature Store.
- **Adapter models:** The landscape of techniques like LoRA tuning for adapter models is constantly evolving. . You can leverage established data storage solutions (e.g. cloud storage) to manage and version these assets effectively.

Continuous integration of gen AI systems

In a continuous integration framework, every code change goes through automatic testing before merging to catch issues early. Here, unit and integration testing are key for quality and reliability. Unit tests act like a microscope, zooming in on individual code pieces, while integration testing verifies that different components work together.

The benefits of continuous integration in traditional software development are well-understood. Implementing a CI system helps to do the following:

1. **Ensure reliable, high-quality outputs:** Rigorous testing increases confidence in the system's performance and consistency.
2. **Catch bugs early:** Identifying issues through testing prevents them from causing bigger problems downstream. It also makes the system more robust and resilient to edge cases and unexpected inputs.
3. **Lower maintenance costs:** Well-documented test cases simplify troubleshooting and enable smoother modifications in the future, reducing overall maintenance efforts

These benefits are applicable to gen AI Systems as much as any software product. Continuous Integration should be applied to all elements of the system, including the prompt templates, the chain and chaining logic, and any embedding models and retrieval systems.

However, applying CI to gen AI comes with challenges:

1. **Difficult to generate comprehensive test cases:** The complex and open-ended nature of gen AI outputs makes it hard to define and create an exhaustive set of test cases that cover all possibilities.

2. Reproducibility issues: Achieving deterministic, reproducible results is tricky since generative models often have intrinsic randomness and variability in their outputs, even for identical inputs. This makes it harder to consistently test for specific expected behaviors.

These challenges are closely related to the broader question of how to evaluate gen AI systems. Many of the same techniques discussed in the Evaluation section above can also be applied to the development of CI systems for gen AI. This is an ongoing area of research, however, and more techniques will undoubtedly emerge in the near future.

Continuous delivery of gen AI systems

Once the code is merged, a continuous delivery process begins to move the built and tested code through environments that closely resemble production for further testing before the final deployment.

As mentioned in the "Develop and Experiment" segment, chain elements become one of the main components to deploy, as they fundamentally constitute the gen AI application serving users.

The delivery process of the gen AI application containing the chain may vary depending on the latency requirements and whether the use case is batch or online:

1. Batch use cases require deploying a batch process executed on a schedule in production. The delivery process should focus on testing the entire pipeline in integration in an environment close to production before deployment. As part of the testing process, developers can assert specific requirements around the throughput of the batch process itself and checking that all components of the application are functioning correctly (e.g., permissioning, infrastructure, code dependencies).

2. Online use cases require deploying an API, in this case, the application containing the chain, capable of responding to users at low latency. The delivery process should involve testing the API in integration in an environment close to production, with tests to assert that all components of the application are functioning correctly (e.g., permissioning, infrastructure, code dependencies). Non-functional requirements (e.g., scalability, reliability, performance) can be verified through a series of tests, including load tests.

Deployment of foundation models

Because foundation models are so large and complex, deployment and serving of these models raises a number of issues – most obviously, the compute and storage resources needed to run these massive models successfully. At a minimum, a foundation model deployment needs to include several key considerations: selecting and securing necessary compute resources, such as GPUs or TPUS; choosing appropriate data storage services like BigQuery or Google Cloud Storage that can scale to deal with the large datasets; and implementing model optimization or compression techniques.

Infrastructure validation

One technique that can be applied to address the resource requirements of gen AI systems is infrastructure validation. This refers to the introduction of an additional verification step, prior to deploying the training and serving systems, to check both the compatibility of the model with the defined serving configuration and the availability of the required hardware. There are a number of optional infrastructure validation layers that can perform some of these checks automatically. For instance, TFX¹⁹ has an infrastructure validation layer that checks

whether the model will run correctly on a specified hardware configuration, which can help catch configuration issues before deployment. Nevertheless, the availability of the required hardware still needs to be verified by hand by the engineer or the system administrator.

Compression and optimization

Another way of addressing infrastructure challenges is to optimize the model itself. Compressing and/or optimizing the model can often significantly reduce the storage and compute resources needed for training and serving, and in many cases can also decrease the serving latency.

Some techniques for model compression and optimization include quantization, distillation and model pruning. Quantization reduces the size and computational requirements of the model by converting its weights and activations from higher-precision floating-point numbers to lower-precision representations, such as 8-bit integers or 16-bit floating-point numbers. This can significantly reduce the memory footprint and computational overhead of the model. Model Pruning is a technique for eliminating unnecessary weight parameters or by selecting only important subnetworks within the model. This reduces model size while maintaining accuracy as high as possible. Finally, distillation trains a smaller model, using the responses generated by a larger LLM, to reproduce the output of the larger LLM for a specific domain. This can significantly reduce the amount of training data, compute, and storage resources needed for the application.

In certain situations, model distillation can also improve the performance of the model itself in addition to reducing resource requirements. This happens because the smaller model can combine the knowledge of the larger model with labeled data, which can help it to generalize better to new data on a limited use case. The process of distillation usually involves training a large foundational LLM (teacher model) and having it generate responses to certain tasks,

and then having the smaller LLM (student model) use a combination of the LLMs knowledge as well as task specific supervised dataset to learn. The size and complexity of the smaller LLM can be adjusted to achieve the desired trade-off between performance and resource requirements. A technique known as step-by-step distillation²⁰ has proven to achieve great results.

Deployment, packaging, and serving checklist

Following are the important steps to take when deploying a model on Vertex AI.

- ☐ **Configure version control:** Implement version control practices for LLM deployments. This allows you to roll back to previous versions if necessary and track changes made to the model or deployment configuration.
- ☐ **Optimize the model:** Perform any model optimization (distillation, quantization, pruning, etc.) before packaging or deploying the model.
- ☐ **Containerize the model:** Package the trained LLM model into a container.
- ☐ **Define target hardware requirements:** Ensure the target deployment environment meets the requirements for optimal performance of the LLM model, such as GPUs, as well as TPUs and other specialized hardware accelerators.
- ☐ **Define model endpoint:** Define the endpoint configuration using Vertex AI's endpoint creation interface or the Vertex AI SDK. Specify the model container, input and output formats, and any additional configuration parameters.
- ☐ **Allocate resources:** Allocate the appropriate compute resources for the endpoint based on the expected traffic and performance requirements.

- ❑ **Configure access control:** Set up access control mechanisms to restrict access to the endpoint based on authentication and authorization policies. This ensures that only authorized users or services can interact with the deployed LLM.
- ❑ **Create model endpoint:** Create a Vertex AI endpoint to deploy²¹ the LLM as a REST API service. This allows clients to send requests to the endpoint and receive responses from the LLM..
- ❑ **Configure monitoring and logging:** Establish monitoring and logging systems to track the endpoint's performance, resource utilization, and error logs.
- ❑ **Deploy custom integrations:** Integrate the LLM into custom applications or services using the model's SDK or APIs. This provides more flexibility for integrating the LLM into specific workflows or frameworks.
- ❑ **Deploy Real-time Applications:** For real-time applications, consider using Cloud Functions and Cloud Run in combination with LLMs hosted in Vertex AI to create a streaming pipeline that processes data and generates responses in real time.

Logging and monitoring

Monitoring gen AI applications and, as a result, their components, presents unique challenges that require additional techniques and approaches on top of those in traditional MLOps. The use of gen AI requires the chaining of components in order to produce results for practical applications. Additionally, to your application user, all the components are hidden. Therefore, the interface they have to your application is their input and the final output. This creates the need to log and monitor your application end-to-end: that is, logging and monitoring the input and output of your application overall as well as the input and output of every single component.

Logging is necessary for applying monitoring and debugging on your gen AI system in production. An input to the application triggers multiple components. Imagine the output to a given input is factually inaccurate. How can you find out which of the components are the ones that didn't perform well? To answer this question it is necessary to apply logging on the application level and at the component level. We need lineage in our logging for all components executed. For every component we need to log their inputs and outputs. We also need to be able to map those with any additional artifacts and parameters they depend on so we can easily analyze those inputs and outputs.

Monitoring can be applied to the overall gen AI application and to individual components. We prioritize monitoring at the application level. This is because if the application is performant and monitoring proves that, it implies that all components are also performant. You can also apply the same practices to each of the prompted model components to get more granular results and understanding of your application.

Skew detection in traditional ML systems refers to training-serving skew that occurs when the feature data distribution in production deviates from the feature data distribution observed during model training. In the case of Gen AI systems using pretrained models in components chained together to produce the output, we need to modify our approach. We can measure skew by comparing the distribution of the input data we used to evaluate our application (the test set as described under the Data Curation and Principles section above) and the distribution of the inputs to our application in production. Once the two distributions drift apart, further investigation is needed. The same process can be applied to the output data as well.

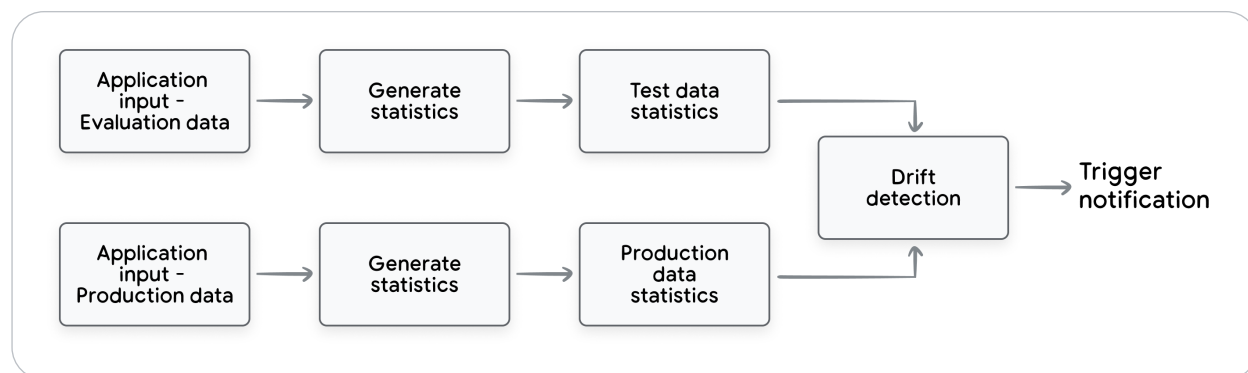


Figure 11. Drift/skew detection process overview

Like skew detection, the drift detection process checks for statistical differences between two datasets. However, instead of comparing evaluations and serving inputs, drift looks for changes in input data. This allows you to check how the inputs and therefore the behavior of your users changed over time. This is the same as traditional MLOps.

Given that the input to the application is typically text, there are a few approaches to measuring skew and drift. In general all the methods are trying to identify significant changes in production data, both textual (size of input) and conceptual (topics in input), when compared to the evaluation dataset. All these methods are looking for changes that could potentially indicate the application might not be prepared to successfully handle the nature of the new data that are now coming in. Some common approaches are calculating embeddings and distances, counting text length and number of tokens, and tracking vocabulary changes, new concepts and intents, prompts and topics in datasets, as well as statistical approaches such as least-squares density difference,²² maximum mean discrepancy (MMD),²³ learned kernel MMD,²⁴ or context-aware MMD.²⁵ As gen AI use cases are so diverse, it is often necessary to create additional custom metrics that better capture abnormal changes in your data.

Continuous evaluation is another common approach to GenAI application monitoring. In a continuous evaluation system, you capture the model's production output and run an evaluation task using that output, to keep track of the model's performance over time. One approach is collecting direct user feedback, such as ratings (for example thumbs up/down), which provides immediate insight into the perceived quality of outputs. In parallel, comparing model-generated responses against established ground truth, often collected through human assessment or as a result of an ensemble AI Model approach, allows for deeper analysis of performance. Ground truth metrics can be used to generate evaluation metrics as described in the Evaluation section. This process provides a view on how your evaluation metrics changed from when you developed your model to what you have in production today.

As with traditional monitoring in MLOps an alerting process should be deployed for notifying application owners when a drift, skew or performance decay from evaluation tasks is detected. This can help you promptly intervene and resolve issues. This is achieved by integrating alerting and notification tools into your monitoring process.

Monitoring expands beyond drift, skew and evaluation tasks. Monitoring in MLOps includes efficiency metrics like resources utilization and latency. Efficiency metrics are as relevant and important in gen AI as they are in any other AI application.

Vertex AI provides a set of tools that can help with monitoring. Model Evaluation for gen AI²⁶ tasks can be used for classification, summarization, question answering, and text generation tasks. Vertex Pipelines can be used to allow the recurrent execution of evaluation jobs in production as well as running pipelines for skew and drift detection processes.

Govern

In the context of MLOps governance encompasses all the practices, and policies that establish control, accountability, and transparency over the development, deployment, and ongoing management of machine learning (ML) models, including all the activities related to the code, data and models lifecycle.

As mentioned in the Develop & Experiment section the chain element and the relative components become a new type of assets that need to be governed over the full lifecycle from development to deployment, to monitoring.

The governance of the chain element lifecycle extends to lineage tracking practices as well.

While for predictive AI systems lineage focuses on tracking and understanding the complete journey of a machine learning model, in gen AI, lineage goes beyond the model artifact extending to all the components in the chain. This includes the data and models used and their lineage, the code involved and the relative evaluation data and metrics. This can help auditing, debugging and improvements of the models

Along with these new practices, existing MLOps and DevOps practices still apply to MLOps for gen AI:

1. The need to govern the data lifecycle; see “**Data Practices**”.
2. The need to govern the tuned model lifecycle; see “**Tuning and Training**”.
3. The need to govern the code lifecycle; see “**Deployment of GenAI System components**”.

The next segment will introduce a set of products that allow developers to perform governance of the data, model and code assets. We will discuss products like Google Cloud Dataplex, which centralizes the governance of model and data, Vertex ML Metadata and Vertex Experiment, which allows developers to register experiments, their metrics and artifacts.

The role of an AI platform for gen AI operations

Alongside the explosion of both predictive and gen AI applications, AI platforms, like Vertex AI,¹¹ have emerged as indispensable tools for organizations seeking to leverage the power of Artificial Intelligence (AI). These comprehensive platforms provide a unified environment that streamlines the entire AI lifecycle, from data preparation and model training to deployment, automation, continuous integration/continuous delivery (CI/CD), governance, and monitoring.

At the heart of an AI platform lies its ability to support diverse AI development needs. Whether you seek to utilize pre-trained AI solutions, adapt existing models through tuning or transfer learning, or embark on training your own large models, AI platforms provide the infrastructure and tools necessary to support these journeys. The advent of these platforms has revolutionized the way organizations approach AI, enabling them to productionize AI applications in a secure, enterprise-ready, responsible, controlled and scalable manner. These platforms accelerate innovation as well as foster reproducibility and collaboration while reducing costs and maximizing Return on Investment (ROI).

The new gen AI paradigm discussed in prior sections demands a robust and reliable AI platform that can seamlessly integrate and orchestrate a wide range of functionalities. These functionalities include model tuning for specific tasks; leveraging paradigms like retrieval augmented generation³ (RAG) to connect to internal and external data sources; and pre-training or instruction fine-tuning large models from scratch. Complex applications also often require chaining with other models, such as classifiers to route inputs to the appropriate LLM/ML model, extraction of customer information from a knowledge base, inclusion of safety checks, or even creation of caching systems for cost optimization.

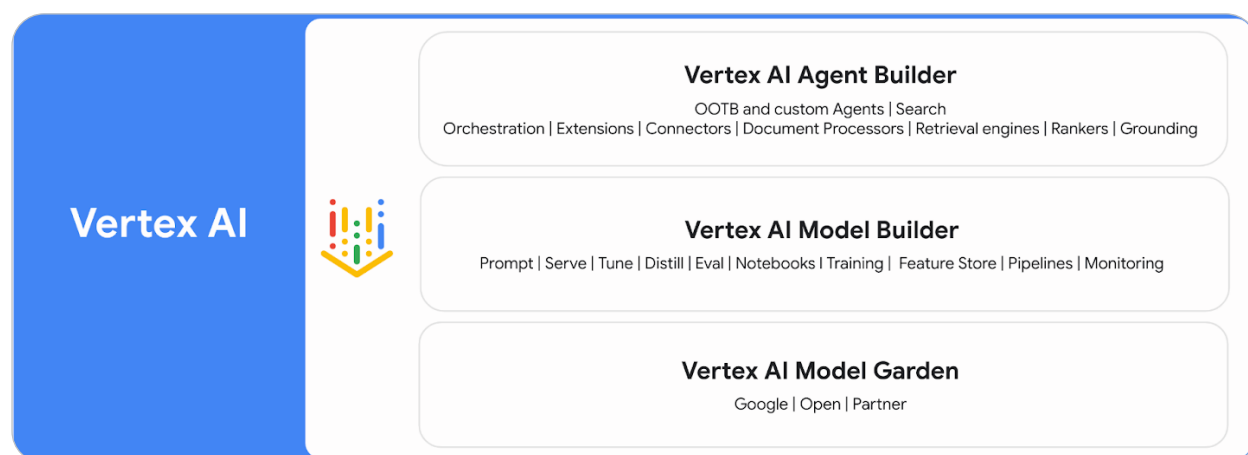


Figure 12. Key components of Vertex AI for gen AI

Key components of Vertex AI for gen AI

Vertex AI eliminates the complexities of managing the entire infrastructure required for AI development and deployment. Instead, Vertex AI offers a user-centric approach, providing on-demand access to the needed resources. This flexibility empowers organizations to focus on innovation and collaboration, rather than infrastructure management, and up-front hardware purchase. The features of Vertex AI that support gen AI development can be grouped into eight areas.

Discover: Vertex Model Garden

As discussed before, there is already a wide variety of available foundation models, trained on a broad range of datasets, and the cost of training a new foundation model can be prohibitive. Thus it often makes sense for companies to adapt existing foundation models rather than creating their own from scratch. As a result, a platform facilitating seamless discovery and integration of diverse model types is critical.

Vertex AI Model Garden¹ supports these needs, offering a curated collection of over 150 Machine Learning and gen AI models from Google, Google partners, and the open-source community. It simplifies the discovery, customization, and deployment of both Google's proprietary foundational models and diverse open-source models across a vast spectrum of modalities, tasks, and features. This comprehensive repository permits developers to leverage the collective research on artificial intelligence models within a single streamlined environment.

Model Garden encompasses a diverse range of modalities such as Language, Vision, Tabular, Document, Speech, Video, and Multimodal data. This broad coverage enables developers to tackle a multitude of tasks, including generation, classification, regression, extraction, recognition, segmentation, tracking, translation, and embedding. Model Garden houses Google's proprietary and foundational models (like Gemini,²⁷ PaLM 2,²⁸ Imagen²⁹) alongside numerous popular open source and third-party partner models like Llama 3,³⁰ T5 Flan,³¹ BERT,³² Stable Diffusion,³³ Claude 3 (Anthropic),³⁴ and Mistral AI.³⁵ Additionally, it offers task-specific models for occupancy analysis, watermark detection, text-moderation, text-to-video, hand-gesture recognition, product identification, and tag recognition, among others. Every model³⁶ in Vertex Model Garden has a model card which includes a description of the model, the main use cases that can cover, and the option (if available) to tune the model or deploy it directly.

Model Garden fosters experimentation by facilitating access to Google's proprietary foundational models through the Vertex AI Studio UI,³⁷ a playground where you can play around with prompts, models, and open-source models using provided Colab notebooks. One-click deployment is available for some external models, and there are more than 40 models available for fine-tuning for specific needs. Furthermore, the platform allows users to leverage technologies like vLLM³⁸ and quantization techniques for optimizing deployments for efficiency and reduced costs. We present below an overview of some of the models in Model Garden. For an up-to-date list, please visit.³⁶

Model Type	Description	Details
First-party models	Foundation models Leverage multimodal models from Google across vision, dialog, code generation, and code completion.	Gemini ³⁹ and Palm2 ⁴⁰ Imagen for text-to-image ⁴¹ Codey for code generation and completion ⁴² Chirp for speech-to-text ⁴³
First-party models	Pre-trained APIs Build and deploy AI applications faster with our pre-trained APIs powered by the best Google AI research and technology.	Text-to-Speech ⁴⁴ Natural Language processing ⁴⁵ Translation ⁴⁶ Vision ⁴⁷
Open models	Open source models Access a wide variety of enterprise-ready open source models	Google’s Gemma, ⁴⁸ PaliGemma, ¹⁶ CodeGemma ⁴⁹ Meta’s Llama ³⁰ TII’s Falcon ⁵⁰ Mistral AI ⁵¹ BERT, ³² T-5 FLAN, ³¹ ViT, ⁵² EfficientNet ⁵³
Third-party models	Third-party models Model Garden will support third-party models from partners with foundation models.	Anthropic’s Claude 3 Haiku, Sonnet and Opus ^{54,55}

Table 2. An overview of some of the models in Model Garden [Last Updated: March 18th, 2024]

Prototype: Vertex AI Studio & Notebooks

Rapid development and prototyping capabilities are also essential for developing gen AI applications. Vertex AI prioritizes inclusivity and flexibility in its development environments, catering to a wide range of developer preferences and proficiency levels. This platform provides options for both console-driven and programmatic development workflows. Users can leverage the intuitive web interface for end-to-end application creation or utilize various APIs for deeper customization and control. These include the REST API⁵⁶ and dedicated SDKs for Python,⁵⁷ NodeJS⁵⁸ and Java,⁵⁹ ensuring compatibility with diverse programming languages and ecosystems. Developers can choose to use the tools and IDEs of their choice for interacting with the platform, or take advantage of Vertex-native tools like Vertex Colab Enterprise or Vertex Workbench to explore and experiment with code within familiar notebook environments.

Vertex AI Studio⁶⁰ provides a unified console-driven entry point to access and leverage the full spectrum of Vertex AI's gen AI services. It facilitates exploration and experimentation with various Google first party foundation models (for example, PaLM 2, Gemini, Codey, Imagen, and Universal Speech Model). Additionally, it offers prompt examples and functionalities for testing distinct prompts and models with diverse parameters. It's also possible to adapt existing models through various techniques like supervised fine-tuning (SFT), reinforcement learning tuning techniques, and Distillation, and deploy gen AI applications in just a few clicks. Vertex AI Studio considerably simplifies and democratizes gen AI adoption, catering to a variety of users, from business analysts to machine learning engineers. You can see the homepage of Vertex AI Studio in Figure 13.

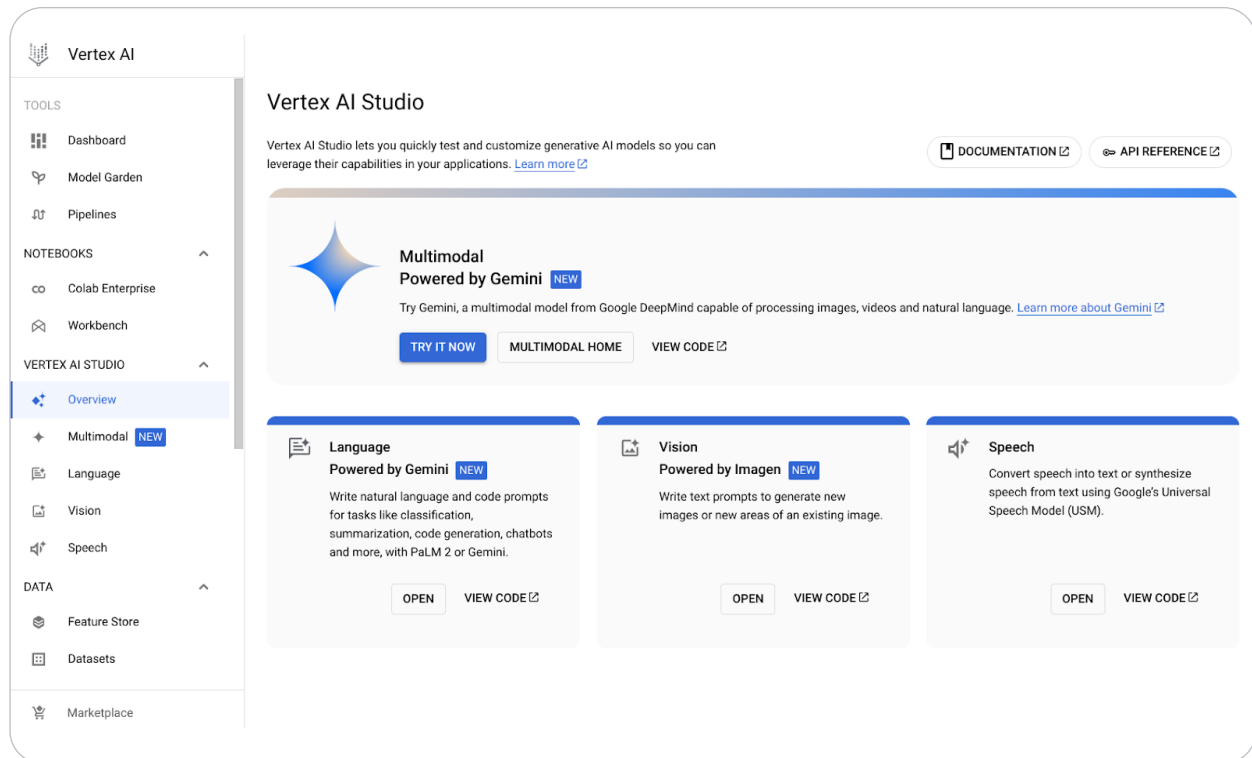


Figure 13. Vertex AI Studio - Homepage

Customize: Vertex AI training & tuning

While prompt engineering and augmentation are sufficient for some gen AI use cases, other cases require training, tuning and adapting the models to get the best results. Vertex AI provides a comprehensive platform for training and adapting LLMs, supporting a range of techniques and approaches from prompt engineering to training models from scratch.

Train

For full-scale LLM training, TPUs and GPUs are vital because of their superior processing power and memory capacity compared to CPUs. GPUs excel at parallel processing, enabling faster model training. TPUs, specifically designed for machine learning tasks, offer even faster processing and higher energy efficiency. This makes them ideal for large-scale, complex models. Google Cloud provides a range of offerings to support LLM training, including TPU VMs with various configurations, pre-configured AI platforms like Vertex AI, and dedicated resources like Cloud TPU Pods for scaling up training. These offerings allow users to choose the right infrastructure for their needs, accelerating LLM development and enabling cutting-edge research and applications.

Tune

Vertex AI also provides a comprehensive solution for adapting pre-trained LLMs. It supports a spectrum of techniques from a non-technical prompt engineering playground at inference time, to data-driven approaches involving tuning, reinforcement learning and distillation methods during the development or adaptation phase. The following five techniques – many of which are unique to Vertex AI – enable users to explore and implement them effectively. This applies to both proprietary and open-source LLMs, allowing you to achieve superior results while optimizing for costs and latency requirements.

- **Prompt engineering**⁶¹ leverages carefully crafted natural language prompts, potentially chained and enriched with external knowledge and examples, to nudge the LLM towards desired outputs without necessitating further training. Vertex AI through Vertex AI Studio offers a dedicated playground for crafting, testing, comparing and managing diverse prompts and techniques. Users can access various pre-built prompt templates within the platform and leverage public prompting guidelines⁶² for Google's proprietary large models.

- **Supervised fine-tuning (SFT)**⁶³ on Vertex AI facilitates model adaptation by leveraging a set of labeled examples (even a few hundred is enough) to tune a model on specific tasks and contexts within domain-specific datasets. The required examples resemble the one-shot example structure employed in the construction of a prompt. This effectively extends the few-shot learning approach for enhanced optimization. This focused tuning enables the model to encode additional parameters in the model necessary for mimicking desired behaviors such as improved complex prompt comprehension, adaptation to specific output formats, correcting errors, and learning new tasks. The SFT tuning approach on Vertex AI, minimizes computational overhead and time while yielding an updated model that integrates the newly acquired parameters with the original model's core parameters.
- **Reinforcement learning with human feedback (RLHF)**,⁶⁴ available on Vertex AI for foundational models like PaLM 2, and open-source models like T5 (s-xxl) and Llama2, leverages human feedback to train large models to align with human preferences. This technique is well-suited in complex tasks involving preference modeling and optimizes LLMs on intricate, sequence-level objectives not easily addressed by traditional supervised fine-tuning. The process involves first training a reward model using a human preference dataset, then utilizing it to score the output from the LLM, and finally applying reinforcement learning to optimize the LLM. This approach is recognized as a key driver of success in conversational large language models.
- **Distillation step-by-step**²⁰ is an advanced distillation technique transferring knowledge from a significantly larger model (known as teacher model) to a smaller task-specific model (known as student model), preserving important information while reducing model size. Step-by-Step Distillation²⁰ surpasses common techniques by requiring significantly less data. This method, accessible on Vertex AI,⁶⁵ significantly reduces inference costs and latencies while minimizing performance impact in the resulting smaller LLM.⁶⁶

Orchestrate

Any training or tuning job you run can be orchestrated and then operationalized using Vertex Pipelines,¹³ a service that aims to simplify and automate the deployment, management, and scaling of your ML workflows.

It provides a platform for building, orchestrating, scheduling and monitoring complex and custom ML pipelines, enabling you to efficiently translate your models from prototypes to production.

Vertex Pipelines is also the platform behind all the managed tuning and evaluation services for the Google Foundation Models on Vertex AI. This ensures consistency as you can consume and extend those pipelines easily, without having to familiarize yourself with many services.

Getting started with Vertex Pipelines is simple: you define the pipeline's step sequence in a Python file utilizing Kubeflow SDK.⁶⁷ For further details and comprehensive onboarding, consult the official documentation.⁶⁸

Chain & Augment: Vertex AI Grounding, Extensions, and RAG building blocks

Beyond training, tuning and adapting models and prompts directly, Vertex AI offers a comprehensive ecosystem for augmenting LLMs, to address the challenges of factual grounding and hallucination. The platform incorporates emerging techniques like RAG and agent-based approaches.

RAG overcomes limitations by enriching prompts with data retrieved from vector databases, circumventing pre-training requirements and ensuring the integration of up-to-date information. Agent-based approaches, popularized by ReAct prompting, leverage LLMs as mediators interacting with tools like RAG systems, APIs, and custom extensions. Vertex AI facilitates this dynamic information source selection, enabling complex queries, real-time actions, and the creation of multi-agent systems connected to vast information networks for sophisticated query processing and real-time decision-making.

Vertex AI function calling⁶⁹ empowers users by enhancing the capabilities of language models (LLMs). It enables LLMs to access real-time data and interact with external systems, providing users with more accurate and up-to-date information. To do that, users need to provide function definitions such as description, inputs, outputs to the gen AI model. Instead of directly executing functions, the LLM intelligently analyzes user requests and generates structured data outputs. These outputs propose which function to call and what arguments to use.

Vertex AI Grounding⁵ helps users connect large models with verifiable information by grounding them to internal data corpora on Vertex AI Agent Builder⁷⁰ or external sources using Google Search. This enables two key functionalities: verifying model-generated outputs against internal or external sources and creating RAG systems using Google's advanced search capabilities that produce quality content grounded in your own or web search data.

Vertex AI extensions⁶ let developers integrate Vertex Foundation Models with real-time data and real-world actions through APIs and functions, enabling task execution and allowing enhanced capabilities. This extends to leveraging 1st party extensions like Vertex AI Search⁷ and Code Interpreter,⁷¹ or 3rd party extensions for triggering and completing transactions. Imagine building an application that leverages the LLM's knowledge to plan a trip and seamlessly utilizes internal APIs to book hotels and flights, all within a single interface. Additionally, Vertex Extensions facilitate function calling with the gemini-pro model, enabling you to generate descriptions, pass them to the large model, receive JSON with function arguments, and automatically call the function.

Vertex AI Agent Builder⁷⁰ is an out-of-the-box solution that allows you to quickly build gen AI agents, to be used as conversational chatbots or as part of a search engine. With Vertex AI Agent Builder, you are able to easily ground your agents by pointing to a diverse range of data sources, including structured datastores such as BigQuery, Spanner, Cloud SQL, unstructured sources like website content crawling and cloud storage as well as connectors to Google drive and other APIs. Agent Builder utilizes a robust foundation of Google Search technologies, encompassing semantic search, content chunking, ranking, algorithms, and user intent understanding. Under the hood it optimizes document loading, chunking, embedding models, and ranking strategies. It abstracts away these complexities and allows users to simply specify their data source to initiate the gen AI-powered agent. This approach is ideal for organizations seeking to build robust search experiences for standard use cases without extensive technical expertise.

Vector databases are specialized systems for managing multi-dimensional data. This data, encompassing images, text, audio, video, and other structured or unstructured formats, is represented as vectors capturing its semantic meaning. Vector databases accelerate searching and retrieval within these high-dimensional spaces, enabling efficient tasks like finding similar images from billions or extracting relevant text snippets based on various

inputs. For a deeper dive into these topics, refer to 4 and 19. Vertex AI offers three flexible solutions for storing and serving embeddings at scale, catering to diverse use cases and user profiles.

Vertex AI Vector Search⁷ is a highly scalable low-latency similarity search and fully managed vector database scaling to billions of vector embeddings with auto-scaling. This technology, built upon ScaNN⁷² (a Google-developed technology used in products like Search, YouTube, and Play), allows you to search from billions of semantically similar or related items within your stored data. In the context of gen AI, the most common use cases where Vertex Vector Search can be used are:

1. Finding similar items (either text or image) based solely on their semantic meaning, in conjunction with an embedding model.
2. Creating a hybrid search approach that combines semantic and keyword or metadata search to refine the results.
3. Extracting relevant information from the database to feed into LLMs, enabling them to generate more accurate and informed responses.

Vertex AI Vector Search primarily functions as a vector database for storing pre-generated embeddings. These embeddings must be created beforehand using separate models like Vertex Embedding models⁷³ (namely `textembedding-gecko`, `text-embedding-gecko-multilingual`, or `multimodalembd`). Choosing Vertex Vector Search is optimal when you require control over aspects like the chunk, retrieval, query and models strategy. This includes fine-tuning an embedding model for your specific data. However, if your use case is a standard one requiring little customization, a readily available solution like Vertex Search might be a better choice.

Vertex AI Feature Store⁷⁴ is a centralized and fully managed repository for ML features and embedding. It enables teams to share, serve, and reuse machine learning features and embeddings effortlessly alongside other data. Its native BigQuery²³ integration eliminates duplication, simplifies lineage tracking and preserves data governance. Vertex AI Feature Store supports offline retrieval and an easy and fast online serving for machine learning features and embeddings. Vertex AI Feature Store is a good choice when you want to iterate and maintain different embedding versions alongside other machine learning features in a single place.

Vertex AI offers the flexibility to seamlessly create and connect various products to build your own custom grounding, RAG, and Agent systems. This includes utilizing diverse embedding models (multimodal, multilingual), various vector stores (Vector Search, Feature Store) and search engines like Vertex AI Agent Builder, extensions, grounding, and even SQL query generation for complex natural language queries. Moreover, Vertex AI provides SDK integration with LangChain⁹ to easily build and prototype applications using the umbrella of Vertex AI products. For further details and integration information, consult the official documentation⁷⁵ and official examples.⁷⁶

Evaluate: Vertex AI Experiments, Tensorboard, & evaluation pipelines

In the dynamic world of gen AI, experimentation and evaluation are the cornerstones of iterative development and continuous improvement. With a multitude of variables influencing Gen AI models (prompt engineering, model selection, data interaction, pretraining, and tuning), evaluation goes hand-in-hand with experimentation. The more seamlessly experiments and evaluations can be integrated into the development process, the

smoother and more efficient the overall development becomes. Vertex AI provides cohesive experimentation and evaluation products permitting connected iterations over applications and models alongside their evaluations.

Experiment

The process of selecting, creating, and customizing machine learning (including large models) and its applications involves significant experimentation, collaboration, and iteration.

Vertex AI seamlessly integrates experimentation and collaboration into the development lifecycle of AI/ML and gen AI models and applications. Its **Workbench Instances**⁷⁷ provide Jupyter-based development environments for the entire data science workflow, connected to other Google Cloud services and with GitHub synchronization capabilities. **Vertex Colab Enterprise**⁷⁸ accelerates the AI workflow by enabling collaborative coding and leveraging code completion and generation features.

Vertex AI also provides two tools for tracking and visualizing the output of many experiment cycles and training runs. **Vertex AI Experiments**⁷⁹ facilitates meticulous tracking and analysis of model architectures, hyperparameters, and training environments. It logs experiments, artifacts, and metrics, enabling comparison and reproducibility across multiple runs. This comprehensive tracking permits data scientists to select the optimal model and architecture for their specific use case. **Vertex AI TensorBoard**⁸⁰ complements the experimentation process by providing detailed visualizations for tracking, visualizing, and sharing ML experiments. It offers a range of visualizations, including loss and accuracy metrics tracking, model computational graph visualization, and weight and bias histograms, which - for example - can be used for tracking various metrics pertaining to training and evaluation of gen AI models with different prompting and tuning strategies. It also projects embeddings to lower-dimensional space, and displays image, text, and audio samples.

Evaluation

Vertex AI also provides a comprehensive set of evaluation tools for gen AI, from ground truth metrics to using LLMs as raters.

For Ground Truth-based metrics, **Automatic Metrics in Vertex AI**⁸¹ lets you evaluate a model based on a defined task and “ground truth” dataset. For LLM-based evaluation, **Automatic Side by Side (Auto SxS)** in Vertex AI⁸² uses a large model to evaluate the output of multiple models or configurations being tested, helping to augment human evaluation at scale.

In addition to that, users can also leverage Rapid Evaluation API, which offers a set of pre-built metrics for evaluating gen AI applications and relative SDK, integrated into the Vertex AI Python SDK for rapid and flexible, notebook-based, prototyping. To get started with Rapid Evaluation Vertex AI SDK see example in the official documentation.⁸³

Predict: Vertex AI endpoints & monitoring

Once developed, a production gen AI application must be deployed, including all its model components. If the application uses any models that have been trained or adapted, those models need to be deployed to their own serving endpoints. You can serve any model in the Model Garden through Vertex AI Endpoints²¹, which acts as the gateway for deploying your trained machine learning models. They allow you to serve online predictions with low latency, manage access controls, and monitor model performance easily through Model Monitoring. Endpoints also offer scaling options to handle varying traffic demands, ensuring optimal user experience and reliability.

Along with the prediction service, Vertex AI offers the following features for all Google managed models:

- **Citation checkers:** Gen AI on Vertex performs Citation checks⁷¹. Citations are important for LLMs and gen AI for several reasons. Citing sources ensures proper acknowledgment of sources and prevents plagiarism and demonstrates transparency and accountability. Citing sources is essential for LLMs and gen AI also because they help identify, understand potential biases, and enable reproducibility and verification. For example in Google Cloud,⁸⁴ the gen AI models are designed to produce original content, limiting the possibility of copying existing contents. If this happens, Google Cloud provides quotes for websites and code repositories.
- **Safety scores:** Safety attributes are crucial for LLMs and gen AI to mitigate potential risks like bias, lack of explainability, and misuse. These attributes help detect and mitigate biased outputs and mitigate misuse, enabling these tools to be used responsibly. As LLMs and gen AI evolve, incorporating safety attributes will be increasingly essential for responsible and ethical use. For example, Google Cloud added safety scores in Vertex AI PaLM API and Vertex AI Gemini API⁸⁵: content processed through the API is checked against a list of safety attributes, including "harmful categories" and sensitive topics. Each attribute has a confidence score between 0.0 and 1.0, indicating the likelihood of the input belonging to that category. These safety filters can be used in conjunction with all models: be it proprietary ones like Palm2 and Gemini or OSS ones like the ones available in Model garden.
- **Watermarking:** With AI-based tools becoming increasingly popular for creation of content, it's very important to identify if an image has been created using AI. Vertex AI offers digital watermarking and verification for AI-generated images⁸⁶ using the algorithm SynthID⁸⁷ developed by Google DeepMind.

- **Content moderation and bias detection:** By using the Content moderation⁸⁸ and Bias⁸⁹ detection tools on Vertex AI, you can add an extra layer of security on the responses of the LLMs to mitigate the risk that the model training and tuning may sway a model to generate outputs that aren't fair or appropriate for the task.

Govern: Vertex AI Feature Store, Model Registry, and Dataplex

Addressing the multifaceted requirements of data and model lineage and governance in gen AI requires a comprehensive strategy that tackles both conventional challenges and novel regulatory or technical complexities associated with large models. By adopting robust governance, observability, and lineage practices in the development of gen AI solutions, organizations can ensure comprehensive tracking, iteration, and evolution of data. They can also track the large models used, prompt adaptations, tuning, and other artifacts. This facilitates reproducibility of results, transparency and understanding of generated content sources, troubleshooting, compliance enforcement, and enhanced reliability and security. These practices collectively enable the ethical and responsible development and deployment of gen AI solutions. This fosters internal and external trust and fairness in gen AI models and practices. Vertex AI and Google Cloud offer the following comprehensive suite of tools for unified lineage, governance and monitoring, effectively addressing these critical concerns.

In the context of governance and lineage, **Vertex AI Feature Store**⁷⁴ offers:

- Track feature and embeddings versions and lineage, ensuring transparency
- Monitor feature (prompt) and embedding, response drift, and identify potential issues proactively
- Store feature formulas and discover relevant features or embeddings for different use cases

- Utilize feature selection algorithms to optimize model performance
- Consolidate and unify all machine learning data within a singular repository encompassing numerical data, categorical data, textual data, and embeddings representations

Vertex AI Model Registry¹² serves as a centralized repository for comprehensive lifecycle management of both Google proprietary foundational and open-source Machine Learning models. This includes gen AI models in addition to predictive models. This unified platform enables registration, storage, and version control of diverse model types, including various iterations of tuning for large models. Vertex AI Model Registry seamlessly integrates with **Vertex Pipelines**,¹³ facilitating orchestration and management of training and tuning jobs while leveraging lineage capabilities for recording and documenting the lineage from datasets to models and associated artifacts. It also couples with **Vertex AI Experiments**⁷⁹ and **Vertex AI Model Evaluation**,⁹⁰ enabling performance monitoring and comparison of different model versions alongside their artifacts – all within a single interface. Furthermore, Vertex AI Model Registry bolsters observability by providing integrated configuration and access to **Vertex AI Model Monitoring**⁹¹ and logging functionalities. This enables proactive identification and mitigation of both training-serving skew and prediction drift, ensuring reliability and accuracy of deployed models. Users can directly assign desired model versions to endpoints for one-click deployment from Vertex Model Registry or leverage aliases for simplified deployment.

Google Cloud Dataplex¹⁴ provides an organization-wide lineage across product boundaries in Google Cloud. Within the domains of AI and gen AI (and more broadly across data analytics and AI/ML) Dataplex seamlessly integrates with BigQuery and Vertex AI. Dataplex facilitates the unification, management, discovery, and governance of both data and models. Through comprehensive data lineage, quality, and metadata management capabilities it provides actionable insights for comprehensive data and model understanding. This promotes compliance, facilitates data analysis, and guarantees the training of machine learning models on trusted data sources. This in turn leads to enhanced accuracy and reliability. This

integration permits users across an organization to identify ‘champion models’ and ‘golden datasets and features’ across projects and regions in a secure way by adhering to identity access management (IAM)⁹² boundaries. In short, Dataplex encapsulates a framework within an organization that governs the interaction between people, processes and technology across all the products in Google Cloud.

Conclusion

The explosion of gen AI in the last several years introduced fundamental changes in the way AI applications are developed – but far from upending the MLOps discipline, these changes have only reinforced its basic principles and processes. As we have seen, the principles of MLOps that emphasize reliability, repeatability, and dependability in ML systems development are comfortably extended to include the innovations of gen AI. Some of the necessary changes are deeper and more far-reaching than others, but nowhere do we find any change that MLOps cannot accommodate.

As a result, many tools and processes built to support traditional MLOps can also support the requirements of gen AI. Vertex AI, for instance, is a powerful platform that can be used to build and deploy machine learning models and AI applications. It provides a comprehensive suite of functions for developing both Predictive and gen AI systems, encompassing data preparation, pre-trained APIs, AutoML capabilities, training and serving hardware, advanced fine-tuning techniques and deployment tools, and a diverse selection of proprietary and open-source foundation models. It also offers evaluation methods, monitoring capabilities, and governance tools, all unified within a single platform to streamline the AI development lifecycle. It’s built on Google Cloud Platform, which provides a scalable, reliable, secure and compliant infrastructure for machine learning. It’s a good choice for organizations that want to build and deploy machine learning models and AI applications.

The next few years will undoubtedly see gen AI extended in directions that today are unimaginable. Regardless of the direction these developments take, it will continue to be important to build on solid engineering processes that embody the basic principles of MLOps. These principles support the development of scalable, robust production AI applications today, and no doubt will continue to do so into the future.

Endnotes

1. Model Garden on Vertex AI. Available at: <https://cloud.google.com/model-garden>
2. Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, William Fedus. 2022. Emergent Abilities of Large Language Models. Available at: <https://arxiv.org/pdf/2206.07682.pdf>
3. Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, Douwe Kiela. 2022. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. Available at: <https://arxiv.org/pdf/2005.11401.pdf>
4. Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, Yuan Cao, Department of Computer Science, Princeton University, Google Research, Brain team, REACT: SYNERGIZING REASONING AND ACTING IN LANGUAGE MODELS. Available at: <https://arxiv.org/pdf/2210.03629.pdf>
5. Grounding in Vertex AI. Available at: <https://cloud.google.com/vertex-ai/docs/generative-ai/grounding/ground-language-models>
6. Vertex Extensions. Connect models to APIs by using extensions. Available at: <https://cloud.google.com/vertex-ai/docs/generative-ai/extensions/overview>
7. Overview of Vertex AI Vector Search. Available at: <https://cloud.google.com/vertex-ai/docs/vector-search/overview>
8. What is Vertex AI Agent Builder? Available at: <https://cloud.google.com/generative-ai-app-builder/docs/introduction>
9. LangChain. Get your LLM application from prototype to production. Available at: <https://www.langchain.com/>
10. Introduction to the Vertex AI SDK for Python. Available at: <https://cloud.google.com/vertex-ai/docs/python-sdk/use-vertex-ai-python-sdk>
11. Introduction to Vertex AI. Available at: <https://cloud.google.com/vertex-ai/docs/start/introduction-unified-platform>
12. Introduction to Vertex AI Model Registry. Available at: <https://cloud.google.com/vertex-ai/docs/model-registry/introduction>

13. Introduction to Vertex AI Pipelines. Available at: <https://cloud.google.com/vertex-ai/docs/pipelines/introduction>
14. Dataplex. Available at: <https://cloud.google.com/dataplex>
15. BigQuery. Available at: <https://cloud.google.com/bigquery?hl=en>
16. PaLi-Gemma model card. Available at: <https://ai.google.dev/gemma/docs/paligemma/model-card>
17. Version Control. Available at: https://en.wikipedia.org/wiki/Version_control
18. Continuous integration. Available at: https://wikipedia.org/wiki/Continuous_integration
19. TFX is an end-to-end platform for deploying production ML pipelines. Available at: <https://www.tensorflow.org/tfx>
20. Cheng-Yu Hsieh, Chun-Liang Li, Chih-Kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alexander Ratner, Ranjay Krishna, Chen-Yu Lee, Tomas Pfister. 2023. Distilling Step-by-Step! Outperforming Larger Language Models with Less Training Data and Smaller Model Sizes. Available at: <https://arxiv.org/pdf/2305.02301.pdf>
21. Vertex Endpoints. Use private endpoints for online prediction. Available at: <https://cloud.google.com/vertex-ai/docs/predictions/using-private-endpoints>
22. Tuan Duong Nguyen, Marthinus Christoffel du Plessis, Takafumi Kanamori, Masashi Sugiyama, 2014. Constrained Least-Squares Density-Difference Estimation. Available at: <https://www.ms.k.u-tokyo.ac.jp/sugi/2014/CLSDD.pdf>
23. Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, Alexander Smola, 2012. A Kernel Two-Sample Test. Available at: <https://jmlr.csail.mit.edu/papers/v13/gretton12a.html>
24. Oliver Cobb, Arnaud Van Looveren, 2022. Context-Aware Drift Detection. Available at: <https://arxiv.org/pdf/2203.08644.pdf>
25. Google Gemma Model. Available at: <https://gemini.google.com/>
26. Perform metrics-based evaluation. Available at: <https://cloud.google.com/vertex-ai/docs/generative-ai/models/evaluate-models>
27. Gemini Team, Google, 2023. Gemini: A Family of Highly Capable Multimodal Models. Available at: https://storage.googleapis.com/deepmind-media/gemini/gemini_1_report.pdf
28. Anil, Dai et al., 2023. PaLM 2 Technical Report. Available at: <https://arxiv.org/abs/2305.10403>

29. Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J Fleet, Mohammad Norouzi, 2022. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. Available at: <https://arxiv.org/abs/2205.11487>
30. Build the future of AI with Meta Llama 3. Available at: <https://llama.meta.com/llama3>
31. Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, Jason Wei. 2022. Scaling Instruction-Finetuned Language Models. Available at: <https://arxiv.org/abs/2210.11416>
32. Jacob Devlin, Ming-Wei Chang, Kenton Lee, Kristina Toutanova, 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. Available at: <https://arxiv.org/abs/1810.04805>
33. Stable Diffusion. Available at: <https://github.com/CompVis/stable-diffusion>
34. Vertex AI Function Calling. Available at: <https://cloud.google.com/vertex-ai/generative-ai/docs/multimodal/function-calling>
35. Mistral AI. Available at: <https://mistral.ai/>
36. Models available in Model Garden. Available at: <https://cloud.google.com/vertex-ai/docs/start/explore-models#available-models>
37. Vertex AI Studio. Customize and deploy generative models. Available at: <https://cloud.google.com/generative-ai-studio>
38. vLLM. Easy, fast, and cheap LLM serving for everyone. Available at: <https://github.com/vllm-project/vllm>
39. Overview of multimodal models. Available at: <https://cloud.google.com/vertex-ai/docs/generative-ai/multimodal/overview>
40. Text models. Available at: <https://cloud.google.com/vertex-ai/docs/generative-ai/model-reference/text>
41. Imagen on Vertex AI | AI Image Generator. Available at: <https://cloud.google.com/vertex-ai/docs/generative-ai/image/overview>
42. Code models overview. Available at: <https://cloud.google.com/vertex-ai/docs/generative-ai/code/code-models-overview>

43. Convert speech to text. Available at: <https://cloud.google.com/vertex-ai/docs/generative-ai/speech/speech-to-text>
44. Text-to-Speech AI. Available at: <https://cloud.google.com/text-to-speech>
45. Natural Language AI. Available at: <https://cloud.google.com/natural-language>
46. Translate docs, audio, and videos in real time with Google AI. Available at: <https://cloud.google.com/translate>
47. Vision AI. Available at: <https://cloud.google.com/vision>
48. Git. Available at: <https://git-scm.com/>
49. CodeGemma model card. Available at: https://ai.google.dev/gemma/docs/codegemma/model_card
50. TII's Falcon. Available at: <https://falconllm.tii.ae/>
51. Mistral AI. Available at: <https://mistral.ai/>
52. Hugging Face, 2024. Vision Transformer (ViT) Documentation. Hugging Face, [online] Available at: https://huggingface.co/docs/transformers/en/model_doc/vit
53. Mingxing Tan, Quoc V. Le, 2019. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. Available at: <https://arxiv.org/abs/1905.11946>
54. Anthropic Claude 3. Available at: <https://www.anthropic.com/news/claude-3-haiku>
55. Anthropic Claude 3 on Google Cloud Model Garden. Available at: <https://cloud.google.com/blog/products/ai-machine-learning/announcing-anthropics-claude-3-models-in-google-cloud-vertex-ai>
56. Vertex AI API. Available at: <https://cloud.google.com/vertex-ai/docs/reference/rest>
57. Vertex AI: Python SDK. Available at: <https://cloud.google.com/python/docs/reference/aiplatform/latest/vertexai>
58. Vertex AI: Node.js Client. Available at: <https://cloud.google.com/nodejs/docs/reference/aiplatform/latest/overview>
59. Vertex AI for Java. Available at: <https://cloud.google.com/java/docs/reference/google-cloud-aiplatform/latest/overview>
60. Customize and deploy generative models. Available at: <https://cloud.google.com/generative-ai-studio>

61. Design text prompts. Available at: <https://cloud.google.com/vertex-ai/docs/generative-ai/text/text-prompts>
62. Introduction to prompt design. Available at: <https://cloud.google.com/vertex-ai/docs/generative-ai/learn/introduction-prompt-design>
63. Supervised tuning. Available at: <https://cloud.google.com/vertex-ai/docs/generative-ai/models/tune-models#supervised-tuning>
64. RLHF model tuning. Available at: <https://cloud.google.com/vertex-ai/generative-ai/docs/models/tune-text-models-rlhf>
65. Vertex AI Distillation. Available at: <https://cloud.google.com/vertex-ai/generative-ai/docs/models/tune-text-models>
66. Create distilled text models. Available at: <https://cloud.google.com/vertex-ai/docs/generative-ai/models/distill-text-models>
67. Pipeline Basics. Available at: <https://www.kubeflow.org/docs/components/pipelines/v2/pipelines/pipeline-basics/>
68. Build a pipeline. Available at: <https://cloud.google.com/vertex-ai/docs/pipelines/build-pipeline>
69. Vertex AI Search extension. Available at: <https://cloud.google.com/vertex-ai/generative-ai/docs/extensions/vertex-ai-search>
70. What is Vertex AI Agent Builder? Available at: <https://cloud.google.com/generative-ai-app-builder/docs/introduction>
71. Generative AI on Vertex AI, Citation Check. Available at: https://cloud.google.com/vertex-ai/generative-ai/docs/learn/overview#citation_check
72. Ruiqi Guo, Philip Sun, Erik Lindgren, Quan Geng, David Simcha, Felix Chern, and Sanjiv Kumar, 2020. Accelerating Large-Scale Inference with Anisotropic Vector Quantization. Available at: <https://arxiv.org/pdf/1908.10396.pdf>
73. Get text embeddings. Available at: <https://cloud.google.com/vertex-ai/docs/generative-ai/embeddings/get-text-embeddings>
74. About Vertex AI Feature Store. Available at: <https://cloud.google.com/vertex-ai/docs/featurestore/latest/overview>

75. Google Cloud Vertex AI. Available at: https://python.langchain.com/docs/integrations/llms/google_vertex_ai_palm
76. Generative AI - Language - LangChain. Available at: <https://github.com/GoogleCloudPlatform/generative-ai/tree/main/language/orchestration/langchain>
77. Introduction to Vertex AI Workbench, Workbench Instances. Available at: <https://cloud.google.com/vertex-ai/docs/workbench/introduction>
78. Introduction to Colab Enterprise. Available at: <https://cloud.google.com/colab/docs/introduction>
79. Introduction to Vertex AI Experiments. Available at: <https://cloud.google.com/vertex-ai/docs/experiments/intro-vertex-ai-experiments>
80. Vertex AI TensorBoard Introduction to Vertex AI TensorBoard. Available at <https://cloud.google.com/vertex-ai/docs/experiments/tensorboard-introduction>
81. Perform metrics-based evaluation. Available at: <https://cloud.google.com/vertex-ai/docs/generative-ai/models/evaluate-models>
82. Perform automatic side-by-side evaluation. Available at: <https://cloud.google.com/vertex-ai/docs/generative-ai/models/side-by-side-eval>
83. Rapid Evaluation Vertex AI. Available at: <https://cloud.google.com/vertex-ai/generative-ai/docs/models/rapid-evaluation>
84. Citation metadata. Available at: https://cloud.google.com/vertex-ai/docs/generative-ai/learn/responsible-ai#citation_metadata
85. Responsible AI. Available at: <https://cloud.google.com/vertex-ai/docs/generative-ai/learn/responsible-ai#filters-palm-api>
86. Imagen on Vertex AI | AI Image Generator. Available at: <https://cloud.google.com/vertex-ai/docs/generative-ai/image/overview>
87. SynthID. Identifying AI-generated content with SynthID. Available at: <https://deepmind.google/technologies/synthid/>
88. Moderate text. Available at: <https://cloud.google.com/natural-language/docs/moderating-text>
89. Model bias metrics for Vertex AI. Available at: <https://cloud.google.com/vertex-ai/docs/evaluation/model-bias-metrics>

90. Model evaluation in Vertex AI. Available at: <https://cloud.google.com/vertex-ai/docs/evaluation/introduction>

91. Introduction to Vertex AI Model Monitoring. Available at: <https://cloud.google.com/vertex-ai/docs/model-monitoring/overview>

92. Identity and Access Management (IAM). Available at: <https://cloud.google.com/iam/docs>