Amazon Titan Text Embeddings

AWS AI Service Cards

AWS AI Service Cards: Amazon Titan Text Embeddings

Table of Contents

Amazon Titan Text Embeddings

An AWS AI Service Card explains the use cases for which the service is intended, how machine learning (ML) is used by the service, and key considerations in the responsible design and use of the service. A Service Card will evolve as AWS receives customer feedback, and as the service progresses through its lifecycle. AWS recommends that customers assess the performance of any AI service on their own content for each use case they need to solve. For more information, please see [AWS Responsible Use of AI Guide](#) and the references at the end. Please also be sure to review the [AWS Responsible AI Policy](#), [AWS Acceptable Use Policy](#), and [AWS Service Terms](#) for the services you plan to use.

This Service Card applies to the releases of Titan Text Embeddings (models G1 and V2) that are current as of December 3, 2024. In some documentation, model G1 is also referred to as model V1.

Overview

[Amazon Titan Text Embeddings](#) (TTE) is a family of text embedding models designed for enterprise use cases. Each TTE model converts natural language text input, including words, phrases, or paragraphs, into a vector of numbers (technically, points within an embedding space). Customers can use these vectors in downstream AI systems to solve a variety of use cases, such as question answering, search optimization, and text content grouping. TTE models are available via the [Amazon Bedrock API](#). Each TTE model is a managed sub-service of Amazon Bedrock; customers can focus on processing text input without having to provision or manage any infrastructure such as instance types, network topology, and endpoints.

A text embedding model is said to be "effective" if it maps blocks of text with similar meanings close together in the embedding space, while texts with dissimilar meanings are positioned farther apart. Otherwise, the model is said to be "ineffective". Cosine similarity is one industry-standard way of measuring semantic similarity. It ranges from [-1,1], where -1 indicates two text strings have opposite meanings, 0 indicates dissimilar meanings, and 1 indicates identical meanings. Titan Text Embeddings do not provide a confidence score for the embeddings they generate; a customer's workflow must decide if the model is effective using human judgment, whether human judgement is applied on a case-by-case basis or is applied via the customer's choice of an acceptable score on an automated test.

The "overall effectiveness" of a specific TTE model for a specific use case is based on the percentage of use-case specific text block pairs for which a specific model returns an effective

result. Customers should define and measure effectiveness for themselves for two reasons. First, the customer is best positioned to know which pairs will best represent their use case, and should therefore be included in an evaluation dataset. Second, properties like robustness and fairness are a function of the model and the evaluation dataset together, not the model by itself.

Titan Text Embeddings convert text to numerical vectors using the statistics learned from training datasets (see design section below). Like traditional ML solutions, TTE must overcome issues of intrinsic and confounding variation in the text data. Intrinsic variation refers to text features that differ semantically, and must be captured by the model. Confounding variation, on the other hand, involves variations in text that the model should ignore, such as typos or punctuations that do not alter the meaning . For example, TTE should recognize that the pair of sentences *'The cat sat on the mat'* and *'It will rain tomorrow'* have little semantic overlap, and that the pair of sentences *'I am very tired and need to get some sleep'* and *'I'm very tired, I need to get some sleep'* convey the same idea. TTE does do this. For the sentence pairs above, TTE V2 returns cosine similarities of 0.08 and 0.90 respectively, which are near 0 and 1, as expected.

With today's technology, text embedding models vary in their ability to represent text in a vector space while handling intrinsic and confounding variations. For a given model, higher-dimensional embeddings provide richer representations, but increase storage cost. TTE provides customers with options for embedding dimensions and numeric type (float and binary) so customers can optimize trade-offs between precision, cost, and latency according to their needs.

Intended use cases and limitations

Titan Text Embedding models are optimized for Retrieval-Augmented Generation (RAG) applications, but they serve a range of use cases, either on their own or with an LLM, such as search optimization, question answering, and content grouping:

- **Question answering** (QA): Customers can use TTE in RAG-based question answering systems. Embeddings enable relevant document retrieval from a large dataset by mapping both the query and documents into a vector space. This allows for quick similarity search, where the system retrieves the most relevant information from a knowledge base and then feeds it into an LLM for a more comprehensive answer.

- **Search optimization:** Customers can augment keyword matching, in which query terms are directly compared to keywords in library content, with semantic matching, in which the embedding of a query is directly compared with embeddings of passages from the library content. This enables searches to retrieve results that are conceptually aligned with a query, even if exact terms do not match.

- **Content grouping:** Customers can discover content categories by clustering the embeddings of passages from library content such as emails or customer reviews.

Amazon Titan Text Embeddings G1 supports input texts of up to 8192 tokens (roughly 6000-6500 words). Each input is converted into a vector of 1536 floating point numbers, that is, a 1536-dimensional embedding.

Amazon Titan Text Embeddings V2 supports input texts of up to 8192 tokens as well. However, each input can be converted into embeddings of 256, 384, or 1024 floating point dimensions. These smaller sizes reduce cost over G1. Finally, V2 also offers a binary option, which allows each floating-point number to be stored in binary format, further saving storage.

When assessing the text embedding model for a particular use case, we encourage customers to specifically define the use case by considering at least the following factors: the **business problem** being solved; the **stakeholders** in the business problem and deployment process; the **workflow** that solves the business problem, with the model and human oversight as components; **key system inputs and outputs**; the expected intrinsic and confounding **variation**; and the types of **errors** possible and the relative impact of each.

Consider the following example use case of using text embeddings in a customer service chatbot for an e-commerce company. The **business problem** is to reduce the time and cost of answering customer service questions. The **stakeholders** include customer service agents, customers, the IT team of the company, third-party merchandise vendors, and the company management. The **workflow** involves using TTE to turn the customer inquiries into embeddings, retrieving the most relevant information based on a pre-indexed embedding database (by converting all documentations into embeddings using TTE ahead of time, and only updating when the documents are updated), and then passing the results to an LLM to either provide automated responses or route them to human agents. The key **system inputs** are customer queries, previous conversation history, and a document database that provides customer service policies and merchandise information. The **key outputs** are retrieved documents, automated responses, and agent recommendations. The **intrinsic variation** of the input includes different customer inquiry topics, different English dialects, and varying inquiry length; the **confounding factors** include ambiguous wording, grammatical mistakes, and typos. The types of **errors** include 1/retrieved documents are not relevant to the query; 2/ the most relevant query is not retrieved; 3/ retrieved document contains the opposite or contradictory information of the query; 4/ inconsistent retrieval results on similar queries. When users ask a question, the embedding model converts the query into a vector representation and retrieves the most relevant support documents from the embedding database based on semantic similarity, even if the phrasing differs from the indexed

documents. For example, given the user query: *'how do I remove my old payment method?'*, the embedding model would match the query with a document titled with *'Delete saved payment methods'*, even though the query uses different wording and phrasing with the retrieved document.

Titan Text Embedding models have a number of limitations requiring careful consideration.

Appropriateness for Use

Because embedding spaces are a compressed representation of the input text, a Titan Text Embedding model may produce embeddings that position pairs of unrelated or related concepts as too similar or dissimilar. TTE models do not filter text input based on content. Customers should evaluate outputs for accuracy and appropriateness for their use case. Additionally, if a TTE model is used in customer workflows that produce consequential decisions, customers must evaluate the potential risks of their use case and implement appropriate human oversight, testing, and other use case-specific safeguards to mitigate such risks. See the [AWS Responsible AI Policy](#) for additional information.

Natural Languages

TTE G1 is trained on [25 languages](#) and TTE V2 is trained on [112 languages](#). While TTE models work in all trained languages, our internal testing has been most extensive for English text. We recommend that customers perform their own testing to determine the utility of a TTE model, especially when the text for a use case could include professional jargon, idiomatic language, or words that vary in meaning across language dialects.

Programming Languages

TTE G1 is not trained on programming languages. However, TTE V2 is trained on 13 programming languages: Python, Java, C, Ruby, Rust, Go, JavaScript, C++, C#, Kotlin, PHP, SQL, and Shell. While Titan Text Embedding models work in all trained programming languages, we recommend customers perform their own testing to determine if TTE models are suitable for their specific use case.

Input size

The maximum input text size is measured in tokens. For English, a token is approximately [4.7](#) characters (English words average about five characters). For all models, the maximum input text size is 8192 tokens. For longer text documents, we recommend that customers segment documents into logical segments, such as paragraphs or sections.

Design of Titan Text Embeddings

Machine Learning

Titan Text Embedding models are built using machine learning. Specifically, TTE models are neural network models built on a transformer architecture. Our runtime service architecture works as follows: 1/ TTE receives a text string via the API; and 2/ TTE converts the text into a numerical vector.

Controllability

Our primary control levers for impacting the utility of the TTE models are the unlabeled pre-training data corpus and the contrastive learning corpus. Our development process exercises these control levers as follows: 1/ We pre-train TTE using curated data from a variety of sources, which may include licensed and proprietary data, open-source datasets, and publicly available data where appropriate. By constructing a well-balanced corpus with varied vocabulary and contexts, the model can learn generalizable linguistic patterns. 2/ During the contrastive learning stage, we shape the embedding space so that the embeddings of semantically similar or dissimilar texts are correspondingly nearer or further apart.

Performance Expectations

Customer applications differ in the kinds of text processed and in the degree of similarity or dissimilarity expected between text strings, implying that text embedding model performance is likely to vary across applications. Consider two applications: Application A uses TTE to classify online customer reviews into sentiment categories. With this application, inputs often come from diverse customers, featuring varied language, tone, and informal phrasing. Consequently, Application A needs to handle diverse vocabulary, inconsistent grammar, and occasionally ambiguous sentiment expressions. Application B, on the other hand, leverages TTE to group employee feedback into topic categories. Here, inputs tend to be less linguistically diverse but contain context-dependent terminology unique to internal processes and workplace culture. Application B must cope with specialized jargon, subtleties in feedback tone, and occasional ambiguities in topic boundaries. Because performance results depend on a variety of factors including TTE, the deployment workflow, and the evaluation dataset, we recommend that customers test TTE using their own content.

Test-driven Methodology

We use multiple datasets to evaluate the performance of TTE. No single evaluation dataset provides an absolute picture of performance. This is because evaluation datasets vary based on

task, language, intrinsic and confounding variation, the types and quality of labels available, and other factors. Our development testing involves automated benchmarking against publicly available datasets (see below), automated benchmarking against proprietary datasets, benchmarking against proxies for anticipated customer use cases and more. Our development process examines TTE performance using all these tests, and takes steps to improve the model and/or the suite of evaluation datasets. In this AI Service Card, we provide examples of test results to illustrate our methodology. Customers should perform their own testing on datasets specific to their own use cases.

Fairness

Text embedding models could inherit biases from their training data, leading to unfair associations, such as linking certain groups with stereotypes. For example, a model might more frequently associate "engineer" with "men" and "nurse" with "women", reinforcing gender stereotypes in professions. These biases may impact downstream applications, potentially leading to unfair treatment of certain groups. We design TTE models to work for a diverse set of customers. We examine the extent to which TTE exhibits stereotypical biases. For example, we use the StereoSet dataset to measure stereotypes across categories of gender, race, profession, and religion. In this dataset, each context sentence is paired with a stereotypical sentence and an anti-stereotypical sentence, with each context associated with a bias type. To quantify bias, we calculate the absolute difference in association between the context and its stereotypical and anti-stereotypical sentences. The bias score ranges from 0 to 2. A score close to 0 indicates that the model does not strongly associate contexts with either stereotypical or anti-stereotypical sentences, while a maximum score of 2 indicates the model consistently makes such associations. The highest bias (higher score indicates more bias) score across all 4 groups is 0.13 for TTE G1 with its 1536-dimensional embedding, and 0.09, 0.10, and 0.10 for TTE v2 with its 1024-, 512-, and 256-dimensional embeddings respectively.

We also use the TREC 2022 Fair Ranking Track dataset to measure the retrieval fairness performance for TTE. This dataset tests embedding models on fair ranking of Wikimedia articles. For a given query (WikiProject Topic), the task is to produce a list of 500 ranked documents that are relevant to the query and provide a fair exposure to articles that are associated with 8 attributes such as geographic location (article topic), gender (biographies only), occupation (biographies only), and age of the article. T he test set contains 46 unique topics in total. We measure retrieval fairness using the Attention Weighted Ranking Fairness (AWRF) metric, which compares the exposure of groups in a ranked list with a target fair distribution, and which ranges from 0 to 1, with 1 indicating perfect fairness. On this dataset, the lowest fairness score (lower score indicates more bias)across all 8 groups is 0.88 for TTE G1

with its 1536-dimensional embedding, and 0.89, 0.90, and 0.90 for TTE v2 with its 1024-, 512-, and 256-dimensional embeddings respectively.

Explainability

Customers wanting to check whether similar texts are positioned closely in the TTE embedding space can measure cosine similarities between the output vectors. For customers wanting to verify the effectiveness of the TTE in a RAG use case, we recommend using RAG with Amazon Bedrock knowledge bases, where end users can see attribution of information in a completion.

Robustness

TTE models are designed to perform consistently well across a wide range of use cases. We evaluate the overall embedding performance using the MTEB (Massive Text Embedding Benchmark) benchmark. This benchmark includes 58 datasets covering 112 languages and 8 embedding tasks: bitext mining, classification, clustering, pair classification, reranking, retrieval, semantic textual similarity, and summarization. Each task is evaluated using specific accuracy metrics suited to its purpose. To assess the retrieval effectiveness of TTE, we use the Normalized Discounted Cumulative Gain at 10 (NDCG@10) metric. NDCG@10 measures the quality of the top 10 ranked items retrieved for a given query. It considers both the position and the relevance of each item in the list. The metric ranges from 0 to 1, where 1 indicates a perfect ranking (most relevant items at the top). A higher score suggests better alignment between the model's ranking and an ideal ranking where the most relevant items are positioned first. On MTEB, TTE G1 with 1536 dimensions scores 0.59 overall (average accuracy across all datasets) and 0.47 on retrieval (average NDCG@10 across retrieval datasets only), while TTE V2 with 1024 dimensions scores 0.60 overall and 0.51 on retrieval. We additionally evaluate the binary embeddings and reduce dimension features of V2 on the retrieval tasks within MTEB. Binary embeddings retain 98.5% of the full-precision retrieval performance (when using the standard method of reranking with the full-precision query). Reducing dimension size to 512 and 256 retain 99.0% and 96.8% of the 1024-dimension retrieval performance, respectively.

Privacy

TTE is available in Amazon Bedrock. Amazon Bedrock is a managed service and does not store or review customer prompts or customer prompt completions (including the embeddings returned by TTE), and the model input and output are never shared between customers, or with Amazon Bedrock partners. AWS does not use inputs or outputs generated through the Amazon Bedrock service to train Amazon Bedrock models, including TTE. For more information, see Section 50.3 of the AWS Service Terms and the AWS Data Privacy FAQs. For service-specific privacy information, see Security in the Amazon Bedrock FAQs

Security

All Amazon Bedrock models, including TTE, come with enterprise security that enables customers to build generative AI applications that support common data security and compliance standards, including GDPR and HIPAA. Customers can use AWS PrivateLink to establish private connectivity between customized Titan models and on-premises networks without exposing customer traffic to the internet. Customer data is always encrypted in transit and at rest, and customers can use their own keys to encrypt the data, for example, using AWS Key Management Service (AWS KMS). Customers can use AWS Identity and Access Management (IAM) to securely control access to Amazon Bedrock resources. Also, Amazon Bedrock offers comprehensive monitoring and logging capabilities that can support customer governance and audit requirements. For example, Amazon CloudWatch; can help track usage metrics that are required for audit purposes, and AWS CloudTrail can help monitor API activity and troubleshoot issues as TTE is integrated with other AWS systems. Customers can also choose to store the metadata, prompts, and completions in their own encrypted Amazon Simple Storage Service (Amazon S3) bucket. For more information, see [Security](#).

Transparency

TTE provides information to customers in the following locations: this Service Card, AWS user documentation, AWS educational channels (e.g., blogs, developer classes), and the AWS Console. We accept feedback via the AWS Console and through traditional customer support mechanisms such as account managers. Where appropriate for their use case, customers who incorporate TTE in their workflow should consider disclosing their use of machine learning to end users and other individuals impacted by the application, and customers should give their end users the ability to provide feedback to improve workflows. In their documentation, customers can also reference this AI Service Card.

Governance

We have rigorous methodologies to build our AWS AI services responsibly, including a working backwards product development process that incorporates Responsible AI at the design phase, design consultations, and implementation assessments by dedicated Responsible AI science and data experts, routine testing, reviews with customers, best practice development, dissemination, and training.

Deployment and performance optimization best practices

We encourage customers to build and operate their applications responsibly, as described in [AWS Responsible Use of AI Guide](#). This includes implementing Responsible AI practices to address key dimensions including controllability, safety, fairness, veracity, robustness, explainability, privacy, security, transparency, and governance.

Workflow Design

The performance of any application using TTE depends on the design of the customer workflow, including the factors discussed below:

1. **Effectiveness Criteria:** Customers should define and enforce criteria for the kinds of use cases they will implement, and, for each use case, further define criteria for the inputs and outputs permitted, and for how humans should employ their own judgment to determine final results. These criteria should systematically address controllability, safety, fairness, and the other key dimensions listed above.

2. **Configuration:** TTE provides configuration options on embedding dimensions, the embedding types (float or binary), and whether to normalize the embeddings. Customers should consider which parameter choices will provide the most effective results for their specific use case. For more information, see [TTE model parameters](#) in the *Amazon Bedrock User Guide*.

3. **Human Oversight:** If a customer's application workflow involves a high risk or sensitive use case, such as a decision that impacts an individual's rights or access to essential services, human review should be incorporated into the application workflow where appropriate.

4. **Performance Drift:** A change in the types of text input that a customer submits to TTE, or a change to the service, may lead to different outputs. To address these changes, customers should consider periodically retesting the performance of TTE, adjusting their workflow if necessary.

5. **Updates:** We will notify customers when we release a new version, and will provide customers time to migrate from an old version to the new one. Customers should consider retesting the performance of updated TTE models on their use cases.

Further information

- For service documentation, , see [Amazon Titan Text Embedding models](#).

- For details on privacy and other legal considerations, see the following AWS policies: [Acceptable Use](#), [Responsible AI](#), [Legal](#), [Compliance](#), and [Privacy](#).

- For help optimizing workflows, see [Generative AI Innovation Center](#), [AWS Customer Support](#), [AWS Professional Services](#), [Ground Truth Plus](#), and [Amazon Augmented AI](#).

- If you have any questions or feedback about AWS AI service cards, please complete [this form](#).

Glossary

Controllability: Steering and monitoring AI system behavior.

Privacy & Security: Appropriately obtaining, using and protecting data and models.

Safety: Preventing harmful system output and misuse.

Fairness: Considering impacts on different groups of stakeholders.

Explainability: Understanding and evaluating system outputs.

Veracity & Robustness: Achieving correct system outputs, even with unexpected or adversarial inputs.

Transparency: Enabling stakeholders to make informed choices about their engagement with an AI system.

Governance: Incorporating best practices into the AI supply chain, including providers and deployers.