

ServiceNow Text-to-Text LLM Model Card

ServiceNow Text-to-Text LLM Model Card

Intended Use and Functionality

Purpose of the Model

The Model supports ServiceNow users by enabling text-to-text Now Assist skills within application workflows on the ServiceNow Platform.

It is designed to enhance automation by facilitating assistive text generation, including question answering and summarization, helping users generate structured content efficiently while maintaining alignment with enterprise requirements.

Autonomy Level

The Model operates at an Assistive autonomy level, providing AI-generated text suggestions that require human review before implementation.

Users must validate AI-generated content to ensure its accuracy and appropriateness for the intended use case. This human-in-the-loop approach ensures that users retain full control over the output, preventing fully autonomous decision-making.

Optimization Scope and Limitations

The Model has been fine-tuned with open-source data and ServiceNow Platform-specific data to optimize performance for enterprise applications.

It supports text-to-text processing tasks such as summarization, content generation, and structured automation within workflows. Additionally, it is based on the [Mixtral-8x7B](#) large language model, incorporating enhancements tailored to ServiceNow use cases.

Model Name

ServiceNow Text-to-Text LLM ("the Model")

Model Version

v1.0

Model Release Date

August 2024

Model Distribution Method

ServiceNow Platform

Model License

Please refer to your agreement with ServiceNow for all license information.

Products Using this Model

Now Assist skills within application workflows on the ServiceNow Platform

User Benefits

The Model is fine-tuned for several natural language processing (NLP) tasks including intent classification, summarization, and question answering (Q&A).

Risks

The Model may generate text summaries that may be inaccurate, omit key information, or include irrelevant or redundant text. The Model may, at times, produce undesirable text, even if the prompt itself does not include anything explicitly offensive. For more information, refer to section 3.3 of the [Mixtral paper](#).

See Section 3.3 of the [Mixtral paper](#) for a discussion of how Mixtral safety testing was conducted.

Factors and limitations

- The Model has been primarily trained on English data. While the model may respond to non-English text, the responses are unlikely to be correct or consistent.
- The Model needs sufficient context in a prompt to create an acceptable response. If the prompt is too generic, the Model will not be able to follow the instruction.

Ethical considerations

The Model has been fine-tuned with the intention of reducing bias, toxicity, and hallucinations, although such limitations may still exist.

Text LLMs can produce harmful text based on how they are prompted, and the Model is not free from such limitations similar to the behavior of other industry LLMs.

When using the model customers should follow ServiceNow's guidelines on intended use available on docs.servicenow.com as well as ServiceNow's [AI Acceptable Use Policy](#).

Please report instances of unintended hallucinations, harmful text (e.g. toxicity, profanity, etc.), or unexpected PII occurrences in the model output so that we can evaluate for remediation.

Supported Languages

Primary Language: English

Multilingual Capabilities: The model may respond to non-English inputs, but accuracy is not guaranteed.

Model Architecture

Base Model: [Mixtral-8x7B](#)

Fine-Tuning: The Model is finetuned with next token prediction on open-source and ServiceNow proprietary datasets.

Number of Parameters

7B

Maximum Input and Output Size

- **Maximum Input Size:** 32,000 tokens (shared between input and output)
- **Maximum Output Size:** 4,096 tokens
- **Shared Context Window:** 32,000 tokens

Input and Output Modalities

Modality Type

Single-modality: The Model processes and generates **text-to-text** outputs.

Inputs

Input Type: The Model accepts **plain text** input, typically structured as questions, commands, or prompts in natural language (e.g., "Summarize this article" or "Translate this sentence to French").

Input Constraints:

- The Model is designed primarily for natural language inputs and may perform sub-optimally with **non-text inputs** such as raw numerical data or unstructured code.
- Some **long-form inputs** may be truncated based on token limits, affecting response completeness.
- The Model may struggle with **highly ambiguous or domain-specific jargon** if not adequately trained in those areas.

Outputs

Output Type: The Model generates **plain text responses** based on the input prompt. This can include answers to questions, summarizations, translations, code snippets, or structured text.

Output Constraints: Outputs are constrained by **token length limits**, which may truncate longer responses.

Input and Output Formats

- **Input Format:** Inputs must be formatted as **plain text** with no additional structuring required.
- **Output Format:** Outputs are generated as **plain text** with no additional structuring unless explicitly prompted.

Training Data

ServiceNow did not conduct further pre-training.

Fine-Tuning Data

- The Model has been further tuned by instruction fine tuning for ServiceNow tasks using both opensource datasets and ServiceNow platform-specific text data.
- The platform data is synthetic, which represents IT, CSM, and HRSD domains.
- Data anonymization is conducted before use.
- The model has also undergone further alignment designed to improve helpfulness and harmlessness.

Evaluation Data

- The Model has been evaluated on a variety of public benchmarks and ServiceNow skill-specific datasets.
- The Model has also been evaluated on skill-specific synthetic datasets, which represent IT, CSM and HRSD domains.

Metrics

The Model was evaluated with the following metrics: F1, Rouge, BERT Score, Faithfulness, Correctness (Human eval)

Technical Means for Integration

All interactions and processing occur within the platform's secure architecture.