

2023

ServiceNow text-to-flow LLM Model Card

ServiceNow text-to-flow LLM Model Card

Intended Use

The ServiceNow text-to-flow LLM Model ("the Model") is designed to support ServiceNow users to generate flows from text inputs in Flow Designer on the ServiceNow Platform.

The Model can generate or complete flows according to the intent expressed in a user's prompt. Users manually initiate a flow-generation request with a prompt, then review and edit the AI-generated flow as needed for their use case before accepting and using the generated flow. It is important to have a human reviewing the AI-generated flow (i.e. human-in-the-loop) to ensure the flow is appropriate and functions correctly for the intended use case.

The Model is based on a model in the [StarCoder](#) large language model family, and has been further fine-tuned on ServiceNow Platform-specific flow data.

The license for the StarCoder family of models includes explicit usage restrictions that preclude, among other things, the generation or distribution of malware, the generation or dissemination of personally identifiable information (PII) to harm others, any use that violates applicable laws, any use that results in automated decision making that adversely affects individuals' rights, and any use that is intended to discriminate on the basis of legally protected characteristics.

See Attachment A of the [Open RAIL-M v1](#) license for the license for the StarCoder family of models and a list of usage restrictions.

Model details

Developed by: ServiceNow
Research and Engineering
Released: Nov. 2023, v1.0

Evaluation data

ServiceNow Platform-specific data was used for text-to-flow use-case specific evaluation.

Metrics

ServiceNow internal metrics

Model Architecture

The Model is a finetuned version of a model in the StarCoder model family.

Training data

The Model is based on a model in the StarCoder model family, which was trained on data from publicly available sources.

Fine-tuning was done with ServiceNow Platform-specific flow data.

User Benefits

The Model generates flows to help users reduce the time required to build them, especially for novice users.

The Model can help build more efficient and reusable flows as it suggests flow structures and components following "best practices".

Risks

The Model may generate flows that are inefficient, or that do not fit a person's preferences or that may be deemed simplistic. For more information, refer to section 10 of the [StarCoder paper](#).

Factors and limitations

The Model needs a prompt that is explicit enough to create an acceptable flow.

The Model currently does not support filling inputs, outputs, and data pills for the generated flows.

The Model may struggle with generating custom actions and sub-flows.

The Model may struggle with complex and ambiguous prompts. The Model relies on its ability to understand the context of the flow and semantics of the prompt from the user.

Ethical considerations

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

Some PII and malicious code was removed in training data for the StarCoder family of models, but some may still exist. Additional PII and malware detection in pre- and post-processing can help mitigate this issue. Code LLMs can produce harmful code based on how it is prompted, and the Model is not free from such limitations similar to the behavior of other industry LLMs. Developers should not use the Model to generate malware or in any other way prohibited by the usage restrictions for the StarCoder family of models.

Please report instances of hallucinations, malicious text, or PII in output so that we can evaluate for remediation.

