

Flowstorm: Open-Source Platform with Hybrid Dialogue Architecture

Jan Pichl,¹² Petr Marek,¹² Jakub Konrad¹² Petr Lorenc¹²
Onřej Kobza¹ Tomas Zajicek² Jan ředivy³²

¹ Faculty of Electrical Engineering, Czech Technical University, Prague, Czech Republic

² PromethistAI, Prague, Czech Republic

³ CIIRC, Czech Technical University, Prague, Czech Republic

{pichljan, marekp17, konrajak, lorenpe2, kobzaond}@fel.cvut.cz,

{tomas.zajicek, jan.sedivy}@promethist.ai

Abstract

This paper presents a conversational AI platform called Flowstorm. Flowstorm is an open-source SaaS project suitable for creating, running, and analyzing conversational applications. Thanks to the fast and fully automated build process, the dialogues created within the platform can be executed in seconds. Furthermore, we propose a novel dialogue architecture that uses a combination of tree structures with generative models. The tree structures are also used for training NLU models suitable for specific dialogue scenarios. However, the generative models are globally used across applications and extend the functionality of the dialogue trees. Moreover, the platform functionality benefits from out-of-the-box components, such as the one responsible for extracting data from utterances or working with crawled data. Additionally, it can be extended using a custom code directly in the platform. One of the essential features of the platform is the possibility to reuse the created assets across applications. There is a library of prepared assets where each developer can contribute. All of the features are available through a user-friendly visual editor.

1 Introduction

With the increasing popularity of conversational systems in various areas, there is an increasing demand for human-like aspects of the systems. The human-like aspects cover not only the content of the conversation but also the conversational style and tone of voice. Moreover, many conversational constructs are applicable to various situations and are therefore suitable for sharing and reusing. We introduce the Flowstorm platform¹, where conversational designers can easily design an application using shared assets created by a community, create an application on their own, or use a combination of both approaches. The platform does not require any installation as it can be used as a web app.

¹<https://app.flowstorm.ai>

However, advanced developers can use the open-source code² to run and modify their own version of the platform.

Our platform works with a novel dialogue architecture we have proposed. It uses a combination of dialogue trees/graphs (scenarios) specific to a domain (or only a part of the domain) and generative models handling the aspects of the conversation that are not domain-specific. The Flowstorm platform targets users with limited coding skills and conversational AI knowledge as well as the experts in the area who can easily extend an out-of-the-box functionality with custom logic. Using the provided tools, documentation³, and library of the shared assets, one can create a dialogue application in minutes. The Alquist (Konrad et al., 2021) socialbot, which was the first place winner in the Alexa Prize Socialbot Grand Challenge 4 (Hu et al., 2021), was created using the Flowstorm platform, proving the platform to be suitable even for such a complex system.

Flowstorm platform supports English, Czech, German and Spanish languages.

2 Related Work

There are two categories of related software: open-source frameworks or libraries and online platforms. Rasa (Bocklisch et al., 2017) and DeepPavlov (Burtsev et al., 2018) are the most famous representatives of the framework category. They provide the developer with a set of programmatic and declarative interfaces for defining and training natural language understanding (NLU) and dialogue management (DM) models. They also allow creating a pipeline of the components where the developer can create a sequence of models or components for utterance processing. Both libraries require a local installation and, in addition, inte-

²<https://gitlab.com/promethistai/flowstorm>

³<https://docs.flowstorm.ai/>

grating either of them with the rest of the conversational application requires considerable effort.

Alexa Skills Kit (ASK)⁴, Google Dialogflow⁵, and Voiceflow⁶ are representatives of the category of online platforms. ASK and Dialogflow offer an online application for the definition of NLU models (intents and slots/entities). For dialogue management, ASK introduced Alexa Conversations (Acharya et al., 2021) allowing the developer to specify the conversation flow using sample conversations. Dialogflow CX (a part of Dialogflow providing a new way of working with conversational agents) offers a visual view of the conversation agents to see the possible flows of the conversation. Voiceflow uses a visual dialogue editor as an approach to building a conversational application. Dialogue blocks can be connected via transitions to create a flow of blocks where each block represents a specific functionality such as keyword matching or API calls.

The comparison of the key aspects of the described libraries and platforms is shown in Table 1.

	Flowstorm	Rasa	DeepPavlov	Amazon ASK	Google Dialogflow	Voiceflow
SaaS	✓		✓	✓	✓	
Open-source	✓	✓	✓			<i>p</i>
Visual editor	✓				<i>p</i>	<i>p</i>
Custom code	✓	✓				
ML-based NLU	✓	✓	✓	✓	✓	✓
Sub-dialogues	✓					
Sharing assets	✓					✓
Built-in NRG	✓					
Analytic tools	✓		✓	✓	✓	

Table 1: Comparison of frameworks and online conversational platforms, *p* means partial functionality.

3 Architecture

The Flowstorm platform combines two sets of functions—design and analytics. The design functions allow users to create individual sub-dialogues, train entity models, write shared code, and combine the sub-dialogues into complex conversational applications. Each of the created assets can then be shared with the community. The analytic functions

⁴<https://developer.amazon.com/alexa/alexa-skills-kit>

⁵<https://cloud.google.com/dialogflow>

⁶<https://www.voiceflow.com/>

can be used to inspect the applications’ traffic and show basic metrics such as session transcripts or visualizations of specific parameters.

The following sections describe a novel conversational system architecture. The novel aspect consists of a combination of manually created sub-dialogues (that can be easily combined into a complex dialogue structure) and neural generative models. First, we describe the blocks essential for creating an application, and then we describe each component used in runtime in detail.

3.1 Conversational Application

Conversational application is a set of dialogue trees we call sub-dialogues. Each sub-dialogue is focused on a small subset of a domain (e.g., in the movie domain, one sub-dialogue can be focused on a user’s favorite movie). Each application has at least one sub-dialogue, but typically consists of several sub-dialogues. The sub-dialogue which is triggered when the application is launched is called “main”. The other sub-dialogues are triggered when a specific situation occurs during the conversation. It depends on the specific design of the application, and it is described in detail in the following subsection.

Additionally, each application has configurable parameters such as language, Skimmer rules, or voice.

3.2 Sub-dialogue

Sub-dialogue is a graph structure consisting of connected nodes representing the flow of the conversation. The key node types are described in the following list. The basic structure and node types are shown in Figure 1.

Enter – Entry point of the sub-dialogue. Each sub-dialogue must have exactly one Enter.

Speech – Speech node contains one or more natural language responses that are presented to a user. If there are multiple responses, one is randomly selected. It can also contain slots that are eventually filled with variables based on the context. The responses can be customized using Kotlin-based DSL language.

User Input – Point in a conversation where the bot waits for the user utterance. It is typically connected to multiple intent nodes. Each User Input node has its own underlying intent recognition model trained separately. It has as many intent classes as there are intent nodes connected

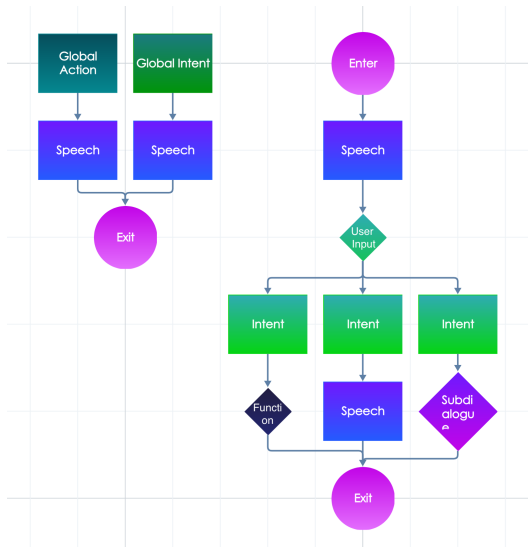


Figure 1: Nodes connected to the sub-dialogue structure. The names of the node types are shown in the corresponding boxes in the picture.

to it. The default behavior can be customized using Kotlin-based DSL language.

(Global) Intent – Intent node contains examples of the users’ possible utterances which serve as training data for the intent recognition model. Each intent represents a single class in the intent recognition model. There are two types of intent nodes: contextual Intent and Global Intent. Contextual intents can be recognized only at a specific point of the conversation—when the conversation flow reaches the User Input node they are connected to. Global Intents, however, can be recognized at any point of the conversation.

Function – Function node contains code written in the Kotlin language (JetBrains, 2011). The code can contain arbitrary logic (usually working with attributes — see the attributes subsection). The function returns the transition to the next dialogue node. It is suitable for branching the dialogue based on the attribute values, API results, etc.

(Global) Action – Actions represent specific situations in the conversation flow. There are several predefined situations: *Silence*, *Error* and *Out of domain*. *Silence* action is triggered when no speech is recognized. *Error* action is triggered when the logic written in the Function nodes fails (connection errors or bad design). *Out of domain* action represents the utterance that does not belong to any of the intent classes. This will be described in detail in the following sections.

Sub-dialogue – Sub-dialogue node introduces the ability to reference another sub-dialogue. When

the conversation flow reaches this node, it triggers the referenced sub-dialogue, processing its logic from the Enter node to the Exit.

Exit – Exit node represents the end of the sub-dialogue flow. If the sub-dialogue is the “main” dialogue, the session is ended. If it is a lower-level sub-dialogue, the conversation flow jumps to the parent sub-dialogue (i.e., the one where the current sub-dialogue was referenced).

Init code – Each sub-dialogue has a declarative code part that typically serves as a place to define attributes (see below) and methods that can be called from the function nodes. As the code is written in the Kotlin language, the type control is done in the build phase, which lowers runtime errors.

Attributes – There are four scopes of the attributes that can be defined in the Init code. The attribute stores values that can be used to modify the conversation flow, or they can be presented directly as part of a response. Each attribute must have its default value which is used to initialize the attribute based on its scope.

1. *Turn* – the value is reset to default at the beginning of each turn.
2. *Session* – the value is reset to default at the beginning of each session.
3. *User* – the default value is used for each new user. Once the value of a user attribute is set, it is valid for all sessions of the user.
4. *Community* – the default value is used for each community namespace. Once the value of a community attribute is set, it is valid for all users using the same community namespace.

4 System Components

The Flowstorm conversational platform consists of several components crucial for processing the user utterance and driving the dialogue. First, NLU gathers the information required by the DM. Afterward, DM traverses the corresponding sub-dialogue tree structure based on the NLU output.

4.1 Natural Language Understanding

The NLU consists of the main components—entity recognition and intent recognition—that are commonly used in conversational systems. Additionally, the out-of-domain detection enhances the intent recognition, allowing the system to handle a situation that was not expected during dialogue design.

4.1.1 Entity Recognition

There are two modules used for Entity Recognition. Regex-based tool Duckling⁷ and a sequence labeling neural network. The neural network uses the Bi-LSTM-CRF architecture (Huang et al., 2015). The architecture is simple enough to provide reasonable response time when run on CPU. Duckling is used to recognize such entities that have a specific structure and can be easily described using regular expressions (date, time, numbers, URLs, amount of money, etc.). Additionally, it parses (not only) the date and time values from natural language representation into a structured format. We enhanced the Duckling component with rules for numbers and time values in the Czech language.

During the process of dialogue creation, one can define custom entity types that can be recognized in the conversation. The entities are defined using example sentences where they can occur, e.g. `My favorite movie is [Matrix]{movie}`, meaning the word `Matrix` is an entity with type `movie`. These examples are used as training data for the sequence labeling model.

4.1.2 Entity Masking

The recognized entities are masked out in the user utterance prior to the intent recognition phase. This step allows the intent recognition model to focus only on the types of entities rather than the actual values. The intent recognition module receives, for example, the sentence `My favorite movie is {movie}` instead of `My favorite movie is Matrix`. Only the entity types included in the intent examples (in the intent nodes) are masked out.

4.1.3 Intent Recognition

A key part of each sub-dialogue are the intent nodes. Each sub-dialogue may contain multiple global intents and multiple user input nodes with various numbers of local intents connected to it. Let the $\mathcal{G} = (g_1, \dots, g_n)$ be a set of global intents of the sub-dialogue S and let the $\mathcal{U} = (u_1, \dots, u_m)$ be the set of User Input nodes where each node u_i has its own set of local intents $\mathcal{L}_i = (l_1, \dots, l_{k_i})$. There is one model trained for each $u_i \in \mathcal{U}$ that classifies the user utterance into $|\mathcal{L}_i|$ classes. Additionally, there is one extra model trained for global intents which has $|\mathcal{G}|$ classes.

⁷<https://github.com/facebook/duckling>

When the conversation flow reaches the User Input node, the system starts listening to the user (or waits for a text input). Then the actual intent recognition process consists of two steps:

First, the user utterance is embedded using the Universal Sentence Encoder (Cer et al., 2018) and the embedding is compared with the training examples from both global and local intent classes. This step only decides whether the final intent will be selected from local intents, global intents or whether the utterance is out of domain given the context of the current User Input node (see Subsection *Out of Domain Detection* and Figure 2). Note that not only global intents from the current sub-dialogue are considered but also the global intents from all parent sub-dialogues which are on the path from the current sub-dialogue to the main dialogue.

Second, the corresponding model (either global intent model or a model corresponding to the current User Input node) is selected. This model classifies the utterance into the final class.

The intent recognition experiments and a comparison with other NLU tools is described in the paper by Lorenc et al. (2021).

4.1.4 Out of Domain Detection

During each dialogue turn, the utterance is classified into one of the intent classes. However, the utterance does not have to correspond to either of the classes. We call such utterances out of domain (OOD). Whether the utterance is OOD or not depends purely on the sample utterances in the global intents and the corresponding local intent classes given in the current User Input node. The OOD is triggered using the confidences of the corresponding intent classes and the empirically estimated threshold. When the OOD is recognized, the dialogue flow continues with a local OOD action (if connected to the current User input node) or with a global OOD action.

4.2 Skimmer

The Skimmer is a component that extracts relevant data from the user utterance and stores it for later usage. It is done on a turn basis regardless of the dialogue context. An example of such information is: the information about the user having a brother mentioned “by the way” in the conversation (e.g., in a movie-related conversation, the user may mention “*I went to the cinema with my brother yesterday.*”). The component skims through each utterance and saves the values into the attributes.

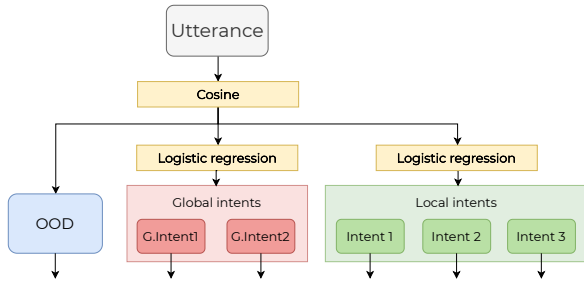


Figure 2: Classification algorithm for intent and OOD detection. The utterance is first classified by a cosine similarity into a class of local intents or a class of global intents. Next, the corresponding logistic regression makes the final intent classification. Additionally, cosine similarity can predict the OOD class if the similarity score falls below the threshold.

The dialogue creator can then use the information in the conversation. The Skimmer is defined by a set of rules containing the following attributes:

- *Regular expression* - a set of patterns which must be contained in the utterance.
- *Attribute name* - the name of the attribute where the value will be stored.
- *Value* - the value stored in the attribute, typically *true*, *false*, or a matched group of the regular expression.

This component is inspired by a similar functionality used in the Alquist socialbot (Konrád et al., 2021).

4.3 Dialogue Management

Dialogue management (DM) of a conversational application created in Flowstorm is divided into two stages: *Sub-dialogue DM* and *Dialogue selection*. The former stage operates on the graph structure of the corresponding sub-dialogue. The DM simply follows the directed edges beginning in the *Start* node. There are two types of nodes that may have multiple outgoing edges: User Inputs and Functions. In the User Input nodes, the next edge is selected using Intent Recognition results, i.e., the edge leading to the most probable intent or global intent is selected. The Function nodes contain code that returns one of the connected edges.

4.3.1 Dialogue selector

The dialogue selector stage is triggered whenever a previous sub-dialogue is ended, and there is not a direct transition to a different one. The overall

goal is to select a sub-dialogue that is related to the context of the conversation. The most suitable sub-dialogue is selected from a list of eligible sub-dialogues, which is defined by the developer who creates the dialogue application. Each sub-dialogue has several properties that are considered during the process of the dialogue selection.

- *Label* – one or more labels can be assigned to each sub-dialogue.
- *Entity* – a sub-dialogue can use the information about an entity and work with the entity attributes. E.g., a sub-dialogue works with the information about a person, including the information about their birth date, occupation, etc. In that case, the sub-dialogue is tagged with the entity type Person.
- *Starting conditions* is a custom logic that tells whether the corresponding sub-dialogue can be selected given the current context. E.g., in the case of the sub-dialogue discussing the user’s favorite movie, it needs to check whether the information about the favorite movie is already stored. Additionally, most of the sub-dialogue can be triggered only once per session. Hence, the starting condition typically checks that as well.

The intuition behind the dialogue selector is to choose a dialogue whose starting condition is met, and its label and entity sets have the biggest overlap with the label and entity sets currently being discussed in the session. Formally, let \mathcal{D} be a set of eligible sub-dialogues, S be the current session. $\text{label}(d, S)$ returns the overlapping labels between the dialogue d and session S , and, analogically, $\text{entity}(d, S)$ returns overlapping entities, and \mathcal{C} is a set of starting conditions. Then the dialogue selector selects a sub-dialogue as follows:

$$\arg \max_{d \in \mathcal{D}} |\text{label}(d, S) \cup \text{entity}(d, S)|$$

$$s.t. \quad \forall \text{cond}_i \in \mathcal{C} : \text{cond}_i(d, S)$$

4.4 Neural Response Generation

All the principles described in the previous sections are suitable for creating high-quality content for various scenarios with maximum control over the conversation flow. A natural language conversation can, however, easily deviate from these scenarios. The Flowstorm platform comes with a novel approach using a combination of manually created

content and generated content for this type of situation. During the dialogue creation process, the user may plug the generated response into any part of the sub-dialogue.

The platform currently uses a GPT-2-based (Radford et al., 2019) architecture for generating responses. The model takes a sequence of user utterances and bot responses as input. Additionally, the input contains the information about the dialogue act of the desired response. Since the model expects only natural language as the input, it is suitable for situations where the next response cannot be determined using only the tree-like dialogue structure.

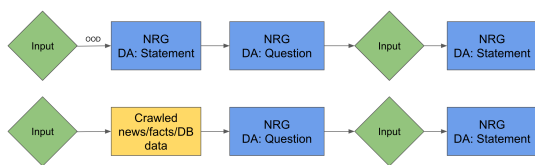


Figure 3: Two basic NRG usage scenarios. The upper part uses NRG response when OOD is detected. NRG generates a statement and a question, and then the system waits for the user’s utterance. Afterward, another NRG statement is used, and the flow is returned to the dialogue graph. The lower part illustrates the scenario where the crawled text or information from a DB is used. A single NRG generated question follows it. The next turn starts with the NRG statement, and the flow is returned to the dialogue graph.

A typical use case for a generative model is to use it whenever an out-of-domain utterance is recognized. Since the dialogue designer cannot easily predict all out-of-domain variants and design a proper reaction, the generative approach is a suitable solution as it operates purely with the dialogue history.

In Figure 3, we show two typical use cases of the generative model—generating response after an out-of-domain utterance is detected and asking additional questions based on free text (news article, fun-fact, etc.).

5 Analytics

The set of analytic tools allows developers to inspect the traffic on their conversational applications. There are tools for inspecting session transcripts, community and user attributes, and creating metric visualizations. The *Session transcripts* contains a list of the dialogue turns of each session with all the annotations and log messages. One can easily

watch the NLU results, identify a weak part of the corresponding application, inspect the Automatic Speech Recognition (ASR) hypotheses, duration of each turn, and session attributes stored during the conversation. The community and user attributes can be viewed in individual tables to watch the values gathered during conversations.

The metric visualization tool allows creating plots of values defined as a metric. One can easily filter the specific values, time ranges, applications, users and specify the granularity of the visualization. An example visualization is shown in Figure 4.

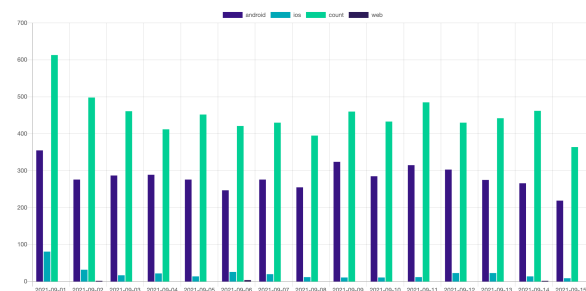


Figure 4: Example metric visualization. The columns show the count of sessions per client (Android, iOS, web and total).

6 Conclusion

The open-source platform Flowstorm allows users to easily create conversational applications with low effort while allowing them to modify each part with custom functionality. Developers can visually design an application and extend it with custom code thanks to an easy-to-use interface. A library of shared assets contains common dialogue structures that can be used in various use cases. A novel hybrid architecture allows the system to handle unexpected utterances and make the conversation more robust. Analytic tools show the conversation transcripts, attributes, and metrics directly in the Flowstorm app. Based on the analytic data, the developer can quickly identify the weak parts of the application and fix them by modifying the dialogue structure or intent examples.

The Flowstorm platform is in active development, and we are primarily focusing on the generative part of the dialogue architecture to make it more suitable for various dialogue situations.

References

- Anish Acharya, Suranjit Adhikari, Sanchit Agarwal, Vincent Auvray, Nehal Belgamwar, Arijit Biswas, Shubhra Chandra, Tagyoung Chung, Maryam Fazl-Zarandi, Raefer Gabriel, et al. 2021. *Alexa conversations: An extensible data-driven approach for building task-oriented dialogue systems*. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Demonstrations*, pages 125–132, Online. Association for Computational Linguistics.
- Tom Bocklisch, Joey Faulkner, Nick Pawlowski, and Alan Nichol. 2017. *Rasa: Open source language understanding and dialogue management*.
- Mikhail Burtsev, Alexander Seliverstov, Rafael Airapetyan, Mikhail Arkhipov, Dilyara Baymurzina, Nickolay Bushkov, Olga Gureenkova, Taras Khakhulin, Yurii Kuratov, Denis Kuznetsov, et al. 2018. *Deppavlov: Open-source library for dialogue systems*. In *Proceedings of ACL 2018, System Demonstrations*, pages 122–127.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Céspedes, Steve Yuan, Chris Tar, et al. 2018. *Universal sentence encoder*. *arXiv preprint arXiv:1803.11175*.
- Shui Hu, Yang Liu, Anna Gottardi, Behnam Hedayatnia, Anju Khatri, Anjali Chadha, Qinlang Chen, Pankaj Rajan, Ali Binici, et al. 2021. *Further advances in open domain dialog systems in the fourth alexa prize socialbot grand challenge*. In *4th Proceedings of Alexa Prize (Alexa Prize 2021)*.
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. *Bidirectional lstm-crf models for sequence tagging*. *arXiv preprint arXiv:1508.01991*.
- JetBrains. 2011. *Kotlin programming language*. <https://kotlinlang.org/>. Accessed on 14.09.2021.
- Jakub Konrád, Jan Pichl, Petr Marek, Petr Lorenc, Van Duy Ta, Ondřej Kobza, Lenka Hýlová, and Jan Šedivý. 2021. *Alquist 4.0: Towards social intelligence using generative models and dialogue personalization*. In *4th Proceedings of Alexa Prize (Alexa Prize 2021)*.
- Petr Lorenc, Petr Marek, Jan Pichl, Jakub Konrád, and Jan Šedivý. 2021. *Benchmark of public intent recognition services*. *Language Resources and Evaluation*, pages 1–19.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. *Language models are unsupervised multitask learners*.