

# BMI: A Platform To Build Intelligent Dialogue Agents

**Francisco Iacobelli**

Northeastern Illinois University  
5500 N. St. Louis St, Chicago, IL 60625

**Gibran Cruz and David Ximenez**

Ibsen 40 Piso 2 Colonia Polanco  
Delegacin Miguel Hidalgo  
Mxico, DF CP 11560

**Salvador Suárez and Juan Cascajares and Andres C. Rodriguez**

Ibsen 40 Piso 2 Colonia Polanco  
Delegacin Miguel Hidalgo  
Mxico, DF CP 11560

## Abstract

AAAI creates proceedings, working notes, and technical reports directly from electronic source furnished by the authors. To ensure that all papers in the publication have a uniform appearance, authors must adhere to the following instructions.

## Introduction

The ability to communicate with customers is fundamental to a successful business model. Interactive websites, phone menus, chatting with a representative, etc. are all part of a communication strategy. Intelligent dialogue systems are becoming popular as part of the communications toolbox. However, they have been restricted to: (a) simple question answering, where the answer is usually an information retrieval algorithm that points to a website that may contain the answer; or (b) dialogues that are guide the customer through a very strict path and can be clumsy agents in terms of their linguistic abilities. **cite**. Dialogue systems that can communicate with users in natural language about more general topics usually get lost very quickly or...**complete this idea with examples** Lastly, although there are common documents to specify dialogue structure across multiple domains **citexml** they are catered to dialogue systems that intend to fill out a form that, when filled, signals the completion of a task and then the dialogue system triggers an action. These systems usually operate under the paradigm of *system initiative*. The *BMI platform* aims to facilitate the design of multi domain dialogue systems to be flexible enough so that users can communicate with them in natural language and under the paradigm of mixed initiative **cite this**. The platform also enables users to specify business actions to be taken at certain steps of the conversation or to take these steps depending on how the conversation is going. Lastly, the *BMI platform* enables companies to easily add/transition from different dialogue domains. **blah blah mejorar**. In this paper we present the basic architecture of the *BMI platform*, a tool that allows for a simple design of dialogue systems, and we discuss three actual implementations of this system. Two are in commercial use and the third is a prototype.**terminar esto**

Copyright © 2014, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

## Background

**write this whole section** Dialogue systems are a challenging

There are three main paradigms of architecture for dialogue systems **cite**. First, a form filling linear ones, then there's information state and then there's information retrieval based ones. In addition, there is a language to configure dialogue systems. However it is more geared towards the form type ones. Dialogues can be mixed, system or user initiative.

## Architecture

The BMI Platform is a dialogue system that is highly configurable. As such, one of its inputs is a configuration file. Figure 1 shows how the Dialogue Manager is in charge of interpreting the configuration file and retrieving all the data needed for the system to work. This data lives in *catalogs* or custom data tables. These catalogs can be created from delimited files as well as from other data sources. Then, when a textual input is processed, the Dialogue Manager coordinates the Natural Language Processing (NLP) components, Business Actions components and the Natural Language Generation (NLG) components with the World State, which acts as a unified information state –according to the information state dialogue system paradigm **cite traum**. Next we describe these components in more detail.

**Text Input** : This is any textual input. This input currently comes from SMSs as well as web applications and facebook. This input is converted into a message object and handed over to the DialogueManager.

**Catalogs** : The catalogs contain tables with information that is relevant to the domain and language of the dialogue system. These catalogs are in a custom database format, but any delimited file or database table can be imported easily into this format via *BlueForms* **cite the website**.

**Configuration** : This is a file that contains the sequence of steps –the flow– that a message object will go through and the updates to the WorldState on each step of the way. The configuration also contains pointers to the catalogs that will be used by the system, such as a list of products and features, machine learning training data, etc. Lastly, these sequences of actions can be divided in sequences that specialize on (a) inference (such as detecting the intention or

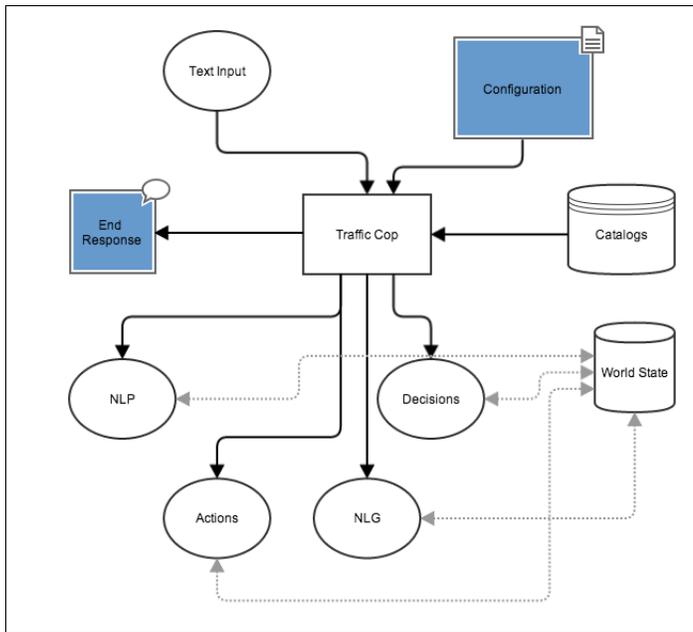


Figure 1: BMI Platform Architecture

sentiment of the message, or detecting entities relevant to the domain), (b) information retrieval from the catalogs given what has been inferred (such as retrieving the price of a product that matches a given query), (c) business rules that will result in an action such as a response (via natural language generation), a business action (such as an email or alert to customer service) or both.

**End Response** This is not a module per se, but the final response textual response, business action or both. This is the output of the system at every turn.

**NLP** : These is a set of natural language processing techniques that can be used by the dialogue system. Such techniques include Named Entity Recognition, Removing stopwords and other more sophisticated configurable techniques such as anaphora resolution.

**Actions** : These are a set of actions that can be triggered depending on parameters in the WorldState. Some actions include sending emails, connecting users to a support phone number or updating databases on the customer's side.

**NLG** : These are a set of natural language generation techniques that can be used in the dialogue system. In practice, the most frequent technique involves a series of response templates with markup to indicate which portions of the template are variables and which are constant.

**Decisions** : The Decisions module is a toolkit with machine learning techniques that are optimized for our dialogue system and can read and update the WorldState. For example, there are algorithms to decide whether or not to execute a business rule, or to respond with a given dialogue act.

**World State** : This component serves as an information state **cite** although not as formal as the literature describes them **cite**. The WorldState saves the output of every component invoked in the configuration. Those components, in turn, can read from the WorldState and take action. For example, if one of the NLP algorithms has detected the name of a product, that name is stored in the world state under a specific key. It follows that the structure of the WorldState is a dictionary that contains a *snapshot* of the memory and status of the dialogue system at any given point.

**Traffic Cop** : The traffic cop coordinates the actions specified in the configuration file in an orderly way. It is the module that coordinates the input and output to/from the WorldState and the other components. At every turn, the traffic cop knows where it is in the dialogue and knows when the dialogue starts and ends. The traffic cop is also responsible for instantiating each component that has not been instantiated. Under the paradigm of a dialogue agent, the traffic cop carries out duties that are normally assigned to a dialogue manager. However, in our system the management of the dialogue is specified in the configuration and therefore there are no components specialized in fully managing the dialogue.

## Implementation Details

Each implementation can, in theory, use a configuration file that uses each of the NLG, NLP, Decision and Actions components as desired. However, after many iterations of this prototype and several actual implementations on clients we share some techniques that we have implemented for some of the most challenging tasks.

### Determining Actions

Formally, the goal of the system is to find the action/dialogue act with the highest posterior probability given the state of the world.

$$A^* = \arg \max_{A \in \{da \cup actions\}} P(A|WorldState)$$

Where  $A^*$  is the optimal action or dialogue act,  $da$  are the individual dialogue acts,  $actions$  are possible business actions (that may include a dialogue act together with the action).

In practice, companies have rules about when to perform an action or output a dialogue act. Therefore, a purely probabilistic model may fail more often than it is desired because it adds probabilities to a deterministic procedure. In order for Ad-Chat to be able to accommodate both the rules established by companies and a probabilistic approach when an unpredicted situation arises, one of the decision components is a decision tree implementation that is able to follow rules, but also act probabilistically when the new situation has no rules associated to it. The implementation is based on the C45 algorithm **cite**. When a new situation is encountered it acts based on the entropy of the relevant fields of the world state that are present.

Entropy is a measure of how well an attribute predicts the classification of a dataset. More formally:

$$E = \sum_{v \in values} P(v) \log(P(v))$$

Lastly, because many of the system beliefs are probabilistic in nature –for example, if entity detection is performed, and entity is detected with a certain probability– the error of these algorithms can carry over and trum the final classification. Therefore, one last step in the determination of actions is to pass all probabilistic outcomes through a regression to compute the overall confidence of the system. That is, how confident is the dialogue system in its own decisions. If the system is not confident enough, a rejection dialogue act may be outputted.

### Anaphora and Ellipsis Resolution

Anaphora is a reference to an entity previously mentioned in text. For example, in the sentence “*John* is happy and *he* is celebrating,” *he* is anaphoric because it is a reference to the entity *John*. Not all anaphora are pronouns e.g. “*John* is happy. *The boy* is so lucky.” Resolving these anaphoric words and determining which entity they refer to is not trivial .

In the same way, Ellipses are omitted words that reference previously grounded information. For example: It is understood that if a person is at the store and asks for a smartphone, the clerk may offer a very expensive one. Thus the person may ask: “Do you have anything cheaper?” instead of asking “Do you have anything cheaper *than the phone you just showed me?*.” The omission of *than the phone you just showed me* is an ellipsis.

Anaphora an ellipsis are part of dialogue and need to be dealt with. Both anaphora and ellipsis resolution are active areas of natural language processing research . the *BMI platform* makes an attempt to solve this challenges with a mix of techniques.

First, On each update to the WorldState, the *BMI platform* keeps a pointer that references relevant entities from the previous update. Second, machine learning algorithms help determine which of these entities is more relevant to the intention of the user. If a reference or omission of an entity is detected, but the intention of the user is clear, the previous entities are used. If there are entities present and the intention of the user is clear, the appropriate entities are replaced with the ones found in the user’s message.

### Mixed Initiative

As discussed in the background section, the initiative of a dialogue system can taken by (a) the system itself –as is the case with reservation dialogues where the system is always in control of the information that the user should provide, (b) the user, by having her ask anything about anything, and (c) both, in a paradigm that is called mixed initiative

Because the WorldState acts as a repository of user beliefs, common ground and system’s beliefs, it is very easy to adapt and does not follow a strict frame-based architecture. Its rules are also designed so that the user can ask many questions and depending on the current WorldState different dialogue acts can be effected.

What follows is an example of a prototype system that is currently in place. This system is built on top of the *BMI platform* and helps users decide on a cellphone. Next we discuss evaluation metrics and accuracy of the responses. We

are also implementing the same system in Spanish and we present those metrics as well. Finally we discuss our result and present conclusions.

### Prototype: AdChat

In order to further showcase the *BMI platform* we developed a dialogue system that helps people find the best smartphone according to their needs. The system knows about 36 smartphones now in the market. The system’s data was gathered and prepared in English. This demo can hold much more complex dialogues and uses heuristics for anaphora resolution, machine learning algorithms and n-gram models for sentiment analysis and entity detection. The business actions are of two kinds: respond to users and send emails. Lastly, the NLG components are simple pattern responses with placeholders for information in the state of the world.

The main capabilities of the *BMI platform* implemented in the AdChat project can be seen in the dialogue in Figures 2,3 and 4.

The dialogue in Figure 2 shows a user that is exploring options about a phone he/she may be interested in. The user freely asks about the camera, prices while at the same time responds questions about wireless speed. In Figure 3 we see a user that knows what he wants to some degree and the system is able to accomodate his need. Lastly, in Figure 4 we see a dialogue in which the user wants to start up a conversation about a different topic, gets frustrated and insults the agent. We see the capabilities of the *BMI platform* to deal with such situations.

**Mixed Initiative** While many other dialogue systems offer mostly system initiative conversations, the *BMI platform* allows users to take the initiative as well. For example, in the dialogue in Figure 2 the system asks the user to provide a priority in terms of what is important on a smartphone and then it continues to guide the user by asking her to provide the internet speed. Then, the system provides one possible smartphone and the initiative is left to the user who asks about the camera on that phone. Then again, it is the user who asks about prices and processing speed without being prompted. At any point in the conversation, the user is free to ask anything about any particular model. Another example can be seen in Figure ?? where the user knows what she wants and disregard the first prompt altogether to ask for a specific product.

**Anaphora and Ellipsis Resolution** A simple ellipsis resolution happens when the system asks the user whether she wants a 4G or 3G device. The user simply enters 4G. However, a less trivial resolution happens when the system provides one phone –namely, the Sony Experia– and the user simply asks “how’s the camera on it” and the system understands that the anaphora should be resolved to the phone just mentioned. Another less trivial resolution happens right after, when the system provides the price and the user asks “do you have anything cheaper?” The system understands the context and resolves the ellipsis, thus providing a phone that is cheaper *than the three mentioned above*. Lastly, another ellipsis that is worth pointing out is when the user asks

BMI :Welcome to Phone Geek. I can find the right smart-phone for you. Please let me know what feature you are looking for (affordable price, high memory, powerful processor, etc.).

USER :i want a phone with a good camera

BMI : We have many phones so I am sure we have the one for you. Are you looking for a 4g or 3g device?

USER : 4G if available

BMI : We have the Sony Xperia Z available in 3G or 4G (LTE) connections. [Watch Video!](#)

USER : how's the camera on it?

BMI : Sony Xperia Z has an incredible camera with 20.7MP.

BMI : Google Nexus 4 has a great camera with 8MP .

BMI : Blackberry Z10 has a great camera with 8MP

USER : and how much is it?

BMI : The Sony Xperia Z prices starts at \$649.

BMI : Google Nexus 4 can be yours for just \$339 .

BMI : Blackberry Z10 prices starts at \$399

USER : do you have anything cheaper?

BMI : Out of our many products, if you are looking for high resolution camera, affordable phones I would suggest Motorola Razr D3 XT919 or Nokia Lumia 505. If you are interested please write your email and I will send you the offer details.

USER : how's the processor on the Lumia 505?

BMI : Nokia Lumia 505 is packaged with Windows Phone 7-8 an a 800 Mhz processor.

USER : and the Razr D3?

BMI : The Motorola Razr D3 XT919 is packaged with Android Jelly Bean 4.1 an a 1.2 Ghz dual-core processor. [Watch Video!](#)

USER : and how much is it?

BMI : Motorola Razr D3 XT919 can be yours for just \$249.00 in Moon Deal 5 or \$429 in Moon Deal 2. [Watch Video!](#)

USER : great, thanks!

BMI : Glad I can be of service. Do not hesitate to contact me if you have further questions

Figure 2: A sample dialogue illustrating complex behavior from the the *BMI platform*

about the processor of the Lumia 505 and after the system responds, the user simply asks “and the Razr D3?” instead of asking “and *how is the processor* on the Razr D3?” Thus, the dialogue is able to keep track of previously mentioned entities as well as attributes.

**Social Interaction In challenging situations** One simple social situation is how to respond to messages that thank the user. However, a slightly more complicated challenge is to respond to off topic conversations such as in Figure 4. Here the system detects when the conversation is off topic and tries to bring it to topic. Upon an insult from the user, the system has other strategies to deal with it. Lastly, is “shuts” the conversation by stressing that if the user wants support, it should call a real person. Another feature implemented in the system is that because the WorldState tracks the previous interaction, the system can count how many times in a row it is unable to help the user. A configurable option al-

BMI :Welcome to Phone Geek. I can find the right smart-phone for you. Please let me know what feature you are looking for (affordable price, high memory, powerful processor, etc.).

USER : hi, I want to know about the Samsung galaxy

BMI : Are you asking about the Samsung I9190 Galaxy S4 mini, Samsung I9295 Galaxy S4 Active or Samsung Galaxy S3?

USER : the S3

BMI : The Samsung Galaxy S3 comes with a Li-Ion 2100 mAh battery, has a great camera with 8MP, available in both 3G or 4G, with JellyBean 4.2.2 and quad-core processor, is a large device, comes with 16 Gb of storage and microSD slot, comes with additional features including social network services & GPS, available on Moon Deal 3 and 6, prices starts at \$549, has a 4.8” display with mid resolution, is the perfect high-end device. I’m sorry but we do not have the Samsung Galaxy S3 mini

USER : great, thanks.

BMI : Glad I can be of service. Do not hesitate to contact me if you have further questions

Figure 3: A sample dialogue illustrating user initiative and handling of products that are not in the catealog.

lows the *BMI platform* to take either a business action or say something sensible to the user.

## Evaluation

### Methodology

### Results

### Discussion

## Ad-Chat en Español

### Evaluation

### Results

### Discussion

### Conclusions

### Acknowledgements

BMI : I can help you find the right smartphone to match your needs. You can start by telling me which feature is most relevant for you (affordable price, high memory, powerful processor, etc.).

USER : are you human?

BMI : I think we are going off topic. I would love to know more, but for the time being I can help choose the right smartphone model for you. What are you looking for? (lots of memory, small size, etc.).

USER : oh, c'mon you suck.

BMI : This kind of language is unnecessary.

USER : sorry.

BMI : Sorry I could not assist you better, I hope I can be more helpful next time. Feel free to contact us any time or contact customer support at (800) 123 - 1234.

USER : are you mad?

BMI : Sorry I could not assist you better. For additional questions please contact our service center.

Figure 4: A sample dialogue illustrating the *BMI platform* handling of chit chat and social interaction.