



Tips for conversational design and user experience in Oracle Digital Assistant

Grant Ronald, February 2019

The power of a conversational interface, such as Oracle Digital Assistant, lies in the ability for a user to communicate using natural language using phrases and terms with which they are familiar and comfortable. So, rather than a user interaction being railroaded by the implementation of an IT system, the Nirvana is a naturally flowing conversation. However, the freedom of an open and loosely structured interface such as spoken and written language may exponentially add to the complexity of how you design these systems. The very nature of language is complex, often imprecise and ambiguous. It is therefore critically important to design conversational interfaces with strategies and techniques to mitigate this complexity.

This paper discusses a number of important tips and lessons for building a conversational interface including designing conversations, voice and personality, natural language intelligence, help, guidance and error handling.



Contents

Why Language can be difficult.....	3
Designing a conversation	3
Understand the objectives of a conversational interfaces	3
Business benefits	3
A conversational interface won't fix a broken CX (Customer Experience)	4
Planning the conversational elements of a project.....	4
It's a conversation, not a form	4
Context is king	5
Voice and personality	6
Design a personality	6
A warm welcome.....	6
Words carry emotions	7
Your digital assistant is not a human	7
Human-like conversation	7
Sympathy, humour and reflective listening.....	7
Handle smalltalk	8
General smalltalk.....	8
Foul and abusive language	8
Natural language intelligence	8
Multiple intents for the same action.....	8
Handle the things you know you don't know	9
Test, fix, repeat	9
Intent classification.....	10
Intent testing.....	10
Help, guidance and error handling.....	10
Offer guidance and help.....	10
Limit the surface are for errors	10
Fail gracefully.....	11
Conclusion.....	11



Why Language can be difficult

Human language has evolved over thousands of years and continues to do so. Like a living organism its evolution is influenced by everything from its geography to even the weather (As a native Scot we in Scotland seem to have a much more colorful palette of words for *rain* which are unknown to our English cousins across the border).

Furthermore, we deliberately use words to confuse and paint abstract concepts. Words are applied in situations where not literally applicable (metaphor, from the Greek meaning to move or carry) to help paint a picture (that phrase itself is a metaphor). The phrase “He seemed to bottle up his emotions” gives us an understanding of someone not sharing or controlling their feelings rather than the fact someone is physically putting something in a glass receptacle.

Metaphors often become so engrained in our language we cease to even think of the literal meaning and hence the birth (see what I did there) of the idiom. Hearing someone say that “the exam was a piece of cake”, on the balance of probability, has got nothing to do with an exam on baking.

And of course humour and newspaper headlines play on the very fact that words have overloaded meanings and can be used in a way to deliberate misdirect the listener or reader

My favorite headline story which it now turns out never actually make it to print related to British politician Michael Foot chairing a committee on nuclear disarmament: the headline read “Foot heads arms body”

All in all, language presents a complex challenge for the conversational designer. Let’s get into the specifics of how to address the challenge of designing a conversation.

Designing a conversation

Designing a conversation is as much a art as it is a science and so like art, there are no hard and fast rules or a guaranteed formula. However, experience gives us some guidance.

Understand the objectives of a conversational interfaces

This first point might seem blindingly obvious but what we do see is that is that a failure to set clear goals and then plan and design a conversation with those goals in mind is one of the primary reasons for project failure. Consider the following points:

Business benefits

In launching a conversational interface project what are the business goal on which this will be measured?

Take the following example: is deflecting customers from having to wait in a queue to speak to a live agent your primary objective? If it is, this might involve setting the thresholds for agent hand off to quite a high level thus you are aiming to have the digital assistant fully address the user’s request.

Conversely your business might dictate that the nature of the interaction with a customer is such that minimizing any incorrect information to the user is the critical factor and so you may be willing to hand off to a live agent much earlier in the conversation rather than risking the bot incorrectly handling the request. The goals as defined by the business will, and should influence the design of the conversation.



Ensure the business stakeholders are involved and set clear objectives as to what the digital assistant should achieve and how this is measured.

A conversational interface won't fix a broken CX (Customer Experience)

A digital assistant is a channel into your company and a digital assistant on top of a broken customer experience is just a new channel into that broken experience.

Whilst your digital assistant itself might be a successful implementation it can become the “touch point” for customer complaints. If the digital assistant is the one telling you that a parcel you ordered 2 weeks ago for a birthday isn't going to arrive in time, or if the digital assistant directs you to speak to an agent then puts you on a queue for 45 minutes, then the perception of failure might be placed incorrectly with the digital assistant.

I've actually heard someone say that a chatbot was a failure because it directed them to a phone number and no one answered. Or the case where the digital assistant correctly tells the user that they can't return the shoes that don't fit after having them for 12 months. The fact was the digital assistant was doing the correct thing, but the process failed, or the customer feels aggrieved, and so the result is the digital assistant becomes the focus for the complaint.

Ensure that in planning a project you need to ensure you are not building on top of a fundamentally flawed process and be able to discriminate between feedback on the conversational channel and feedback on the underlying customer journey.

Planning the conversational elements of a project

A conversational project is not so far removed from traditional software development that the same techniques can't be used, but there are new considerations to take into account

How might you expect the user to access your digital assistant? Voice assistants like Alexa? Facebook? Might the fact that SMS is a ubiquitous channel influence your choice? Or do you think most users will be accessing your conversational interface from a chat client embedded in your corporate website?

Furthermore, tone of voice, conversation flow, help strategies, exit strategies, learning loops, conversation insights also have their own challenges on the conversational channel.

Oracle runs conversation design experience (CDX) workshops which look to address these kinds of challenges. Using role play, personas, scorecards etc. you can start to build out a more concrete plan of the conversation side of the project

It's a conversation, not a form

Another common reason for a conversational design failure is getting into the mindset of simply putting a conversation on your existing systems.

Conversationally enabling an existing back-office or web form style of interface on a like for like basis is a recipe for failure. A use case currently implemented with 50 fields being turned into 50 pieces of information that the digital



assistant either prompts for or tries to interpret is not a conversation – and making this interaction “conversational” adds zero net-benefit to the interaction. Often these types of screens were driven by the data the system required rather than what a user might want to achieve.

Instead of being driven by the design of a set of database tables, look at what the user is hoping to achieve. Does the user need to enter all of that information? What information could realistically be defaulted? Could defaults be based on the context of the conversation? For example in an expenses use case if the digital assistant understands you are inputting expenses in US dollars could that information be assumed in further cases?

Conversation context might also depend on location which could be derived from user profile information. If you reasonably expect users to be inputting expenses “in situ” given the convenience of a conversational interface, then could currency be assumed based on current location and maybe even the expense establishment (e.g. Starbucks) could be derived from the location on a mobile device?

There is also the ability to further rethink the problem for the 21st century. Rather than asking the user to supply different fields of information could photographing the receipt be the source of some sort of optical character recognition?

A like for like implementation fails to take advantage of the power of a conversational interface and instead runs the risk of defaulting to the lowest common denominator whether that is a data-driven web page or an IVR system.

Context is king

As touched on earlier, human conversation is a complex beast and within a conversation we each hold a context within that conversation. That context might be based on things said earlier in the conversation or maybe even social or cultural context.

For example, “get me a taxi home” will mean something different to me as it does to anyone else since home is specific to the person asking. “Can you tell me who his manager is” will depend on the person you have been speaking about. A question such as “what about Friday” seems like a pretty straight forward, unless previously in the conversation you’d asked “can I book an appointment for the first week in March”. The implication is you are not talking about the next Friday coming up but the Friday in the context of the first week in March.

Now, having laid out the problem, what can you do about it? There no “magic AI” that is going to solve all of this for you, however what you can do, having been made aware of the challenge, is to try to mitigate the problem.

Sometimes it’s ok to ask for confirmation; just like a human might do. In the case of an appointment date above you might present an option: “Sure, we have a couple of Fridays free, would you prefer 22nd Feb, 1st March or 8th March.

Equally you could plan that based on your use cases the digital assistant will track the various “actors” in the conversation so that personal pronouns could be resolved. It may also be likely that you are able to identify the user (possibly because you ask them to authenticate) in which cases phrase like “home” can be contextualized.

Context is something you need to decide whether is important in your use case. It may be you only fully understand the need when users start interacting with your digital assistant



Voice and personality

With a website, your colors and logos define the brand of your company and your staff further carry the company brand. With a conversational interface your brand is carried in the words and tone of the conversation. Therefore the design of the digital assistant persona and voice is a critical facet and not something that should be left to the creative whim of your developers.

Design a personality

In our human lives, when we speak to someone we have an implied expectation that a conversation is engaging. And how we feel about a conversation will last beyond when the actual word have disappeared. How do you want the user to feel having engaged on your conversational channel? Does the engagement deserve a certain level of gravitas? Are your target users such that you might expect the conversation to be more informal? Does that match the brand of your company? Is it appropriate for the use case?

Often when a company is designing a conversational channel, they will build a complete persona around the digital assistant; even to the point of giving it a name, a background story, where the persona was brought up or educated –possibly even going as far as employing a voice actor to help give substance to the persona.

Now, that said, that doesn't mean your digital assistant should try to pass itself off as the persona. The current thinking is you make it clear to your users that this is an automated conversation channel and not a human, although it's totally fine to use an appropriately androgynous name and avatar. However the persona is an aid to the design of the voice and personality of the digital assistant.

Think about the brand of your company; think about your target audience; think about the use case. These are all influencing factors in the design of the digital assistant persona.

A warm welcome

A warm welcome and some guidance on what functions the digital assistant can help with is a great way to kick off a conversation.

That said, the conversational channel requires that you keep each turn in the conversation short and to the point. Also, depending on your use case, frequently returning user might tire of the same opening line with the same guidance on topics with which they are already familiar. In this case you may consider firstly varying digital assistant responses or changing the welcome message for a returning user to not only avoid repeating information they already know, but instead to recognize that they are returning users and possibly even shape that opening exchange. For example, “welcome back, are you looking to order another pizza”.

Consider an opening welcome but keep it short and concise.



Words carry emotions

The subtleties of language mean that an incorrectly crafted response might have unintended consequences. For example, “you neglected to supply the order number” carries a certain accusation whereas “Can you let me know the order number and I’ll check that for you” mirrors a much more positive tone.

Sometimes the language of the digital assistant can be inadvertently influence by the backend IT system. “Invalid order ID” might be a quite reasonable error message for a database query but less desirable in a conversation. Instead: “I couldn’t find an order with that number” is saying the same thing, but in a way that shifts a certain implication of blame from the user providing an invalid number, to the digital assistant who is not able to find the order.

Throughout the conversation, but especially for error feedback, ensure the response does not imply any sort of blame.

Your digital assistant is not a human

Current thinking is never try to fool your user into thinking they are talking to a real human. However, that doesn’t mean the digital assistant can take on some of the more positive qualities of human conversations.

Anthropomorphism is that attribution of human-like traits or emotions to an object; and this is something you can consider for your bot. This might include giving the digital assistant a name and an avatar although you should choose something which doesn’t lead the user to assume it is human.

Human-like conversation

One feature of human conversation is it that it varies. Subtle changes in repeated responses or tailored welcomes for returning uses all tend to add to a more engaging conversation that feels less automated.

You may also vary responses based on whether the digital assistant finds the user struggling with a particular step and so give additional information. For example the initial prompt of “Can you let me know your order number” might be followed up with “Let’s try again, please type in your order number. It’s the 16 digit number starting with X”

In a human to human conversation people speak at a speed that the other party can comfortable listen and process. For a digital assistant the equivalent to not send too much information at one time. It could be that breaking up a response into separate messages with a short wait time (and maybe with a typing indicator) may make the response easier to digest for the human.

Finally another hugely important feature of human interaction is that just like a digital assistant, humans can get confused as well; so it’s ok to ask for clarification if there is ambiguity.

Vary responses, if possible tailor responses to the user at a pace it is easy to consume - and when not sure, ask!

Sympathy, humour and reflective listening

Sympathy and humour are, generally speaking, areas to stay clear of for the digital assistant. They are fraught with dangers of misunderstanding or being wide of the mark. That doesn’t mean a well placed “sorry to hear that” when



you are confident, for example, that the user is reporting a parcel was delivered with the contents damaged but no one wants sympathy (or a joke!) when ordering flowers for a funeral.

You may also consider the idea of reflective listening which helps convey and validate the conversation. For example, when the user reports that the parcel contents were damaged the digital assistant might respond with “OK, sorry your parcel was damaged, let’s get that returned for a refund”

Error on the side of caution with any sort of human empathy or humor, however where you are sure, reflective listening can be used as a way of affirming and clarifying

Handle smalltalk

We’d like to think that the humans engaging with a digital assistant know exactly what to ask and engage directly with a clear purpose. The reality is somewhat different.

General smalltalk

When first engaging with a digital assistant you may find questions such as “are you a human”, “is there someone really there” are surprisingly common. It therefore makes sense to prepare to handle those questions rather than responding with a stock “sorry I didn’t understand you” response.

That said, these responses don’t have to be the start of a deep and meaningful conversation, the goal should be to get the conversation back on track. So for example “What’s the weather like where you are” to a flower ordering digital assistant might elicit the response “Flowers love rain and sun so the weather is always good. So would you like to order some flowers now?”

Foul and abusive language

With the best intentions, you may find your customers come to the digital assistant with serious grievances which may result in the use of swear words. Being able to identify that the user may be upset through their use of abusive language can help your digital assistant deal with the situation appropriately. You might acknowledge that you think the conversation could be best handled by escalating this to a human agent.

Dealing with smalltalk is of course not the primary use case for your digital assistant. However, given that smalltalk is often found in conversational interfaces the ability to deal with it can improve the perception of the interface to your users.

Natural language intelligence

The ability for your digital assistant to be involved in a conversation relies on it being able to understand, or more specifically, to classify user utterances to known intents. How you design and train your natural language understand is critical to the success of the conversation.

Multiple intents for the same action

When you start to design the intents that your digital assistant can understand you might typically design one intent for each use case. For example, your digital assistant should be able to handle requests about your company’s



returns policy. For this you might envisage training phrases such as “What is your returns policy”, “do you allow customers to return items” etc. However, the reality is that you may well find different groupings of questions all of which should be handled as a question about the returns policy.

For example, you might find a class of questions such as “I bought this DVD player last week and it is broken” or “I bought this in your store and it doesn’t work”. In which case you might design a new intent for this use case even though the conversation flow and result (i.e. explaining the returns policy) is the same.

By doing this you help with the following

- You have a cleaner classification of training utterances in each intent which may help the NLP disambiguate user input
- You have the ability to respond with appropriate reflective listening: “Sure, I can help you with our returns policy” or “Sorry to hear that let me help you with our returns policy”
- You could use analytics or insights to understand the most common reason for triggering this particular conversation flow

Always consider that each intent may involve different classifications, or types or questions. In which case it is perfectly reasonable to build multiple intents for the same conversation flow

Handle the things you know you don’t know

When building a digital assistant, especially the early release, you will have a limited set of functionality the digital assistant can support. However, that doesn’t mean you can’t at least acknowledge when users ask a question that the bot does not yet handle.

For example, a banking digital assistant might currently handle checking of balance and transfer of money but not the dealing in shares. Rather than responding to question about share dealing with “I’m sorry I don’t understand”, rather you could respond with “Sorry, share dealing is not something I can do yet, however if you go to www.website.com you can buy and sell shares from there.”

To do this you need to identify common classes of question that customers ask and create intents for each of these which respond appropriately without actually performing the requested task. By doing so you are helping the user by acknowledging their question.

Create intents for common questions that you know the digital assistant will be asked, but functionally doesn’t yet support.

Test, fix, repeat

When building natural language understanding (NLU) for your digital assistant, it needs a corpus of well classified and representative utterances; ideally real data from real users. To do this, you need to have a strategy for harvesting, classifying and testing these test utterances. This not only gives you valuable training data, but insight into how the users are interactive with the digital assistant.



Intent classification

The first release of your digital assistant will be the one where the natural language understanding is in its most rudimentary stage. But every release you can improve the NLU. From each testing or launch of your digital assistant you need to capture user utterances. In many cases it is obvious which user utterances should map to which intent. For example “I want to order some flowers” could be clearly seen to be mapping to the ORDER_FLOWERS intent.

However, a phrase such as “I ordered these flowers on Monday and they have not arrived yet” could be regarded as TRACK_ORDER intent or maybe even RAISE_COMPLAINT intent. Only the business can decide how this should be dealt with

Intent testing

Having harvested a thousand new phrases from your conversation logs how should you use those? Firstly and most importantly, you need to understand if these new phrases are going to improve your model or not; and the only way to confirm this is through testing.

A general recommendation is that you have an 80/20 split: 80% of new utterances are used to train the model and 20% are added to your test data.

You are now in a position to add 800 new training phrases and test those with 200 different phrases. Having done this you can then establish if the model is showing a higher accuracy.

Ensure you have a clear plan to harvest new training data from customer logs and insights. Use this to both train and test the model.

Help, guidance and error handling

No user will immediately understand how to use your digital assistant perfectly first time: and that is ok. In a human to human conversation it is equally understandable to ask for help, clarification and to get things wrong. It is how you manage these situations that will leave the user with a lasting impression.

Offer guidance and help

In coming to a new digital assistant it may not immediately be clear which functions it supports. It can be helpful to open with an indication of common functions and ways of asking for help. But as noted earlier you may not want to show that every time and to every user.

Techniques such as quick replies (although not supported on all platforms) can be a way of ensuring the user knows what information is expected. Also making clear that help or reset abilities are available should the conversation become stuck.

Ideally you should also allow the user to explicitly speak to a human agent or detect when the conversation has got to a point that the intervention of a human agent might be helpful.

Plan strategies for allowing the user to request help or exit their current state.

Limit the surface area for errors



The power of a conversational interface is also the thing that is likely to cause the most problems. For example the user asking for things the digital assistant cannot support or long meandering complex input. In order to mitigate against this, you may consider using techniques such as quick replies, cards and carousels which display the exact choice the user has. Of course you have to plan this since a visual carousel isn't going to help if the user is accessing on a voice channel.

Other techniques include avoiding, or at least being careful with, open ended questions that fail to provide guidance of the type of response the bot is expecting. Consider “Anything else I can help with” versus “Do you want to check out or have a final look at the menu?”

Depending on your use case you should also ensure that any important transaction that the digital assistant is going to perform on your behalf should be confirmed with the user before doing so. You could inform the user with a message then ask for them to confirm with a yes or no, although there are some recommendations that suggest rather than asking a question which expects a yes or no, your confirmation should be the verb or noun. For example, displaying two quick reply buttons with “Go ahead and order” and “cancel purchase” in some way gives an extra confirmation above and beyond a yes/no.

Be on the lookout for places in the conversation where you might give the user too much freedom. Try to present the most logical choices to help guide them.

Fail gracefully

Inevitably things will go wrong and it's important to recognize that while you don't set up to deliberate blame the user, the choice of wording might convey a negative connotation even if unintended. For example “That is an invalid order number” versus “I couldn't find an order with that number”. Both sentences essentially mean the same thing but the second is gentler in that the focus is on the digital assistant not finding the order versus the implication that the user failed to provide a correct order number.

An error condition may also be the opportunity to reinforce the digital assistant personality. Rather than a formal “I did not understand your input please try again” something like “Sorry my robot brain is still learning and I didn't quite understand you”. Furthermore, whenever the digital assistant fails to resolve a user input above a particular threshold it can additionally respond with “...did you mean” and then list the top three resolved intents which it thought the answer might be.

It is ok for the digital assistant to not understand or fail at a particular step – that happens with a human to human interaction as well. It is how you handle it that will make the difference.

Conclusion

There are no hard and fast rules for developing conversation between a human and a digital assistant. Instead it is an art based on stories and experiences but most importantly you should plan to measure, evolve and ultimately improve that conversational experience.