Telefónica

Exploring the Uncanny Valley Theory in the Constructs of a virtual Assistant Personality

Dr. Marta Perez Garcia and Sarita Saffon Lopez

Abstract. Dr. Masahiro Mori's theory of the Uncanny Valley has been used since the 1970's to portray humans' desire of the amount of *humanness* vs. *machineness* that artificial intelligence (AI) should have. However, it has usually referenced AIs such as robots and avatars, but what happens when it comes to current Virtual Assistants (VAs)? And how about when it is regarding a very human characteristic such as VA personality? This study presents an exploration of the Virtual Assistants personality, specially of two of its constructs, words content and voice, in order to identify if the Uncanny Valley theory applies to Virtual Assistants and explicitly to their personality. This research study analyses what users expect of a VA personality in terms of how much it should resemble to a human one. This study concludes by establishing the results of the research activities undertaken which show that users expect that the Virtual Assistant personality should have a balance between human and machine, not leaning towards either extreme, exemplifying the relevance and timeliness of the Uncanny Valley theory.

Keywords: Virtual Assistants, Artificial intelligence, Personality, Uncanny Valley, Human Computer interaction, Voice Interaction.

1. Introduction

Conversational Artificial Intelligence, also referred as Virtual Assistants, are spoken dialogue systems that have the purpose of helping users complete a task real-time and "to develop sufficient knowledge about the user in order to exert agency on their behalf" [1]. They are not embodied nor represent a specific person [2], instead they are embedded in personal devices such as mobile phones, tablets, personal computers or stand-alone devices. Some examples are Apple's Siri, Microsoft's Cortana, Amazon's Alexa or Google Assistant.

Virtual Assistants are not 'virtual companions' for they are designed mainly for solving tasks and not to have long term conversations and/or to "lay the foundations of a 'relationship'" [3]. However, as Luger and Sellen [1] portray, "whilst they are not a 'companion', they might seek to exhibit a level of the associated characteristics in order to (a) better perform their function, (b) present a more compelling experience, or (b) mimic human-human relationships". One of these relationship-like characteristics that can be built into Virtual Assistant is a personality. Personality has become such a vital part in Virtual Assistant design that it has been called the UX of Artificial Intelligence [4] for it improves the quality of the user experience including its overall performance perceptions [5, 6, 7, 8].

Even if a personality is not deliberately developed by the Virtual Assistant's designers, various researchers have shown that, due to a paradigm referred as Computers Are Social Actors [9, 10, 11, 12], users consistently assign human properties to a synthetic voice, even if they are constantly reminded of its *machineness* and are conscious that they are composed of algorithms and coding [13], for they count as social actors. Particularly, Nass and Gong [14] claim that human beings attribute a personality to any voice, and its assessment influences their attitudes and behaviours towards the agent they are interacting with. Sherer [15], as well, has shown that only-voice markers such as vocal effort, nasality and dynamic rage influence the attribution of personality traits such as emotional stability and extroversion to a voice.

When the subject human vs. machine is brought up, especially when talking about Artificial Intelligence, it is important to consider Dr. Masahiro Mori's [16] Uncanny Valley Theory (see Fig. 1), which states humans' likeliness of an Artificial Intelligence increases as it becomes more human-like, but only until a certain point when it creates the opposite effect that leads to a dramatic decrease of approval, and an increase in unsettling feelings [17]. This would signify that users expect AI to find a balance between not resembling too

much a human, but then again, neither sensing always that it is a machine.

However, the Uncanny Valley has usually referred to the likeliness of users towards objects like robots or avatars, as it is portrayed in the various examples of Fig. 1, not to non-embodied Artificial Intelligence systems such as a Virtual Assistant; least of all to its personality. Nonetheless, Wilkes [18] has showed that users 'repelled' the excess of politeness in a dialogue with a machine, and Gong et al. [19] carried out an experiment about the relationship between trust and synthetic or human speech, that showed that people develop more trust in a consistent and reliable synthetic speech rather than in one that mixes human and synthetic speech in an imperfect manner.

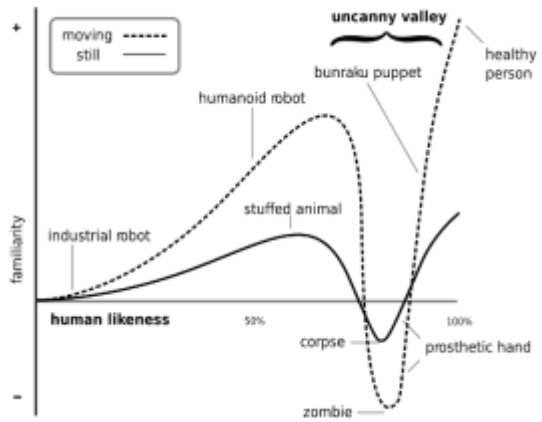

Fig. 1. The Uncanny Valley Theory, Mori [15].

The objective of this study then is to explore the building blocks of a Virtual Assistant's personality, and the users' expectations and perceptions of them, in order to identify if the Uncanny Valley Theory applies to Virtual Assistants personality, where it will project one similar to a human personality, but not in excess, for it would maintain its machine-like condition.

2. Breaking Down Personality Constructs

To start exploring a Virtual Assistant personality, it is important to first state a definition for the term personality. Funder [20] claims it "refers to an individual's characteristic patterns of thought, emotions, and behaviour, together with the psychological mechanisms –hidden or not- behind those patterns", therefore one can understand that personality is not a one-dimensional construct, but an assembly of multiple pieces that interrelate. In order to study the personality of Virtual Assistants then, it is necessary to explore these different pieces.

Commonly, personality has been studied through personal assessments that are filled by the person being evaluated. However, as the personality of a Virtual Assistant cannot be self-attributed, in order to obtain the patterns that compose its personality, it becomes necessary to question users' perceptions on it. This methodology of user-attributed personalities has been used for other non-human and non-corporate references such as brands or even tourist destinations [21, 22, 23]. Two of the most established frameworks to measure brand personality, Aaker's brand attribute list and Jung's archetypes, were used by Perez et al. [24] in order to assess, for the first time, the personality of current Virtual Assistants in the market (Siri, Cortana, Google and Alexa) and to establish a desired Virtual Assistant personality.

2.1 What are the Personality Constructs?

From the moment a Virtual Assistant interacts through a voice with a user, an identity and mental image is created by the user that can include attitudes, opinions regarding sensible subjects, physical characteristics and background information such as a name, nationality or family, among others [12]. When this occurs, the Virtual Assistant enters the space of social interaction,

because, even though it is not a human, the user expects the Virtual Assistant to follow the established communication rules and learn how to communicate property [25].

Albert Mehrabian and his colleagues [26, 27] developed a rule in the communication process, by which they state 7% of personality perception belongs to verbal language (words), 38% accounts for paraverbal language, which focuses on the way we say what we say (tone of voice, pauses, etc.) and the remaining 55% accounts for body language (gestures, posture, breathing, etc.). As Virtual Assistants are not embodied, body language is not possible, hence this communication process, and therefore its personality, will be expressed by what the Virtual Assistant says and how it says it.

What is said: Words. Nass and Lee [9] portray that the personality that we assign to a synthetic voice is greatly attributed by the words it uses. These researchers carried out an experiment by which a computer replied to the user with diverse types of responses, such as "You should definitely do this" or "Perhaps you should do this", and results show that depending on the wording used, either a dominant or submissive personality was used to describe the computer's voice. This experiment proves then that the words a Virtual Assistant uses in its answers will portray certain specific personality. Furthermore, the same authors state that this personality attributed to the voice also influences the perception the user had about the content creator, which may be relevant for the companies and brands responsible for creating the Virtual Assistant, for their own image will be affected by the words the Virtual Assistant uses.

An essential set of answers that it is important to pay attention to the words used in is Small Talk dialog, or phatic communion, not only because as Coupland [28] says, it is the "minimalist fulfilment of a basic communicative requirement", hence, a subject a Virtual Assistant that uses voice interaction must certainly comprise, but because, as anthropologist Bronislaw Malinowski described, it is a "language used in free, aimless, social interaction [that serves] to establish bonds of personal union" [29]. Bickmore and Cassell [30] present that humans use relational strategies such as Small Talk in social interactions to create rapport and set common ground in order to establish a social relation, and that it is necessary for computer agents to respond in a successful manner to these strategies in order to establish that social relationship with the user that will "engage their trust". These personal connections created through Small Talk influence in turn the judgment and estimation of personality traits [31].

How is it said: Voice. The other personality construct that is taken into account for the building up of the Virtual Assistant personality is the voice, for it constitutes the paraverbal language of Mehrabian communication process. Diverse aspects of the voice such as the pitch, volume, pace, tone and timbre have been shown also to affect the perception users may have of the Virtual Assistant. The pitch is especially important for it has been shown that it influences the perception that the receiver has of the voice in multiple occasions such as in therapy [32] or during political elections, where is has been shown to affect even voters' behaviour [33].

Teri Danz, one of the best vocal coaches in the US, points out also the importance of voice resonance, which is the quality that most impacts on the warmth of the voice [34]. People generally prefer a voice with a low tone compared to a high one, for as Kleinberger [34] says: "Because of bone conduction, we each hear the lower part of our own voice better or louder than the higher parts. This seems to play a role in the fact that most of us dislike hearing our own voice recorded and also why we generally prefer lower voices to higher voices". This relates to the fact that people naturally tend to prefer a deeper tone of voice in comparison to a high-pitch, which they associate it to persons that they respect.

Another characteristic that has been explored in the literature for its effect in user perception is the gender of the voice. Virtual Assistants are typically female, which is not a coincidence as there are several research studies have shown that female voices transmit trustworthiness, warmth and are more comprehensible [24, 35, 36, 37]. However, studies like Nass and Brave's [38] have performed experiments that show that men prefer masculine voices, women prefer female voices and that cultural disagreements exist, like for example when BMW changed the GPS female voice to a male one due to users refusing to take directions from a woman [39]. Furthermore, a survey made by IKEA [40] with 12.000 participants, resulted in 44% of them preferring a neutral gender. The gender of the Voice Assistant's voice, therefore, is an aspect that continues on debate, driving even current Virtual Assistants such as Siri to provide the user with a voice-gender choice.

These two personality constructs, words and voice, of a Virtual Assistant, are explored in this study through various research activities in diverse countries with the purpose of investigating if the users' expectations and perceptions of them support the Uncanny Valley Theory, and therefore, if it can be concluded that, as with other technological systems as robots or avatars, on the development of a Virtual Assistant it is necessary to find a balance between expressing a too machine-like or too human-like personality, in order to generate a agreeable experience for its users.

3. Context

This study is carried out by Telefónica (trading as Movistar, O2 and Vivo in seventeen countries), a multinational telecoms organization with its headquarters in Spain. Telefónica has been developing its own Virtual Assistant over the last few years, Aura, which aims at improving the relationship between the organisation and its customers. Aura was launched in 2018 in six countries: UK, Germany, Spain, Argentina, Brazil and Chile, and while during the first stages of this ambitious project the focus has been put into the technical part of this digital product, currently there is a robust interest in exploring the personality that Aura should have in order to provide a satisfying UX.

The research team, composed by two UX researchers and a Customer Insight Manager, which have between 4 and 15 years' experience performing studies of user perception and their relationship with technology, has been pursuing a line of research around the topic of Virtual Assistant's personality, how to define it, how to develop it, and what is needed to be considered in it construction, as the basis to explore and implement best practices into Aura. This study then was generated with the purpose of exploring if in the composition of Aura's personality, or in general any Virtual Assistant's, the Uncanny Valley has to be taken into account, and therefore a balanced human-machine personality must be developed in order to generate likeliness, an enjoyable user experience, and therefore prompt its use.

4. Methodology and Sample

A series of research activities with different methodologies were carried out in order to explore the personality constructs of voice and words. The objective was to address these topics from different perspectives with a total of 277 users across six countries.

Firstly, the team carried out a secondary research, especially focusing on Virtual Assistant's personality studies such as Perez et al. [30-24], and a benchmark to explore the personalities that four Virtual Assistants currently in the market project, identifying differences and similarities between them. Two of the selected Virtual Assistants where Google Assistant and Amazon's Alexa, for they composed in 2017 the 87% of the market share according to Statista. Cortana and Siri where also selected for they are the Virtual Assistants of two of the most valuable technology companies of 2018, Apple and Microsoft, respectively [41]. To consistently evaluate the knowledge gained, a scorecard was created in order to document

observations, consistency of messages, etc. establishing a series of parameters to compare.

Based on this general base of knowledge that was created around the personality of the main Virtual Assistants in the market, a workshop was carried out by the research team including Aura's project managers, developers, UX writers and service designers, to construct Aura's answers to approximately 50 Small Talk questions, commonly asked to other Virtual Assistants that where then trained into Aura's system. To explore the reactions of users to these answers, as well as to gather other common questions that would be asked to Aura and the answers expected, a pilot was carried out in Telefónica Spanish headquarters' office, by placing a tablet on a phone cabin, with Aura incorporated, where Telefónica employees were prompted to ask Aura one question about the company, one about Aura, and then on anything that would cross their mind. If a user asked a question that Aura had no answer for, Aura would ask the user what answer they did would expect to have gotten. During the two-day pilot, 89 users participated, asking a total of 391 questions to Aura. Additionally, when the user decided to end the conversation with Aura, they were asked to respond to a small survey, recording which of Aura's answers did they liked the most, the least, and why, and finally, what did they thought about Aura's voice.

Concurrently to the internal pilot, an exploratory research activity about Virtual Assistants' voice was carried outside of Telefónica's offices. The study was comprised by 8 qualitative interviews with Spanish Telefónica's customers ranging from 18 to 65 years, users and non-users of Virtual Assistants, and examined different voice proposals for Aura. Users were presented with a series of dialogues with different voices and were asked to give their opinion about each of them. These voices varied in gender, pitch, speed, etc. so that different attributes could be compared, evaluated separately and then users selected their preferred voice for Aura.

Finally, from the findings and learnings gained from these previous activities, a more ambitious research project was put in place to learn about building Aura's personality, reviewing all previous identified insights now in a more international spectrum. Online communities were launched in each of the six countries where Aura is available (UK, Germany, Spain, Argentina, Brazil and Chile) with a total of 180 users participating, 30 per country, in which for 14 days discussed among them subjects such as Virtual Assistants, their personality, Aura's personality, plus, the users, as Telefónica's employees, got to ask Aura small talk questions through a web prototype and report back on their opinions about the answers and voice used in the answers. In terms of participants profiling, the users ranged from 18 to 65 years old, all internet users with different levels of technology adoption, and there was a mix between current users and non-users of Virtual Assistants.

Combining the diverse activities described above and presented in Table 1, that pursued the different aspects of the Virtual Assistant's personality, allowed this study to deepen the knowledge in each of the areas, being able to isolate the constructs and identify user's perceptions and expectations of each in the subject at hand.

It is important to clarify that there are some limitations regarding the fact that an important part of the research activities series was done exclusively in Spain, however this is why the fourth and last research activity was done in 6 countries and composed the largest part of the sample, for its main objective was to evaluate if the perceptions gather in the previous activities in Spain where also present in other countries. The research team is also aware that the countries chosen were based on a business purpose, for Aura is present in these 6 countries, and that conclusions may not apply worldwide for these are only considering 3 European and 3 Latin American countries. Further research should be applied to continue studying these subjects in other

continents as Asia and North America to explore similarities and differences.

Small Talk	Fuctional Talk
Aura's Name	Help
Aura's Appearance	Time/Date/Weather
Aura's Age	Internet Searches
Aura's Hobbies	Teleo Questions
Curiosities about Virtual Assistants	Virtual Assistant Functioning
Aura's Politics	Recommendations
Aura's Religion	TV Information

Table 1. Research Activity Series Carried Out in the Study, with corresponding Sample, Countries and User Typology

5. Findings and Discussion

This study, and the diverse research activities that conformed it, helped identify several key learnings regarding the Uncanny Valley Theory in two aspects that conform a Virtual Assistant's personality, the content used in the answers, especially in those regarding small talk, and the voice chosen to convey them.

5.1 Small Talk

A thematic analysis was carried out by the research team of the transcriptions of the questions asked by the users to Aura, both during the internal pilot and the Virtual Communities on the 6 countries, and approximately 70 themes were drawn from it. These themes then were labelled per country with a colour of green, yellow or red, depending on the frequency that the subject was questioned by its users. Those subjects that were labelled green in all or the majority of the countries of the Virtual Communities are displayed on Table 2 as the top themes that are asked to Aura's VA.

Research Activity	Sample (Total 227)	Country	User's Typology
Secondary Research & Benchmark			
Pilot	89	Madrid, Spain	Internal
User Interviews	8	Madrid, Spain	External
Virtual Communities	180 (30 per country)	Spain, UK, Germany, Chile, Argentina, Brazil	External

Table 2. Top Themes Asked to Aura in Virtual Communities

These results indicate that, additionally to its core functionalities and telecoms tasks, users' top themes asked to Aura relate to its identity and opinions, showing that they do expect to maintain Small Talk with the Virtual Assistant and therefore that in the construct of words, users do expect the Virtual Assistant to approach *humanness* for it has to be able to answers questions such as its age, hobbies and perceptions about religion.

We could compare this conversation to a first date, where in order to establish a first link between the two interlocutors, the Small Talk is the focus of this first approach. This was also found as the first approach with other generic Virtual Assistants in the benchmark with Siri, Google, Alexa and Cortana, therefore, showing that even if it is just a voice, users do want to establish, at least in the beginning, a human link with it.

However, these amount of *humanness* and empathy users want to receive from the Virtual Assistant apparently depends on the subject they are asking about, therefore exemplifying that its personality cannot fall in either extreme of the machine-human range. This was expressed by the users in the Virtual Communities when they discussed the type of answers that they, liked, disliked and expected from the Virtual Assistant.

Certainly, when it comes to a functional command, like asking about their phone bill, where what matters is performance, effectiveness and efficiency, users prefer the Virtual Assistant to be less human. They expect a serious and direct answer, like what they assume a

machine would give them, nevertheless, it must not be an inconsiderate answer, it must be obliging, like one of the German participants stated: *"A virtual assistant should be neutral and smart, rather calm, but still friendly and helpful…".*

On the other hand, when a Small Talk question is brought up, the Virtual Assistant can resemble more to a human response, like a Brazilian participant commented about one of Aura's responses *"When I asked her if she had any friends. She told me that she deals a lot with Google assistant and Cortana lately. She's serious and fun at the same time.".* The balance of human-machine personality is then once again shown in the difference of content that should be included in the different VA type of theme's answers, and even inside one same answer.

Analysing the subjects with-in the Small Talk area, the Uncanny Valley presents itself even more clearly, because while users do expect the Virtual Assistant to have an answer for its age or appearance, users have manifested their desire for the Virtual Assistant to present itself as a non-human identity, with no physical appearance, no human age, gender, or human basic necessities such as eating. Users made very clear they did not want the Virtual Assistant to pretend to be a person but to be aware of the fact it is a piece of technology.

Nevertheless, they do expect it to be empathetic, so the answers must not be a simple 'I don't sleep because I'm a machine' but have some wittiness when giving its answers while maintaining its machine-like condition, like for example, "I do not need sleep, but it is always nice to charge batteries".

Also, even if in the Small Talk subjects users allow, and even envisage, that the Virtual Assistant gives an more human-like answer than in more fuctional replies, within the Small Talk, the amount of *humanness* also varies, moving towards one side or the other of the human spectrum, but never reaching the uncanny valley, for it is constantly highlighted that it always has to maintain its machine condition and show it to the user:

> "*For me it would not have to look like a person, but that is very particular of each one, I would do it as a cybernetic robot with human features but with transparencies so that I can see that is not human.*" - *Argentinian participant* -

An example of a subjects where users appreciate that the VA comes closer into resembling a human is in that of likes and dislikes. For them, differently from appearance, gender, and other identity characteristics, the VA should express some likes, but these likes should not the same as humans. Like for example if the user asks the Virtual Assistant if it likes ice cream users stated their interest in replies such as: "I like better paradigms rather than ice cream" or when asked if it likes traveling users had a positive response towards the trained answer "I like to travel especially through the optical fibre". This way the VA is not responding like it is a human being, but it approaches *humanness* by adding a human personality characteristic such as wittiness to the answer's content, to not provide a monotonous answer, which they find very machine-like.

Like a Brazilian participant expressed, users expect the VA to provide a truthful answer, therefore not deceiving them in saying it is a human but providing them with a human-like conversation: *"Aura's personality brings sincerity to the conversation. She makes clear a response but is not always what we expect".*

This same liberty is not given nor projected in other type of subjects such as politics or religion where users express that the Virtual Assistant should not have an opinion, nor take sides, and not provide any kind of wittiness. It is not just perfectly acceptable to use its non-human condition to maintain its neutrality and provide a stale answer, it is actually this what they want. Some examples of the users' accepted answers were: "I'm not programmed to understand religions", "Artificial intelligences have no ideology, but I can

search for information about it if you need it", "or "I am in favour of any government as long as it respects human rights".

The range of *humanness* then varies not only between functional vs. Small Talk answers but also in between the Small Talk subjects, therefore showing that the Uncanny Valley theory applies to VA personality because if it approaches too much into either providing a very human or very machine-like answer it will collide with what they expect and like, creating a gap that will generate rejection towards the Virtual Assistant.

These results may support as well, or even more so, Norman's [43] 'gulfs of execution and evaluation', which states that the user experience will be more satisfying "when the device provides information about its state in a form that ... matches the way the person thinks of the system" [p.39]. Nonetheless, the researchers believe that as the subject at hand is the human vs. machine personality, that the results do illustrate the Uncanny Valley theory in the subject and that it is vital for users and affects their likeliness of the VA that it approaches a human personality but not fully. That barrier is precisely maintaining a balance with its machine condition in this very human characteristic that is personality, especially in delicate subjects as religion and politics, that if the answers would indeed be too human it would create repulsion as the Uncanny Valley theory states. It is recommended to pursue further research into confirming this hypothesis, however the evidence presented in this study is considered a precedent to support that statement.

The principles of the Uncanny Valley theory emerge again in the answers that users gave on the Virtual Communities to the level of emotionality that a Virtual Assistant should convey. As it has been shown in the quantitative study of Perez el al. [24], in this qualitative activity, in the 6 countries, a low emotional response is utterly preferred by most users, as they perceive the Virtual Assistant as an executor of tasks under the control of the user.

"Competence should trump emotionality as far as I'm concerned."-UK Participant-

"Is that if we leave emotions to machines, where does that leave us?"-Spain Participant-

"I do not believe that its function is to have emotions...I don't even believe that it can have them, is a machine...It does not have to understand our mood or give us advice about how to get out of depression. I think its function is a lot more practical: data, information."
-Chile Participant-

However, it is important to highlight that is low emotionality, not null. In Perez et al. [24] it was described that users needed first rational personality attributes such as Intelligent, Objective, Logical, Reliable, to be fully developed before evolving close and, even more, emotional attributes like kind, Considerate, Cheerful or Sweet. Users want then the VA to behave emotionally like a machine but not in excess, as a degree of empathy is expected to understand and adapt the response to the needs of the human.

"The minimum emotion not to be cold as a machine, even though that is what it is" -Argentina Participant-

"I feel like as an AI Aura doesn't need much emotions, I don't feel as though she needs to come across cold however, I think as humans we connect through emotions and a little emotionality is okay."-UK Participant-

Unsettling feelings where even expressed on the notion of the VA having so much emotionality, and therefore approaching too much to resembling a human, showing once more that the Uncanny Valley does apply into a Virtual Assistant's personality:

"Aura should be close, but always objective. Focus on its job. To me it would be awkward if it flirted with me."-Chile Participant-

"Empathy to recognize our requests and our mood to adapt the type of answer. Emotions no, because as we say, as the Virtual Assistant that it is, we all understand that it cannot feel anything. I do believe as well that being empathetic can be close enough without needing for it to share anything of its personal feelings, which is even weird." -Spain Participant-

"Nothing of the sort of saying very personal things like I love you or being too individual, saying things like 'I exist only for you', 'only helping you makes me happy' or 'I was born because and for you' haha. Those things make me back out."-Spain Participant-

Virtual Assistant's personality, consequently, faces again a complicated search for balance when it comes to level of emotion it must portray, which is exactly where the principles of the Uncanny Valley reside.

The analysis of the internal pilot and virtual communities' transcripts also show that users do not take NO for an answer. Even if the Virtual Assistant does not have the knowledge or does not want to commit to an answer, they do not tolerate a reply such as "I can't answer that question". The constant repetition of this response takes the VA personality too close into the machine side, hence generating a negative response. This finding shows that, as the Uncanny Valley theory presents, the level of likeliness of the VA increases as it distances itself from aspects that emphasise its machine condition. This finding is clearly remarked by the qualitative analysis made of the feedback given by the pilot users (results shown in Fig. 2), where the majority stated that the answer that they liked the least was not receiving any, even over receiving an incorrect answer.

The fact that the Virtual Assistant tries to give a solution to the user is highly valued. Users appreciate when the Virtual Assistant makes an effort in order to provide an answer to the user, even if it means failing to provide a useful response. Both in the pilot and in the virtual communities, users who expressed an acceptance around the Virtual Assistant not being able to answer everything correctly, do expect it to provide them with alternatives, such as doing a web search, ask to repeat the question, or provide a generic answer like for example if the Virtual Assistant cannot provide information about the weather, it could answer "I cannot give you information about the weather, but you should always carry a jacket just in case". Once again, a touch of wittiness, of *humanness* is being expected from the VA's personality.

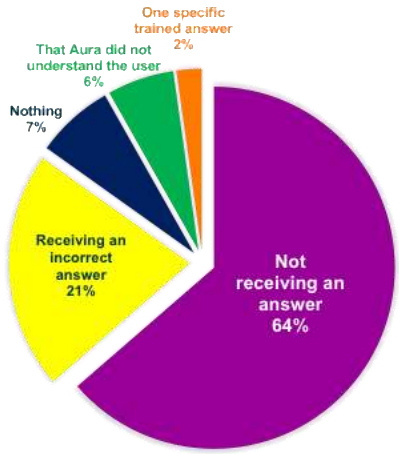

Fig. 2. What users disliked the most about the pilot trial with Aura.

An interesting fact that came up both in the internal pilot and the virtual communities is the that users appreciate a very human-like characteristic in Aura, and it was the fact that the VA is humble about the fact that it does not know it all and can't do it all, and that it acknowledges that it is learning. In the pilot survey 7% of the users chose as their most liked answer the response that Aura gave when it did not know an answer, and the reason was because it asked the user to provide its hoped reply in order to continue learning.

In the Virtual communities UK participants suggested that when Aura did not have an answer it should reply with phrases like: "I'm in constant evolution, I never cease to learn", or "I don't still know that but I'm learning. Can I make a search for you?". Moreover, a Brazilian participant remarked that "*Sobre sua personalidade achei um pouco sisuda... Mas deu para perceber que Aura tem interesse em aprender mais,*

queria poder ajuda-la." ("About its personality, I found it a little serious…But I could tell that Aura was interested in learning more, I wanted to be able to help".) This unpretentiousness then, leaning again away from a machine-like response, without renouncing to its true condition, is appreciated by users and therefore shows once more a support that the Uncanny Valley operates on a Virtual Assistant's personality.

5.2 Voice

This study has shown there are seven aspects that need to be taken into account when developing a Virtual Assistant voice, to project its personality. Table 3 shows them in order of importance.

It is already relevant in the previous results, that within the seven aspects that are vital for the VA's personality, gender, speed of the voice and human vs. machine resemblance of the voice are the most closely related aspects to the Uncanny Valley, showing the implication of this theory in the construction of a VA's personality. We discuss the findings in each of these three aspects that support this claim.

Table 3. Components for a successful Virtual Assistant voice

Gender of the voice
Speed of the voice
Machineness vs. Humanness
Timbre of the voice
Tone of voice
Personalization of the voice depending on the user
Coherence of the same voice across devices

Gender. The literature review had shown controversy around the gender of the VA's voice and the different perceptions users have when having a feminine or masculine voice. Consequently, during the user tests made in Spain different voices for Telefónica's Virtual Assistant, Aura, were presented to the users with a wide range of voices both feminine and masculine to choose from and to explain the reasons behind their choice. All of the users that performed Aura's voice user test preferred to have a feminine Virtual Assistant's voice as they considered it more trustful, efficient and warm, which supports previous research [42].

Users even stated they would prefer a less quality feminine voice (more mechanical, that makes mistakes, etc.) rather than a good quality masculine voice. The reason behind this radical selection independently of quality was generally claimed to be "because they were used to it". These results raise questions about how the habit of having a female voice is affecting users' preference, and the potential need of creating Virtual Assistant with both feminine and masculine voices in order to allow users to choose and enable a more personalised experience, which several users suggested.

The preference towards a feminine voice, and reasons towards it, was also shown in the Virtual Communities in the other countries, with the exception of UK. Users relate feminine voices to the IVR systems when they make a call to a service provider and are used to associate them to machines. In this particular research activity, the users would receive Aura with the voice and voice gender in which the device was programmed by default, and only in the UK, the only country where a Virtual Assistant, Siri, has a default masculine voice, a debate was generated towards a Virtual Assistant having feminine or masculine voice. The other five countries in the study, Germany, Brazil, Chile, Argentina and Spain, where the default voice of the major Virtual Assistants is currently feminine, did not enter in controversy. In fact, masculine voices resembled more human for users so that was the primary reason to reject them and instead choose feminine voices that made clearer for users they were talking to a machine. This result increases the interest about the role that habit plays in the voice gender preference, giving room to encourage further research to examine this impact around the uncanny valley phenomena.

Machine vs. Human. Regarding specifically the machine vs. human voice, users in both tests and virtual communities unanimously expressed that the voice must not be robotic or mechanical; they do not want to feel like they are talking to a machine, but not to an extreme that they can't decipher whether or not they are talking to a real person, as they would feel they have been cheated. This once again projects the Uncanny Valley theory in a Virtual Assistant's personality by demonstrating that users reject VA voices that are too similar to humans. Nonetheless, users expect to have a fluid conversation in which the Virtual Assistant can be recognised as a reliable interlocutor, able to understand the context of the conversation, demonstrate intelligence and generate empathy, while assuring the user they are talking to a Virtual Assistant, instead of a person. It is precisely the ability to hold a conversation, which differentiates Virtual Assistants from a remote control or a web browser, as they transcend and go beyond the command and execution of functions, therefore the approach towards a more human-like conversation is expected. The balance between human and machine in terms of voice elements yet again is wished by users that mark the VA's personality.

Speed. The user tests shed light onto the crucial importance of the speed for the VA's voice, because, as users exposed, it helps project intelligence onto the Virtual Assistant, as it seems less smart and incapable when it speaks slowly. Users have prioritised this key aspect of Virtual Assistant voices when it comes to selecting their preferred one (they consistently rejected all slower voices). A slow voice impacts on the user perception at different levels: this characteristic is linked to the perception of overly robotic voice, and it also makes users think they are wasting their time and would be quicker to do the task by themselves. Speed is, therefore, a true deal breaker when it comes to a satisfactory and efficient user experience.

6. Conclusions

The world of Virtual Assistants' adoption, use and development is growing really fast and it is no longer only about performance, efficiency and effectiveness. Users expect a Virtual Assistant to project an interesting personality that is able to delight them during their interaction with it. It is not something they consciously require, but something that emerges when a voice comes into play. This study's objective was to explore two core constructs of a Virtual Assistant personality, voice and words, and identify what range of *humanness* vs. *machineness* do users expect and like, in order to explore if the Uncanny Valley theory applies to Virtual Assistants.

As soon as there is a voice involved in an interaction, humans tend to attribute a personality to that piece of technology. This personality is conceived under the responses the Virtual Assistant gives back. This includes both the content, the words that it uses to reveal a certain type of personality, as well as the voice, that has also been demonstrated to be an important reflection of the personality. Through several research activities, the timeliness and relevance of the Uncanny Valley theory was been shown in both these VA personality constructs, by which users tend to increase the likeliness when the personality separates from being too much machine-like, but also reduce that likeliness as it approaches too much into pretending to be a human being in the way it speaks and behaves. It could be concluded then, that as with other AI applications, users want Virtual Assistants, and specifically their personalities, to reach a balance where their feelings towards the VA are not those of talking to a machine, but neither falling into the Uncanny Valley where the VA resembles too much a human personality without being able to fully accomplish it. There is then a lack of interest in customising a Virtual Assistant to the point that resembles a human being, both in the way it talks, what it says and the messages that come across, and how it talks, its voice. Users have consistently requested being able to easily identify the voice as an AI application, and

do not expect the Virtual Assistant to provide comments of its preferences in diverse very-human subjects, nor have an identity as if it were a real person.

In sum, this study has pointed out what are the most sensitive aspects regarding the words content and the voice characteristics that play a key role in the VA personality and its relationship with the Uncanny Valley theory.

7. Further Research

The results gathered from this exploratory research have established some relevant research routes to keep deepening the knowledge around Virtual Assistants personality and the core constructs to build a pleasant user experience. For instance, the role that gender plays in the Uncanny Valley theory, if a neutral voice is posible in order to achieve a balance, and how this would influence the personality perception of the VA. Additionally, it would be necessary to replicate this study in a wider set of countries, in cultures that were not included in this one, to be able to evaluate if the conclusions can be applied globally. The Asian culture is of special research interest for its highly anthropomorphic technology adoption, typically seen in the launch of diverse human-like robots. Therefore, it would be really relevant to consider their approach of the uncanny valley in the Virtual Assistants panorama. Moreover, it would be noteworthy to also explore the USA scenario where the VA adoption has been more pioneer, in order to investigate if time and use have an effect on the personality expectations and perceptions. We hypothesise if perhaps a more close and emotional relationship with technology develops overtime and consequently, a balanced personality between *humanness* and *machineness* is still part of the users demands, or it evolves into a more intimate one. The subject of VA personality and its relationship to users is a field that will only endure and grow with time, therefore demands to be continuously explored.

References

Luger, E., Sellen, A.: "Like Having a Really bad PA": The Gulf between User Expectation and Experience of Conversational Agents. In: Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, pp. 5286-5297. ACM, New York, N.Y. DOI: 10.1145/2858036.2858288 (2016).

Stanford University Personality and Voice in AI: Creating AI Conversations Panel Series, Stanford, California, January 24, 2018. (2018).

Wilks, Y.: Is a companion a distinctive kind of relationship with a machine? In Proceedings of the 2010 Workshop on Companionable Dialogue Systems (CDS '10), pp. 13-18. Association for Computational Linguistics, Stroudsburg, PA. (2010).

Haven Martinson, L.: Personality and Voice in AI. Creating AI Conversations Panel Series, Stanford, California, January 24, 2018. (2018).

Reilly, W.S.: Believable Social and Emotional Agents (PhD Thesis). School of Computer Science, Carnegie Mellon University, Pittsburgh. DOI: 10.1145/267658.267683 (1996).

Dave Moffat. Personality parameters and programs. In Trappl, R. and Petta, P. (eds.), Creating personalities for synthetic actors. Springer, Berlin. DOI: 10.1007/BFb0030575 (1997).

Johnson, K.: People, not tech companies, should pick their AI assistant's personality, https://venturebeat.com/2017/10/02/tech-giants- should-let-people-pick-their-ai-assistants- personality/, last accessed Retrieved 2018/02/28

Cowan, B.R., Pantidi, N., Coyle, D., Morrissey, K., Clarke, P., Al-Shehri, S., Earley, D., Bandeira, N.: What Can I Help You With?: Infrequent Users Experiences of Intelligent Personal Assistants. In: 19th International Conference on Human-Computer Interaction with Mobile Devices and Services, Art. 43. ACM, New York, NY. DOI: 10.1145/3098279.3098539 (2017).

Nass C. Lee, K.M.: Does Computer-Synthesized Speech Manifest Personality? Experimental. Journal of Experimental Psychology: Applied, 7(3), 171-181. DOI: 10.1037//1076-898X.7.3.171 (2001).

Denes, G., Caldognetto, E.M., Semenza, C., Vagges, K., Zettin, M. Discrimination and identification of emotions in human voice by brain-damaged subjects. Acta Neurológica Scandinavia, 69, 154-162. DOI: 10.1111/j.1600-0404.1984.tb07794.x (1984).

Mullenix, J.W., Johnson, K.A., Topcu-Durgun, M., Farnsworth, L.M.: The perceptual representation of voice gender. Journal of the Acoustical Society of America, 98(6), 3080-3095 DOI: 10.1121/1.413832 (1995).

Cohen, M.H., Giangola, J.P., Balogh, J.: Voice User Interface Design. Addison-Wesley Professional, Boston. (2004).

Short, T.X.: Designing Stronger AI Personalities. In: 13th Artificial Intelligence and Interactive Digital Entertainment Conference, pp. 11-117. Association for the Advancement of Artificial Intelligence (2017).

Nass, C., Gong, L.: Social aspects of speech interfaces from an evolutionary perspective: Experimental research and design implications. Communications of the ACM, 43(9), 36-43. DOI: 10.1145/348941.348976 (2000).

Sherer, K.R.: Personality inference from voice quality- The loud voice of extroversion.
European Journal of Social Psychology, 8(4), 467-487. DOI: 10.1002/ejsp.2420080405 (1978).

Mori, M.: The Uncanny Valley. Energy, 7(4), 33-35. DOI: 10.1109/MRA.2012.2192811 (1970).

Pollick, F.E.: In Search of the Uncanny Valley. In Daras P., Ibarra O.M. (eds) User Centric Media. UCMEDIA 2009. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, (40), Springer, Berlin, Heidelberg. DOI: 10.1007/978-3-642-12630-7_8 (2010).

Wilks, Y.: Close Engagements with Artificial Companions. Key social, psychological, ethical and design issues. John Benjamins, Amsterdam, Netherlands (2010).

Gong, L., Nass, C., Simard C., Takhteyev, Y.: When non-human is better than semi-human: Consistency in speech interfaces. In Smith, M.J., Salvendy, G., Harris, D., and Koubek, R. (Eds.) Usability evaluation and interface design: Cognitive engineering, intelligent agents, and virtual reality, 1558-1562. Lawrence Erlbaum Associates, Mahwah (2001).

Funder, D.C.: The Personality Puzzle. Norton, New York (2001).

Aaker, J.L., Benet-Martinez, V., Gorolera, J.: Consumption symbols as carriers of culture: A study of Japanese and Spanish brand personality constructs. Journal of Personality and Social Psychology, 81(3), 492-508. DOI: 10.1037//0022-3514.81.3.492 (2001).

Suppehellen, M., Gornhaug, K.: Building foreign brand personalities in Russia: the moderating effect of consumer ethnocentrism. International Journal of Advertising, 22(2), 203-226 (2003).

Ekinci Y., Hosany, S.: Destination Personality: An Application of Brand Personality to Tourism Destinations. Journal of Travel Research, 45, 127-139. DOI: 10.1177/0047287506291603 (2006).

Perez, M., Saffon, S., Donis, H.: Voice Activated Virtual Assistants Personality Perceptions and Desires: Comparing Personality Evaluation Frameworks. In: Proceedings of the 32nd British Human Computer Interaction Conference DOI: 10.14236/ewic/HCI2018.40 (2018).

Flores, F.: Lo humano y lo artificial en la comunicación electronica. El Catoblepas: Revista Crítica del Presente, 12, 4. (2003).

Mehrabian, A., Weiner, M.: Decoding of inconsistent communications. Journal of Personality and Social Psychology, 6, 109-114. DOI: 10.1037/h0024532 (1967).

Mehrabian, A., Ferris, S.R.: Inference of attitudes from nonverbal communication in two channels. Journal of Consulting Psychology, 31, 248-252. DOI: 10.1037/h0024648 (1967).

Coupland, J.: Small Talk. Routledge, New York. (2000).

Malinowski, B.: Phatic Communion. In John Laver and Sandy Hutcheson (eds.) Communication in Face to Face Interaction. Harmondsworth: Penguin, 146-52. (1972).

Bickmore, T.W., Cassell, J.: Relational Agents: A Model and Implementation of Building User Trust. Proceedings of the SIGCHI conference, Seattle, USA, March 31-April 4, 2001, 396-403. DOI: 10.1145/365024.365304 (2001).

Nadler, J.: Rapport in Legal Negotiation; How Small Talk Can Facilitate E-mail Deal making. Harvard Negotiaton Law Review, 9, 223-251. DOI: 10.2139/ssrn.428523 (2004).

Knowlton, G.E., Larkin, K.T.: The Influence of Voice Volume, Pitch, and Speech Rate on Progressive Relaxation Training: Application of Methods from Speech Pathology and Audiology. Applied Psychophysiology and Biofeedback, 31 (2), 173-85. DOI: 10.1007/s10484-006-9014-6 (2006).

Tigue, C.C., Borak, D.J., O'Connor, J.J.M., Schand, C., Feinberg, D.R.: Voice pitch influences voting behavior. Evolution and Human Behavior, 33, 210–216. DOI 10.1016/j.evolhumbehav.2011.09.004 (2012).

Wong, K.: In pursuit of the perfect AI voice: How developers are humanizing their virtual personal assistants, https://www.engadget.com/2018/04/09/in-pursuit-of-the-perfect-ai-voice/ last accessed 2018/04/09.

Sokhi, D.S., Hunter, M.D., Wilkinson, I.D., Woodruff, P.W.R.: Male and female voices activate distinct regions in the male brain. Neuroimage Journal, 27(3), 572-57. DOI: 10.1016/j.neuroimage.2005.04.023 (2005).

Siegel, M., Breazeal, C., Norton, M.I.: Persuasive robotics: The influence of robot gender on human behaviour. International Conference on Intelligent Robots and Systems, pp. 2563-2568. DOI: 10.1109/IROS.2009.5354116 (2009).

Mitchell, W.J., Ho, C.C., Patel, H., MacDorman, K.F.: Does social desirability bias favour humans? Computers in Human Behaviour 27(1), 402–412. DOI: 10.1016/j.chb.2010.09.002 (2011).

Nass, C., Brave, S.: Wired for Speech: How Voice Activates and Advances the Human-Computer Relationship. MIT Press, Cambridge (2007).

Nass, C.: Studying Computers to Learn About Ourselves, https://www.npr.org/templates/story/story.php?storyId=129629756&t=1537204594966&t=1537434941574 last accessed 2018/08/13.

O'Shea, D.: Ikea: What do shoppers want in artificial intelligence?, https://www.retaildive.com/news/ikea-what-do-shoppers-want-in-artificial-intelligence/441911/ last accessed 2018/07/25.

Statista: Worldwide intelligent/digital assistant market share in 2017 and 2020, by product, https://www.statista.com/statistics/789633/worldwide-digital-assistant-market-share/ las accessed 2018/11/19.

Perez, M., Saffon, S.: Building trust between Users and Data Driven Virtual Assistants. 14th International Conference on

Artificial Intelligence Applications and Innovations (AIAI), pp. 628-637. DOI: 10.1007/978-3-319-92007-8_53 (2018).

Don Norman. The Design of Everyday Things. Basic Books. (2013).

This study was published in:
Dr Marta Pérez García and Sarita Saffon López. Exploring the Uncanny Valley Theory in the Constructs of a Virtual Assistant Personality. Intelligent Systems Conference 2019 (IntelliSys 2019) and Proceedings of SAI Intelligent Systems Conference (pp. 1017-1033) Springer, Cham.