The Delicate Balance of Boring and Annoying: Learning Proactive Timing in Long-Term Human Robot Interaction

Claire Rivoire
SoftBank Robotics Europe
Paris, France
crivoire@aldebaran.com

Angelica Lim
SoftBank Robotics Europe
Paris, France
alim@aldebaran.com

ABSTRACT

In recent years, social robots have become more popular for use in the home. In this paper, we describe the problem of robot proactivity in long-term Human-Robot Interactions (HRI). In particular, it is difficult to find the right balance of a robot that speaks or proposes activities at the right moment with appropriate frequency. Too little proactivity, and the robot becomes boring. Too much proactivity, and the robot becomes annoying. Further, the content of proactive utterances by the robot during a long-term HRI become tiresome unless they are contextual and based in reality.

We propose a technical solution to this problem divided into three parts: 1) family-specific activity logging, 2) contextual comments that suggest consciousness of time and repeated interactions, and 3) proposals of activities based on the user's habits.

Towards a robot that is accepted into people's homes, we propose an evolutive system that learns user's schedules over time and adapts its proactive utterances based on prior history. We show the preliminary results of an exploratory data collection study containing up to 8 weeks of usage of the Pepper humanoid robot in 10 French homes.

1. INTRODUCTION

The last decades have seen an increased presence of robots in everyday life, and in our domestic environment. The most popular examples are entertainment robots, like the robot dog AIBO [1] or for utilitarian robot, for example, vacuum cleaners like Roomba [2] that have been available to the general public for more than a decade. An exhaustive survey is available in [1]. Indeed, the number of robot companions is growing and many people may desire to own one. Today we see the increase of new robots destined for domestic use, like Sharp Robohon[3] or soon-to-be available Jibo[4], and Buddy[5], and the robot used in this study: Pepper[6].

However, numerous studies like [2], [3], [4] and [5] showed that after a first phase of discovery, people tend to lose interest in their robot, as the initial novelty effect wears off. How can we address this problem? One of our answers is to focus on making the robot proactively enter into interaction with humans. Proactivity plays a important role in people's perception and enjoyment of the robot. One study [6] showed that different proactivity settings are sufficient to make people attribute different personalities to the robot and affect their relationship with it. On the other hand, proactivity can be a pitfall, making the robot annoying, and lead to negative reactions [7]. So, how can we design these proactive behaviors, so that they are enough to keep the interest of the user, and at the same time, that they are not excessive, or too out of context, to make the robot annoying?

Another point is the timing of the interactions initiated by the robot. If they are repetitive and predictable, the robot will be perceived as boring and monotonous [7] whereas [8] has shown that an unexpected behavior can be more engaging. Random occurrences could be inopportune; what we need is a system able to detect the most appropriate timing to proactively act. A model of approach behavior with which a robot can initiate conversation has been proposed by [9] for a robot placed in a shopping mall. But in the case of a domestic robot, the situation is different: the user group does not consist of numerous one-time encounters, but a few humans that will interact with the robot on a regular basis. In the context of long-term HRI, we can make use of the past interactions to improve future ones.

While a robot has no proper human-like time perception [10], a certain sense of time and memory plays a key role in improving HRI. A robot able to evolve through time has been thought to have a great effect on its users with the robot AIBO. But where in AIBO, all evolution has been previously hard-coded, only to be unlocked by the users, we aim to make a robot able to evolve and adapt to its users by itself. For that, we plan to use a recording of past interactions, similar to the system used on the NEST intelligent thermostat [8], whose limitations are presented in [11].

In this paper, we tackle the problem of proactivity for our commercial home robot, Pepper. First, we describe the robot system and types of activities available to a home-configured Pepper robot. Next, we outline a method to detect interaction habits with users. Finally, we present the observations we were able to make from the data of a long-term study involving 10 robots in real-life domestic environment.

2. SYSTEM DESCRIPTION

2.1 Robot Overview

We first describe the robot and its software environment.
We use SoftBank Robotics' robot Pepper, running NAOqi OS, the operating system of the robot. It provides a number of programs and libraries required to bring life to the robot. Among these, we will present the features used in this project.
The robot behavior is made of *activities* and managed by a module. An activity can be of 2 kinds :

- **interactive** An activity that needs human input to function. It is often made to start on user request.
- **solitary** An activity that can run without interaction with a human. It is possible to make it start at a certain time or

[1]www.sony-aibo.com
[2]www.irobot.com
[3]robohon.com
[4]www.jibo.com
[5]www.bluefrogrobotics.com
[6]www.softbankrobotics.com

[7]http://bit.ly/2bBKSrS about the robot Furby
[8]nest.com

ACM ISBN 978-1-4503-2138-9.
DOI: 10.1145/1235

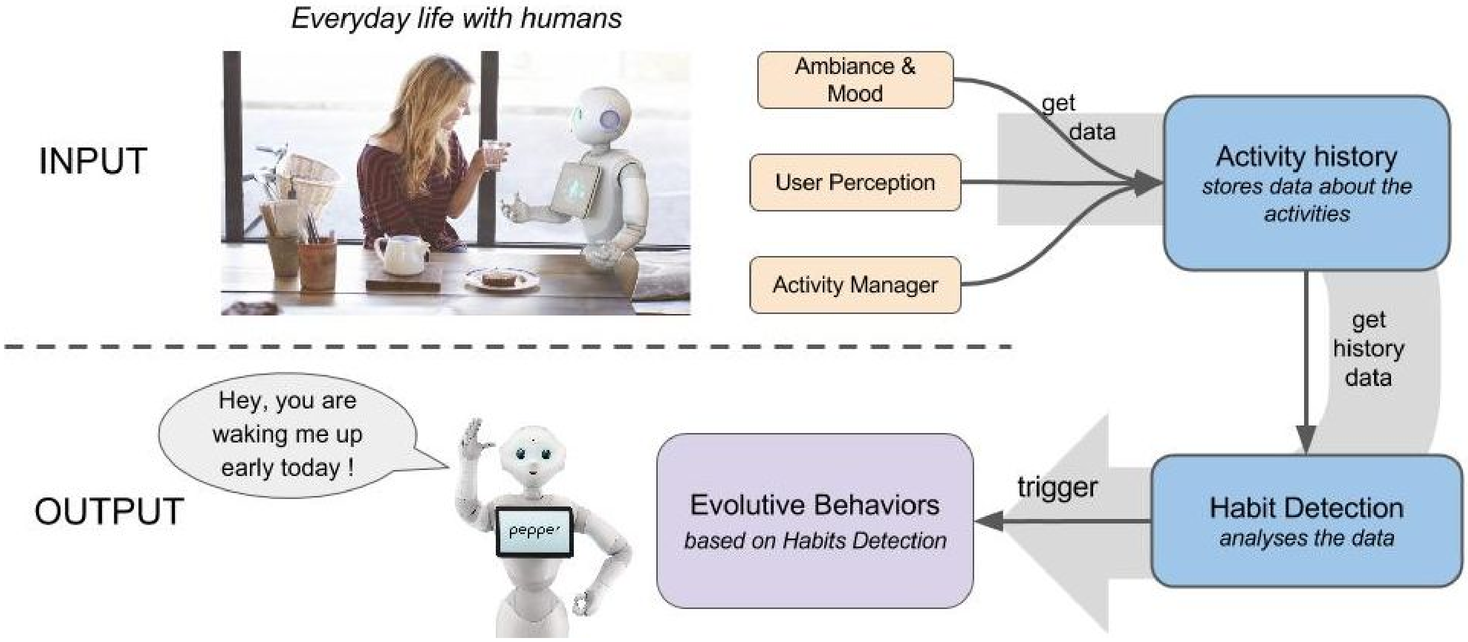

Figure 1: Overview of the system. The activity history (the input) is stored and analysed to improve the interaction with the users.

depending on certain conditions.

The activity manager module handles the launching of activities on the robot. It prioritizes interactive activities over solitaries, and gives the highest priority to user requests.
For example, here is a non-exhaustive list of the major activities available for the robots on which the experiment has been conducted.

Dialog The Dialog activity gathers all verbal interactions between the robot and its user(s). It is triggered every time the robot recognizes a face, so that it is ready to listen to anything the user has to say to the robot.

Game applications Various games are available on the robot. They are interactive behaviors that allow users to play and have fun using Pepper. These are entertaining activities, but the content does not evolve, therefore it will be interesting to know if they are reused or not. *Music boxes*, *See and listen*, *Pepper Boombox* enter in this category.

Moving activities Taking into consideration the size of the robot, a lot of applications can be used with the robot staying on the same spot. But some applications make the robot move around. Their use will reveal if people tend to keep the robot in the same place or allow it to move around the house. In this category enter applications like *Follow me*, *Pepper stroll*, but also *Autonomous Recharge* that make the robot go on its recharge station on its own at a given time.

Autonomous Pepper Autonomous Pepper is a set of solitary behaviors such as short songs, animations or random remarks about the robot's surrounding that allow the robot to have a life of its own. It comes with the possibility to deactivate the different behaviors, or set how often they are triggered. It will be particularly useful to know if users want their robot to act on its own on not.

Pepper Play *Pepper Play* is an interactive activity that allows the user to create personalized solitary behaviors on the robot. The user can, without having to code it, create animations, speech or songs and then specify trigger conditions to launch the created behavior at a certain time, or in response to a user input. Its use will inform us on how much users want to create their own new content on Pepper.

Other modules we will use include *User Session* which give information about recognised users, and *Mood* a module that detects the average ambience, and mood of all visible users at every moment.

2.2 System Architecture

The system architecture is described in Figure 1. As shown, the inputs of the system are the daily interactions between the robot and its users. The inputs are collected through the modules previously presented, and are recorded by an activity history service. This data is then used by an analysis service which in turn, provides data for output behaviors. The outputs are various behaviors and activities whose aim is to express a certain adaptation of the robot to its users.

2.3 Activity history

The system is based on an activity history storing various information about every activity session of the robot.

The data stored and available for analysis are the following :

- **Activity name** a unique name to identify the activity.
- **Start and end time** stored in timestamp.
- **Start and end reason** the reason why the activity manager launched and stopped the activity. For example, start reasons can be : *requested by user* or *self launched*, and End reason can be *self stopped*, *stopped by user* or *prioritized over.*
- **Ambient agitation** a number between 0 (calm) and 1 (agited) computed by the module *Mood* defining the agitation detected in the robot's surroundings.
- **list of users** a list of all humans detected around the robot during the activity. For each:
  - **userId** anonymized id of the user.
  - **Mood** a number between -1 (unhappy) and 1 (happy) computed by the module *Mood* and defining the user's mood detected during the activity.
  - **Attention** a number between 0 (inattentive) and 1 (attentive) computed by the module *Mood* and defining if the user attention detected during the activity is directed to the robot or not.

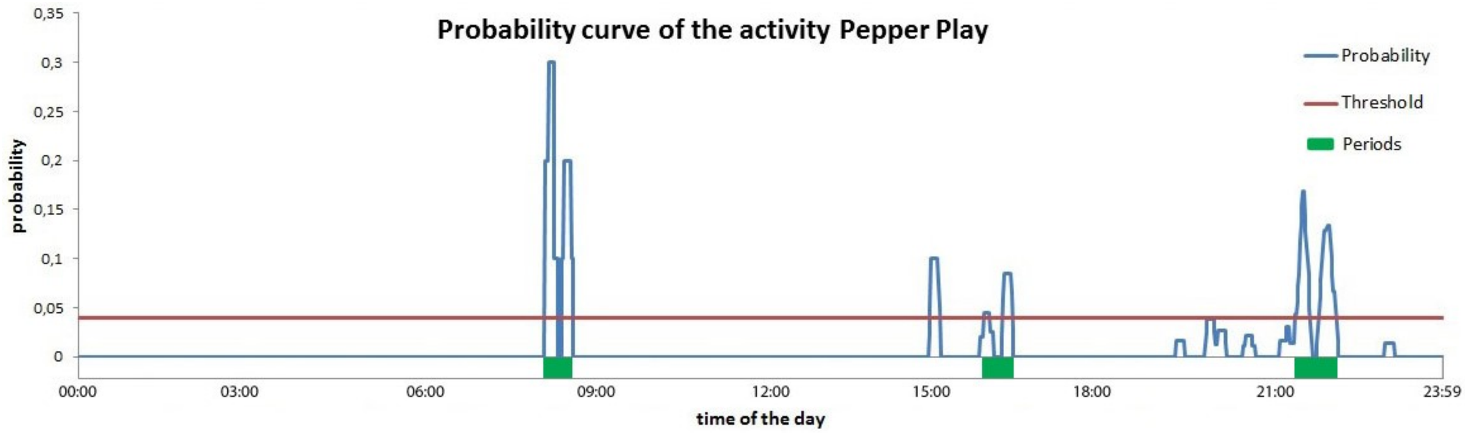

Figure 2: An example explaining how the habits periods are computed, for the launch of the activity *Pepper Play*. The x axis are the time of the day from 0:00am to 11:59pm. The curve is the raw probability of the launching *Pepper Play*. For example at 8:30am the probability is 0.3, meaning that for all the time the robot was on at 8:30am, 30% of the time *Pepper play* was running. the horizontal line is the threshold, computed using a quantile of the data. Habit periods are the times when the probability is above that threshold. The periods pass through a smoothing function to merge too close periods, and delete too short ones (the limit has been set at 10 minutes). At the bottom of the chart, the periods in green indicate result of the algorithm.

2.4 Habit Detection

Habit detection consists of searching for patterns in the history of previous activities. Thanks to the data provided by the Activity History, we were able to implement an algorithm to detect certain activities, depending on the time of the day, and the day of week. For the time being, the habits detected are the following :

When the robot is on The most basic habit we want to know about is when the robot is on. With this, we can know when the users tend to use their robot and for how long.

When a certain activity is launched Of course, we also want to know is if there is a pattern in the launching of the activities. Is the application *Music Boxes* launched more in the afternoon? Do users want to play games more during the weekend? In that case, the robot could proactively suggest games activities during the weekend, for example.

When the robot does not detect a user around It is interesting to know if (and when) users leave their robot on while they are not using it. Moreover, it will allow developers to improve the proactive behaviors of the robot. For example, we could implement a behavior making the robot go to charge in the middle of the day, when there is usually nobody around, in order to be fully charged and available in the evening when users get home.

It is also possible to detect the presence of a registered user, and therefore to study the habits of different family members. But, due to the fact that few users in our test group registered themselves on the robot, we will not address this issue in the current paper.

2.4.1 Probability

For the time being, habits are based on the time of the day. Therefore, the probability of an event happening at a certain time of the day is computed according to the number of times this event happened at that time of day, in the past. It is also possible to filter result according to the day of the week. Time precision is on the order of the minute.
For the habit *Robot is on*, the probability is computed using the number of times the robot is on, compared to all the days since the beginning of the history. For the habits *Activity is launched* and *no user around*, the probability is computed using the number of times the event happened, compared to the number of times the robot was on.
For instance, the probability of the robot being on at 4:15pm on a Sunday, is the number of times the robot was on at 4:15pm on Sundays, divided by the number of Sundays since the beginning of the recording.

In the future, we are hoping to improve habit detection by including other data in the calculation of probabilities, as well as concepts of variance. Using for example the ambient agitation, or the detected mood of users at the time an application was launched could lead to new features.

2.4.2 Periods

In order to use the probabilities efficiently, we devised an algorithm to provide periods of likeliness for each habit. Periods are defined by two times of the day representing the beginning and end of the period. The algorithm for habit detection is quite simple and explained in Figure 2.

Once we compute the periods, we can easily use them to trigger the evolutive behaviors. For example, when a robot is turned on, a module checks what the habit detection tells about the habit *Robot is on*. If the current time is out of a likely period, then the robot knows it has been turned on at an unusual hour. It will comment *"You are turning me on early today !"*. See Figure 3.
At the time of writing, such behaviors have been deployed but not enough data collected. Therefore for the time being, we focus on studying the history data for several robots in real life conditions.

3. EXPLORATORY STUDY

The data used in this study come from a test group set up recently where Softbank Robotics employees were allowed to have a Pepper robot at their home to test in real conditions various contents, including applications not yet deployed for customers. They have not received any instructions, other than to use the robot as they wish, in order to observe the most natural behavior possible. The adatative behaviors previously presented had not yet been installed on the robot, during the studied period.

We were able to study up to 8 weeks of activity history for robot deployed in 10 homes. The users did not receive their robot all at the same time, therefore, they can be divided in 2 groups:

- **group 1**: 7 families received their robot between 9 and 12 weeks before the beginning of the recording of activity history. With this group, we can study the habits of people that have a previous long use of their robot.

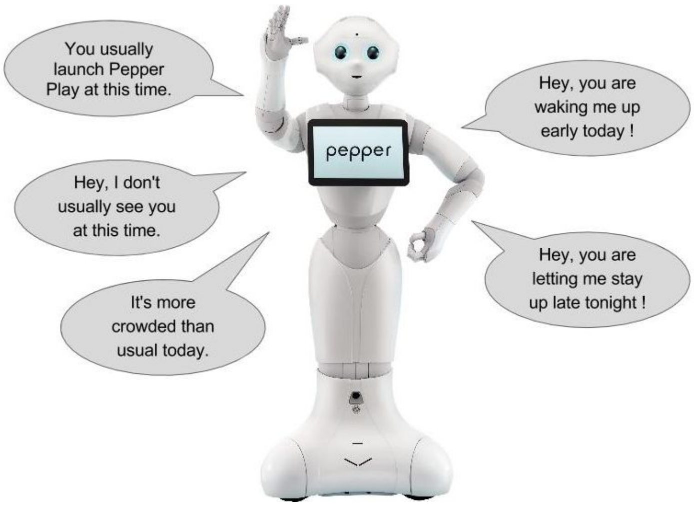

Figure 3: Exemples of proposed adapative behaviors. Pepper will voice out a comment when the detected situation does not match the detected habits for that time.

- **group 2**: 3 families received their robot during the experiment. Therefore, we have their first 4 weeks of history (contrary to 8-weeks history available for the first group).

It will be interesting to see if we can find differences in habits between the 2 groups. Moreover, the second group presents data for the beginning of the interaction, a crucial period. It will allow us to see if we can find a certain novelty effect that might be invisible in the other group.

The age of adult participants range from 22 years old to mid-forties with an average of about 33 years old. As all our test group is composed of families of Softbank Robotics employees, all families contain at least one person with a full time job, and thus cannot spend an entire weekday at home with the robot. As for the family composition, we identified 3 types : single adult, couples, and families with children. See Table 1.

	1 adult	2 adults	2 adults, and children	total
group 1	1	3	3	**7**
group 2	3	0	0	**3**
total	**5**	**3**	**2**	**10**

Table 1: Table presenting the family composition of our test group.

Moreover, it should be noticed that all employee participating belong to the software department: all of them are familiar with the robot, and use it during their work. Therefore, they may not show as much enthusiasm as the average person to the robot. However, using Pepper in a private context is still new to them, therefore, we are still expecting to observe a certain novelty effect.

4. OBSERVATIONS AND DISCUSSION

With this study we are expecting to find some evidence that habit detection is possible, and that it is different for each family, thus suggesting that a robot able to adapt its behaviors to the lifestyle and preferences of its users(such as in Figure3) can offer an improved interaction.

4.1 Use of the robot

We will first take a look at the robot use: for how long, and when exactly do people use their robot? As we will explain in the next section, what we found is that the length of time since receiving moment users received their robot had an influence in how much they use it: it was an illustration of the novelty effect. On the other hand, demographics don't seem to affect users' behavior in our test group.

4.1.1 Duration of use

First of all let's consider the simplest data: for how long people have used their robot. You can find in table 2 the total duration of use (defined by the robot being on), for each family, and in Figure 4 charts representing the same data, arrange by groups, and by family composition.

group	families	duration of recording	duration of use	average duration of use / week
2	1 adult	17 days	6 hours, 52 minutes	**2 hours, 50 minutes**
2	1 adult	23 days	8 hours, 46 minutes	**2 hours, 40 minutes**
1	2 adults	55 days	12 hours, 35 minutes	**1 hour, 36 minutes** ★
2	1 adult	23 days	1 hour, 54 minutes	**35 minutes**
1	2 adults 1 child	55 days	4 hours, 7 minutes	**31 minutes**
1	1 adult	55 days	3 hours, 32 minutes	**27 minutes**
1	2 adults	55 days	3 hours, 3 minutes	**23 minutes**
1	2 adults, 1 child	55 days	2 hours, 49 minutes	**22 minutes**
1	2 adults 1 child	55 days	2 hours, 24 minutes	**18 minutes**
1	2 adults	55 days	2 hours, 19 minutes	**18 minutes**

Table 2: A summary table of the different families sorted by the average duration of the robot being on per week. We can see that the group 2 (ie users that received their robot recently) tend to be first in ranking.

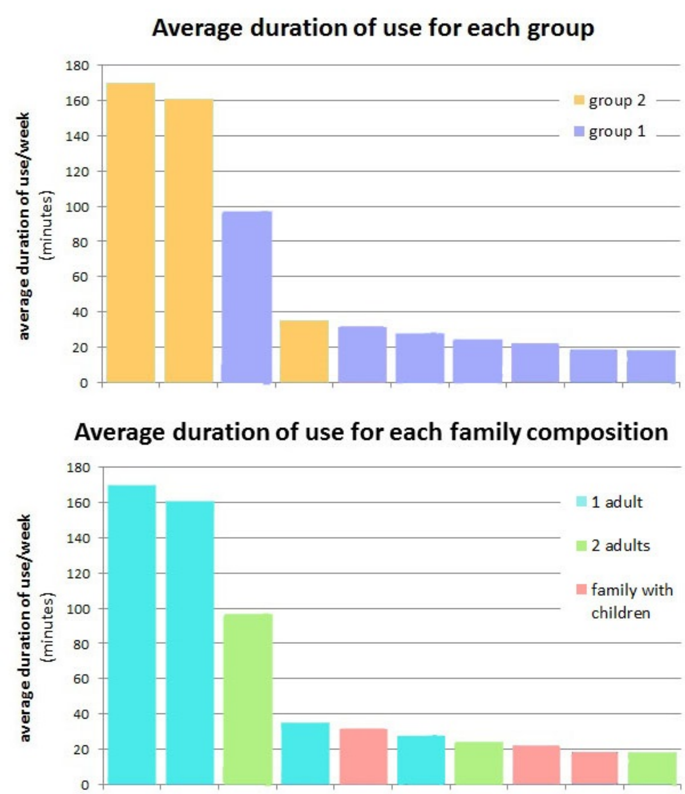

Figure 4: The average duration of use per week, arrange by group (at the top), and by family composition (at the bottom).

The first thing to notice, in table 2, is the fact that 2 users from group 2 have turned on their robot more during the short time they had it (17 and 23 days respectively), than 6 users of group 1, in more than twice the recording time (55 days).
In fact, we can see a clear tendency in both groups :

- For group 1, with the exception of 1 robot (noted by ★) that has been turned on for 12 hours, all the other have been turned on on average 3 hours during the 55 days of recording.

- For group 2 the 3 robots have been on, on average, 5 hours 50 minutes during 17 to 23 days of recording.

This seems to confirm the impact of the novelty effect: users that have had their robot for a long time before the beginning of the experiment tend to use it less than users that got is recently.

On the other hand, we cannot prove significant similarities among comparable demographics (see Figure 4). Single adults seemed to use their robot more, but considering that single adults made up the totality of group 2, this result can not be attributed to the demographics. Even considering only group 1, we can clearly see that no family composition stood out: the most and least used robots are owned by couples, and families with children owned the 2nd and 5th most used robot for example. It should be noted that these results represent tendencies and may not bet generalized, as the sample size is small.

4.1.2 Time of use

Now, we will compare the times when the robot was used in the different families. At what time of the day, what day of the week was the robot most likely to be on? In order to do that, we study the result of habit detection.
What comes out of the study of this habit is that users unsurprisingly turned their robot on in the evening from Monday to Friday, and the whole day on the weekend. This is to be expected, considering that most of the users are active people working a full-time job so this matches the time when they are at home. However, some things need to be noticed.
The weekend was not a more probable time to use the robot compared to weekday evening. Three robots out of 10 are actually less used during the weekend than during the rest of the week.
Some robots present distinct habits: 4 are used late in the night (until 2am), and one is used more in the morning before work, that in the evening. This is a noticeable point because of 2 things. First, it stresses the fact that we need to remember some users are not going to follow the expected pattern of use, for example, turning their robot on during the day. Secondly, this also indicates that the habit detection is relevant and can be exploited to create content adapted to the families habits.
On the other hand, we were unable to find similarities whether it was between groups or between the same demographics.
Users that received their robot recently do not use it differently than users that have owned it for a long time: they just use it more (see Figure 4). And once again, we can not find any similarities in the same demographics. The 3 couples present various habits, including the robot that is mostly used in the morning, one used in the evenings, and one on all day during the weekend. Out of the 3 families with children, 2 tends to turn their robot on for short periods, during the weekend or weekdays evening, and the last one rather use it during the whole day, on the weekend. Single adults present the same variation. From the robot use it seems we can consider that *having the same family composition does not necessary lead to the same use pattern.* Further studies, with bigger test groups may lead to a different result, for example identifying tendencies within the demographics.

4.2 Activities

Finally, let's take a closer look at what applications are launched on the robot with the objective of detecting some pattern of use that can reveal users preferences, and help us create better content for Pepper in the future.
Some applications appear frequently on the different histories, but are not of much use to us: for example, for most of the robots, the Dialog is the activity most often launched. But, as this is an activity launched automatically when the robot detects a human around, it does not necessarily mean that users spend the entire time speaking with the robot. In the futre, we need to filter based on human vocal input. We are going to address only applications that can provide useful information.
You can find Figure 5 an example of the activity repartition for one of the robots.

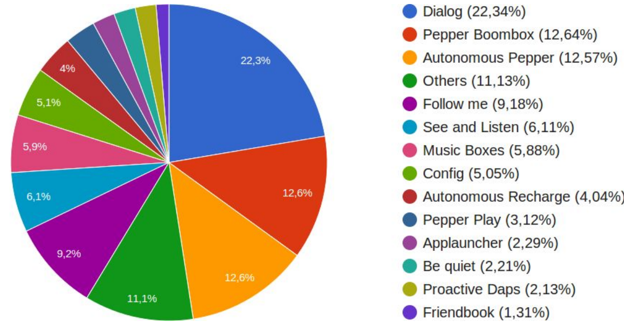

Figure 5: The activity repartition (in duration of activity) cumulated for all users in our test group, for the total duration of the study.

4.2.1 Game applications

First, let's consider the games applications : application made for entertaining, through interactive games with the robot. In terms of duration, they appear often high in the ranking (see Figure 5), but it is most likely due to the fact that their duration is significantly longer than average. In fact, in terms of number of launches, all families have launched a few games at least once, but rarely more than once. This is not really surprising, considering that their content does not evolve, therefore, they do not present much interest the second time. But the fact that all families have used most of the games available at least once, shows that there is still a certain interest in them. This information is particularly interesting for the recommendation system: if the robot keeps proposing an activity users don't want to re-use, it can quickly become annoying.

4.2.2 Moving activities

If we focus on activities that make the robot move around the house, we see that they are not used a lot. Two robots have never launched an activity making them move around: their users probably don't have enough room to use this kind of activities. For the remaining 8 robots, all have used *Autonomous recharge*, but only 3 present a relatively high use of various moving activities. It seems that when users want to, and have enough room to allow the robot to move around, they tend to launch more easily this kind of behaviors. But when they don't, they still have recourse to utilitarian behaviors, leaving aside the ones made for entertainment. This result show the importance to focus on improving these behaviors that are already in frequent use.

4.2.3 Autonomous Pepper

Autonomous Pepper is an optional set of solitary behaviors that allow the robot to have a life of its own. As it is possible for users to set the frequency of those behaviors (or deactivate them) it gives a good indication of what the users want: a robot that keeps silent and still when it is not in interaction, or a robot that acts on its own ?
It seems from the data, that the majority of users in our test group have activated *Autonomous Pepper* with a high frequency of autonomous behaviors: for 9 users out of 10, it is one of the 3 most launched activities (for 3 of them, it is even the most launched). The last robot has few launches of *Autonomous Pepper* early in its history, and none after a certain point: users deactivated it.
We can deduce from this that *users in their majority wanted a robot that acts, moves and speaks on it's own.* But it is important to offer users the possibility to set the frequency of these behaviors, whether manually or automatically.

Indeed, this makes *Autonomous Pepper* a perfect candidate for autonomous evolution: using the history data on previous launches, the robot can automatically set the frequency of solitary behaviors, based on how the users reacted, and the context of previous launches. For example, let's imagine users that set their robot as very expressive and extroverted when they have guests. The more users there are around the robot, the higher the *Agitation* of the history data will be. An improved habit detection could detect that *Autonomous Pepper*'s behaviors are much more likely when agitation is high. The frequency could be automatically increased at that time and turned down when agitation drops again.

4.2.4 Pepper play

The application *Pepper play*, allowing users to create their own behaviors for the robot was expected to be one of the most used. According to data collected in Group 1 before our study, this was the case. During our study however, we observed lower use of the application. One possible explanation is that users habituated to this activity and /or that at the beginning of our study we added *Autonomous Pepper*.
However, 1 family continued to use it enough to be the second most used application. It should be noted that this particular family is the one that deactivated *Autonomous Pepper*. As *Pepper play* can be used to create behaviors and bring life to the robot, we can deduce that this particular user seems to have replaced the behaviors of *Autonomous Pepper* with its own creations. That confirms the deduction that users wanted an proactive lively robot rather than a quiet one.

5. CONCLUSION

In this study, we addressed the problem of long term interaction with domestic robots, and more precisely, the timing of proactive behaviors. We devised an evolutive system for the robot Pepper. This system consisted of a recording of all activity of the robot with various information, such as the time, the agitation around the robot, or the users surrounding it and their mood at that time. Then based on this data, a system could detect user's habits and trigger comments that the robot would make when the current situation does not match with the detect habits. The first part of this system, activity recording, was deployed and tested it in real condition, on a 10-homes test group for 8 weeks.

By studying the activity recordings of the different robots, we found the following observations:
1) Thanks to the comparison of 2 users groups that did not received their robot at the same time, we were able to observe an influence of the novelty effect, on the total duration of use of the robot: users that received their robot recently tended to turn it on more often than users that had it for a long time. This suggests the pertinence of the proactive system that we propose in this paper.
2) When observing habit of use, we were not able to identify a common pattern: users tended to have their own habits that were not necessary similar to the users with the same family composition. Moreover, some users show clearly distinctive habits diverging from the expected use, such as turning their robot on only during the morning before going to work, or late in the night. This should be taken into account when designing new content, and suggest that a system customized to each family is necessary.
3) With their use of solitaries activities, all families showed a preference toward an active and lively robot over a robot that stayed silent and still when not in interaction with a human. It seems to be the direction to take with Pepper's behavior.
4) Finally, we noted that the current games and entertainment applications tended to not be re-used, whereas utilitarian applications (such as the autonomous recharge) are. This information can be exploited in a recommendation system, so that the robot does not propose unwanted activities.

As for the future work, a lot can still be done.
First of all, it would be interesting to collect user's feedback on the robot, and particularly on the behaviors based on the habit detection that have been recently deployed. Has their robot made comments about their habits? What did they think of it?
We also plan to improve the habit detection system, by making it able to detect habits, not only based on the time, but also on all the other information available in the recording, especially the ambient agitation, and the mood of surrounding users.
Finally, we plan to use that improved habit detection to create more complex evolution, such as automatically change proactivity settings according to the current situation.

6. ACKNOWLEDGMENTS

The authors would like to thank Samuel Lascombe, Mathias Le Borgne and Taylor Veltrop for their support of this work.

7. REFERENCES

[1] Iolanda Leite, Carlos Martinho, and Ana Paiva. Social robots for long-term interaction: a survey. *International Journal of Social Robotics*, 5(2):291–308, 2013.

[2] Maartje MA de Graaf, Somaya Ben Allouch, and Jan AGM van Dijk. Long-term acceptance of social robots in domestic environments: In-sights from a user's perspective. 2016.

[3] Rachel Gockley, Allison Bruce, Jodi Forlizzi, Marek Michalowski, Anne Mundell, Stephanie Rosenthal, Brennan Sellner, Reid Simmons, Kevin Snipes, Alan C Schultz, et al. Designing robots for long-term social interaction. In *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1338–1343. IEEE, 2005.

[4] Julia Fink, Valérie Bauwens, Frédéric Kaplan, and Pierre Dillenbourg. Living with a vacuum cleaning robot. *International Journal of Social Robotics*, 5(3):389–408, 2013.

[5] Séverin Lemaignan, Julia Fink, Pierre Dillenbourg, and Claire Braboszcz. The cognitive correlates of anthropomorphism. In *2014 Human-Robot Interaction Conference, Workshop" HRI: a bridge between Robotics and Neuroscience"*, number EPFL-CONF-196441, 2014.

[6] Crystal Chao and Andrea L Thomaz. Controlling social dynamics with a parametrized model of floor regulation. 2013.

[7] Henrik Hautop Lund. Adaptive robotics in the entertainment industry. In *Computational Intelligence in Robotics and Automation, 2003. Proceedings. 2003 IEEE International Symposium on*, volume 2, pages 595–602. IEEE, 2003.

[8] Séverin Lemaignan, Julia Fink, Francesco Mondada, and Pierre Dillenbourg. You're doing it wrong! studying unexpected behaviors in child-robot interaction. In *International Conference on Social Robotics*, pages 390–400. Springer, 2015.

[9] Satoru Satake, Takayuki Kanda, Dylan F Glas, Michita Imai, Hiroshi Ishiguro, and Norihiro Hagita. How to approach humans?-strategies for social robots to initiate interaction. In *Human-Robot Interaction (HRI), 2009 4th ACM/IEEE International Conference on*, pages 109–116. IEEE, 2009.

[10] Michail Maniadakis and Panos Trahanias. Time in consciousness, memory and human-robot interaction. In *International Conference on Simulation of Adaptive Behavior*, pages 11–20. Springer, 2014.

[11] Rayoung Yang and Mark W Newman. Learning from a learning thermostat: lessons for intelligent systems for the home. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, pages 93–102. ACM, 2013.