2019 16th International Conference on Ubiquitous Robots (UR)
Jeju, Korea, June 24-27, 2019

# Lessons Learned from the Development and Deployment of a Hotel Concierge Robot to be Operated in a Real World Environment

Min Sung Ahn, Daniel Sun, Hosik Chae, Itsui Yamayoshi, and Dennis Hong[1]

*Abstract*— **This paper documents the many valuable lessons learned during the development and deployment process of a hotel concierge robot to be operated in a real world environment, which are otherwise difficult to gain in a lab environment or from typical academic research. Real world requirements were explicitly stated by the client before the start of the project, and a significant time constraint of delivering the robot in 1 month was given. Despite the constrained timeline, a functional platform that met all the requirements was created and successfully deployed, where it is run for 24 hours a day. This was possible by relying on rapid prototyping, modifying off the shelf components, and creating a custom codebase based on a scripted language with open-source libraries. As the robot was to be operated in a real world environment, we have encountered numerous problems that are not often encountered in academic research. These include problems such as the robot being repeatedly "exposed" to human interaction, how clothes gradually prevent motions over time, keeping the robot up-to-date after the initial delivery, training of non-roboticists such that they can conduct basic maintenance, and issues that arise as runtime lengthens.**

## I. INTRODUCTION

Developing and operating robots in a controlled environment such as a research lab is difficult, but less so than outside the lab, as assumptions on certain variables can be made and eliminated. When the environment cannot be controlled, such as in an actual hotel, the difficulty becomes unpredictable as not only are most variables uncontrollable, but some cannot even be foreseen.

We see robots outside of labs in places like theme parks, where animatronics is playing a pivotal role in enhancing the visitors' experiences [1]. These robots need to consider the toll that extended runtimes can have on the hardware, but have the benefit that those that maintain the robot are often on-site.

Deploying robots in the real world environment where human-robot interaction is bidirectional presents other difficulties. The diverse end-user population exposes the robot to unpredictable behaviors from humans. Developing and delivering a platform for such scenarios is a non-trivial task for researchers who primarily work in a controlled lab environment with only a limited number of participants to receive feedback from.

For example, mobile kiosk robots have started to be implemented in highly populated areas such as airports, where the robot would have to conduct a daunting task of maneuvering among thousands of people [2]. This is a difficult task not just from a navigation perspective, but also because constant human interaction can lead to a need for more regular maintenance, and close human-robot interaction still has safety issues to consider. These are difficult to test for extended periods, which is why the initial deployment is realistically the first test for researchers. In extreme situations, these tests may suggest reconsidering the initial approach altogether. Also, being mobile robots, the durability is usually better than a robot arm that is moving constantly, as there are less degrees of freedom that can be exposed to wear and tear, which impacts the overall system's performance.

Fig. 1. LARA (Luskin Automated Robot Assistant)

There are some other instances of robots being deployed to interact with humans [3] [4] [5], but the experience and lessons learned from creating and deploying these robots, and the effort required afterwards to maintain their performance, has not yet been documented in much detail. In the industry, most robots created to interact with humans were created by companies whose development and deployment processes and experiences, especially the non-technical details, are kept as trade secrets and therefore are not easily accessible. Generally, to the authors' best knowledge, the details in regards to the presentation, maintenance, and other non-technical details in regards to delivering a robot outside

*This work was supported by UCLA Housing & Hospitality Services.

[1]All authors are with the Robotics and Mechanisms Laboratory (RoMeLa) at the University of California, Los Angeles (UCLA), Los Angeles, CA 90095, USA. The corresponding author can be reached at aminsung@ucla.edu.

the lab, which are as valuable and can be a challenge to overcome, are not well documented anywhere. This paper's objective is to start filling this missing gap in the literature.

### A. Motivation

With the celebration of the opening of the UCLA Meyer and Renee Luskin Conference Center's hotel, the hotel administration wanted a concierge robot that would do multiple things at the lobby. Essentially, it would do a subset of the tasks that a human concierge would do to begin with. Below are the explicit design requirements provided by the hotel:

1) Functional hardware within a month
2) Respond to at least 35 different pre-selected questions
3) Run for 24 hours a day
4) A camera to have the capability to collect data
5) Stay idle sitting down and stand up when someone approaches
6) Easy on/off mechanism of the setup
7) Easy re-location of the setup
8) Easy interface for future development by the hotel staff

With a tight schedule of just a month, development of the hardware and the software began. The plan was to conduct background research of similar implementations for a week, and then develop to be delivered for the remaining three weeks. At the time of the background research, to the best of the development teams knowledge, while there were many similar technologies that had been in the works within a lab, there wasn't one outside the lab doing similar things.

To deliver the product in a limited time span starting from scratch, it was imperative to rely on parts of existing technologies that would promptly get the job done. Also, while satisfying all the requirements, the goal was to initially develop and deliver a platform that can serve as a foundation for further future updates and research by the hotel staff. Many unforeseeable difficulties outside of the technology itself arose during the development stage, especially as the robot was to be stationed without an engineer on site yet among a constant traffic of people. After the delivery, even more occurred.

This paper's primary purpose is to document thoughts, experiences, and lessons learned, which to the authors' knowledge is sparsely available in the literature, but just as valuable as the technical approaches when delivering robots in a limited time frame that also have to perform outside of a controlled environment. The paper is organized where in Section II and Section III, the design and technical approach of LARA (Luskin Automated Robot Assistant), specifically its hardware, software, and human-robot interaction interfaces, are very briefly outlined for future references. Section IV then explains in detail the primary contributions of the paper, which are the experiences and lessons learned with rapid development, delivery, and after-service of a robot that must perform 24 hours a day inside a real hotel. Section V lays out the future direction of the delivered hardware, and Section VI concludes the paper.

## II. DESIGN OVERVIEW

One of the biggest hurdles to this project was time constraint. Therefore, the overall design of the hardware and the software had to heavily rely on available, easy to modify hardware and open-source software. The core hardware was based on DARwIn-OP (Dynamic Anthropomorphic Robot with IntelligenceOpen Platform), a platform our lab RoMeLa (Robotics and Mechanisms Laboratory) originally created with the support of the National Science Foundation (NSF) [6]. Hence, we had the expertise to easily build on this hardware, while the software was built using Python, Node.js, and other easier to prototype languages (as opposed to C++) for reasons explained later on.

### A. Lower-Body

The lower body was fabricated from 80/20 frames and 3D printed ABS. In this application, 3D printed ABS worked well as a way to make custom parts, rather than machining metal or wood, which takes relatively more effort. The legs were wrapped in foam which was then shaped to create a fixed standing motion. The reason behind the modification from the original specification of sitting down when idle, and then standing up when someone approaches the robot is explained in Section IV.

### B. Head Design

The head was designed to portray a robot personality that is intelligent, helpful and charming. The head was made from 3D printed ABS that was sanded and painted to have a finished and professional look. A Raspberry Pi 3 and TFT LCD was mounted in the head for transmitting the camera stream and displaying facial expressions. A large LED was mounted at the top of the head to serve as a status indicator and accentuate whenever LARA speaks.

### C. Visual Modification

The original DARwIn-OP robot does not have identifiable hands or fingers. However, since LARA would often be pointing or gesturing, small 3D printed hands were designed to make the human-robot interaction more natural. To

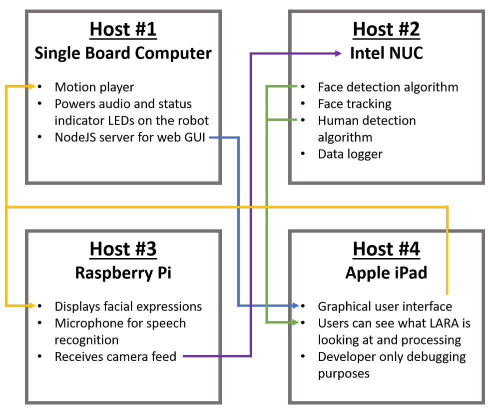


Fig. 2. Block diagram showing how the host computers are interconnected and each of their main roles

smoothen the noticeable layers from the 3D printing process, the hands surface was melted with acetone vapor. Afterwards, the hands underwent the same sanding and painting process as the head pieces.

LARA's overall visual was further modified to better match a hotel setting by using clothing pieces from Build-A-Bear. However, the clothing could not be directly used for a robot doing dynamic motions, and customizations were required to overcome these issues that were unforeseeable at the beginning, which are explained in Section IV.

### D. Software Design

Four separate hosts interconnected through sockets are used to run LARA. To begin with, a motion computer (Host #1) commands the designed joint trajectories in sync with the accompanying voice, while also running a server which runs a web GUI that other hosts can connect to and command the robot. A vision computer (Host #2) runs the necessary vision algorithm for human detection and face detection. A Raspberry Pi 3 (Host #3) powers LARA's facial expressions at the same queue as the motions, while it also retrieves camera feed and passes it on to the vision computer. Lastly, an Apple iPad (Host #4) acts as the interface device to display the web GUI, trigger the motions and the expressions, diagnose LARA, and also visualize LARA's vision streams.

## III. HUMAN-ROBOT INTERACTION

As a concierge robot, LARA interacts with hotel guests, and therefore, it is imperative for the interaction to be welcoming and hassle-free. As a first step, we outlined the potential interaction methods between robots and humans. The two most common methods that seemed viable were interaction over speech and a user interface with buttons that people could physically press and the robot responding accordingly.

To achieve all the design requirements within schedule, a digital user interface on a kiosk was initially decided on. For rapid development, the interface was a web-based GUI built on Javascript and Node.js. As seen in Fig. 3, the GUI has a tabular interface with each page having big beveled buttons, to encourage users to physically press the buttons and navigate through the different tabs—a feature prominent in UI these days. The presentation of the UI was done specifically over an iPad, a well-known device which we believed guests would naturally try to interact with using touch.

At the press of a button on the iPad, LARA was to give a pre-scripted answer, which included a custom motion sequence alongside a synchronized pre-recorded human voice and facial expressions. Because LARA is not simply a virtual chatbot, but is a humanoid that also physically interacts with humans through its motions, rather than using a combination of text-to-speech and random motions, we thought users would be more inclined to extend their interaction with LARA if LARA's motions matched well with her responses. LARA's responses were scripts that the hotel provided beforehand. To maximize the cohesion between her response and her motions, we recorded a human voice for each response, such that both the motion sequences and the vocal response would complement each other and enhance the interaction experience. We chose this direction rather than experimenting with natural language processing algorithms, due to time constraints.

To further enrich the experience, custom facial expressions for each motion was developed, since an important part of human communication happens through expressions. Before designing the facial expressions, a welcoming facial theme had to first be created. As researchers with no experience in character design, an extensive search for previous robot facial designs was done. Fig. 5 shows some of the sample faces that were identified. Ultimately, a face inspired by Disney's Eve was designed. After designing a "base" face, depending on the response, custom animation sequences that were synchronized with the intonation and the motion was created, by varying the face's color and shape. Snapshots of some of the expression primitives are shown in Fig. 4.

Furthermore, in the case with the motions complementing the speech, the head was given an option such that it could move according to the pre-planned motions, but it could also move based on face detection. Because eye to eye contact is an important part of interpersonal communication, the camera mounted on LARA's forehead was used to detect faces. Based on facial detection, a face tracking motion could be executed.

## IV. RESULTS & EXPERIENCES

LARA was built and delivered to the Luskin Conference Center within schedule, but few compromises and adjustments had to be made from the requirements list in Subsection I-A. This was primarily because of the aggressive time constraint and the unforseeable amount of non-technical difficulties once LARA was deployed on-site.

### A. Adjustments due to Population Count

To begin with, unlike mobile robots, legged robots have the advantage to sit down and stand up to create exaggerated motions and grab attention to let guests know that it is available to help. While this was not an issue in the first hour or two of operation, it quickly became a job that was demanding on the robot's actuators during beta testing

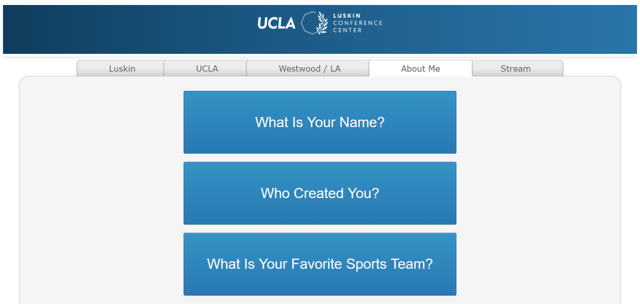


Fig. 3. Graphical user interface for user's to interact with LARA. The tab shown is where users could learn more about LARA.

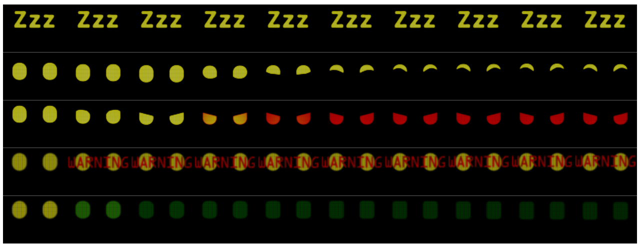


Fig. 4. LARA's facial expressions were custom designed for each motion. The above are some snapshots and primitives that the sequence of expressions are made. Row 1: LARA when idle for an extended period of time. Row 2: Smiling. Row 3: Angry. Row 4: Error mode that signals that something has gone wrong somewhere in the system. Row 5: Special seasonal Halloween update where LARA was setup as a zombie.

Fig. 5. Some of the sample faces that were identified before the development of LARA's face. In order from left to right: Eve, Cozmo, Baxter, Yumi, Buddy

at the hotel. The internal testing of the sit/stand sequence within the lab did not serve as a good predictive test, as the number of people entering and leaving a hotel significantly outnumbered those entering and leaving the lab. Therefore, an alternative, quick solution was sought for, resulting in choosing to go with a constant standing posture. This meant removing the lower body actuators entirely and building legs out of 80/20 frames and 3D printed ABS parts. By deleting the standing up motion, other means to catch the attention of guests was needed.

While this structural modification was quick and practical, which was what was viable in a limited time, it unfortunately made LARA feel more like a kiosk, rather than truly an active robotic personal assistant. Thus, working with the hotel staff as to what type of a standing posture would make LARA look more natural and less like a passive kiosk, foam was used around 80/20 frames such that when LARA's pants were on, it would sport a professional look like those working at a hotel.

As a result of the above modifications, features to grab attention in a different manner had to be available. Simple adjustments included installing louder speakers to trigger people's sense of hearing rather than sight, while catchy sounds and remarks were played at human detection. Because of the heavy traffic in hotels, these remarks also had to be unobtrusive to people working at the lobby, because responding at every human detection would cause the robot to be talking nearly all the time.

### *B. Clothing Dynamics*

In regards to the 80/20 based legs, the adoption of pants for such pair of legs was a simple task. However, in regards to the visuals of the upper body of the robot, many unforeseeable issues with clothing occurred all throughout the development. The most notable of the many is that clothes on robots impose constraints on the motions that the robot can do. This is beyond just the obvious kinematic constraints. Unlike humans' skin, robot frames do not conform to outside pressure, so clothing had to be specifically chosen based on whether we would be able to fit them on LARA. Beyond this, what had to also be considered was the problem that while humans naturally and unconsciously make small movements to keep ourselves look presentable, with LARA, regardless of whether the sleeves were short or long, the more it moved, the more both the sleeves and the hem of the shirt slowly curled up. This ultimately made the robot look unprofessional as runtime extended. This issue was solved by strategically modifying the original clothes using elastic bands such that the clothes would be flexible enough for a variety of motions, but still come back to its original place because of the elasticity.

We can confirm that dynamics and constraints of clothing on robots affect the motions that can be created. Aside from the slow curl up, motions generated without clothes made some clothes get pinched and torn over time due to the motions. This is an interesting issue both from a motion generation and a clothing design perspective that deserves further attention if we want to avoid modifying clothes every time we want to put it on a robot. On a side note, we could have chosen to go with the more common approach of a custom, painted plastic cover, which then would have allowed us to circumvent issues coming from clothes. However, such approaches would have made LARA look too "robotic," making it potentially feel unwelcoming and out of place for the Luskin Center because of its rigid and industrial contours.

### *C. Seasonal Updates*

LARA was initially built in a month with clear deliverables. To achieve this in a limited time, a pre-recorded human voice was used to guarantee an engaging human-robot interaction, rather than developing or implementing a text-to-speech algorithm which had the potential to take too much time to develop and also reduce the engagement level. The facial expression and the motions were, however, delicately hand designed to complement the overall experience of talking with a concierge robot because no suitable technology could be quickly found and implemented.

With this tight integration, LARA received a very positive reaction, despite its limited capabilities compared to a human concierge. However, because of the success in creating this engaging experience, after the delivery of the robot, LARA became a popular figure that not only the hotel, but also the university wanted to use for a lot of their internal events and marketing efforts, thus requiring seasonal updates [7]. Unfortunately, unlike a lab setting where time is often heavily invested until satisfaction, the hotel moved in a strict timeline, which was difficult to meet because every one of LARA's responses take time to craft to maximize the engagement. It was clear that this was impractical after a few updates. Quickly done jobs were noticeably less engaging and unnatural, as the client would sometimes want further modifications. Training a group of hotel staff such that there would be a team of robot caretakers in the hotel was an option, but that had its own difficulties, which is further described in Sub-section IV-D.

### D. *Staff Training*

An integral part of delivering LARA was to allow the hotel staff to be able to conduct basic maintenance of it. Initially, due to time constraints, rather than doing an on-site training session with the robot and the staff, a manual was written with instructions on how to debug the robot in the case of abnormal behavior. This is something that is seldom done in a lab environment, as research publications or internal wikis are instead used. Even then, these are read by an audience that are in a relevant field from a similar background. With LARA's manual, it had to be for a group of people whose prior knowledge was unknown.

Writing down a manual was a non-trivial task mostly due to the nature of robotics. Manuals for other high-tech devices such as cellphones or drones can be easier because of the maturity and abundance of the technology in our everyday lives. Additionally, because there are many similar products, if the end-user was to face difficulties, they may be able to also find a solution by inferring from other similar documentations online. However, because of the scarcity of consumer robots in our everyday lives and since it is a highly interdisciplinary field, capturing the important details of how to debug the robot and avoiding writing a lengthy document was difficult. The hotel staff also preferred a short and concise documentation as lengthy ones can deter people from reading them in the first place. Scripted languages were also chosen for this purpose, as they are relatively easier to comprehend, thus empowering the hotel staff with the ability to conduct updates and basic maintenance by themselves.

However, on-site staff training still had to occur because of an incident where the robot had a software bug (explained in Subsection IV-E) after an extended period of operation, and to fix this issue without the researchers, the hotel staff had disassembled key components of the robot. While disassembly was explained in the manual, because a single independent document could capture only a subset of the technology behind it, the staff had made mistakes such as resetting the router configurations, which is critical to the entire system being functional. However even with the training, new issues continued to rise from time to time, and it was impossible to train for every potential scenario.

This issue seems to imply a need for standardization in robotics from hardware (e.g. on/off buttons and their locations) to software (e.g. GUIs), just like how the smartphone industry has a similar look (both hardware and software) in all devices regardless of the model. Robotics is a multidisciplinary field and it will be difficult for the general public to understand the technological details. Thus, it is the responsibility of roboticists to not only ensure that extensive testing of the product is done such that issues do not rise in the first place, but when they do rise, because of a similar look throughout the industry, users can infer from their past experiences with other robots to conduct initial diagnosis. We believe that such standardization will lead to easier adoption of service robots as the end-user will feel comfortable using them.

### E. *Integration of Available Open-Source Technologies for Prompt Delivery*

During this process, we also did find that creating a complete, purpose-built working humanoid robot is still a monumental task. To some extent, this is being improved upon by projects like ROS, however in practice there are typically between 5-10 software modules that need to all interact with each other. Unfortunately, because of the nature of open-source technology, many packages that were needed were not rigorously tested or updated. Issues were bound to rise when building a humanoid robot based on these packages due to having to deliver in a month, but also needing the robot to run for 24 hours a day.

When the robot was running for days, we found memory leaks in a couple modules that caused the robot to eventually crash. These issues would be difficult to identify had such code never been run in extended lengths, which is usually the case with robotics as demonstrations are either filmed or only a few minutes long. Hence, in our case, patches had to be delivered after the technology had been initially installed. This may have been preventable, had the selected modules been rigorously tested independently, early in the development stage, prior to system integration.

## V. FUTURE WORK

Currently, reliability issues are continuously analyzed and resolved. How can we ensure a near 100% uptime since a hotel does not have closing hours? How can we make the robot more robust to misuse from people that are less tech-savvy? What safety measures do we need to implement to avoid potential casualties? As researchers that primarily conduct experiments in a controlled environment surrounded by other researchers, these questions (and more) that need to be solved are completely different to what we face inside a lab, which requires us to rethink how things need to be implemented and controlled when thinking about deploying a robot in an uncontrollable setting.

From a longer outlook standpoint, with ample time, there are a lot of areas and functionalities that LARA can be improved upon as a robot working outside the lab. To begin with, in a hotel, we believe that a more personal relationship between a guest and a concierge robot can be established if the robot could label unique faces and store them in a secure database and greet them as they enter and leave the lobby. While this brings up privacy issues, we also believe that this collection of data can be used to enhance a guest's experience at the hotel. Additionally, alongside integrating advanced text-to-speech technology that is quickly making progress in the smart speakers industry, we believe that for an intimate interaction between a human and a robot, a new paradigm of text-to-motion (TTM) needs to be investigated.

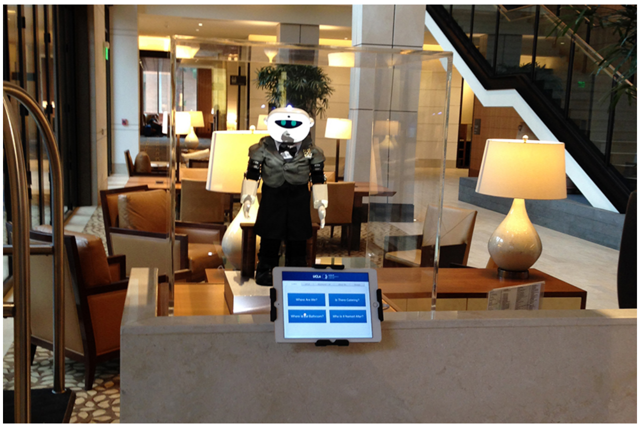

Fig. 6. LARA working at the lobby of the Luskin Conference Center.

Because it is impossible to have a pre-programmed motion plan for every speech output from a service robot running a nautral language processing algorithm, primitive motions and a dataset of poses and speeches (with their intonations) could be used such that motions can be automatically generated based on the output of these algorithms.

Furthermore, we realize that while LARA is a service robot created specifically for a hotel, as shown in Fig. 6, similar robots will make their way into various parts of the society. Thus, before their wide spread adoption, LARA can be used as a general platform for gathering data and conducting human-robot interaction research. To realize this, we have an API on top of our current setup in our lab's repository[1] such that experts from other fields such as machine learning and computer vision can recreate and utilize a pre-made platform that can constantly interact with people in the real world.

## VI. CONCLUSION

In a very limited amount of time, a platform that is acting as a hotel concierge at an actual hotel was developed from scratch. Rapid development was possible by relying on readily available hardware and modifying off-the-shelf components. An amicable presentation of the robot by embedding a character through a cohesion of speech, motion, and facial expression was done to encourage guests to interact with the robot.

However, because this was a robot deployed outside of the lab in a short amount of time, valuable experiences were gained. Specifically with LARA, it was difficult to let people know that the robot is available to help, rather than passively waiting until someone approaches for help. This was because since LARA is the first concierge humanoid, no one expects a robot to be available at the lobby. Visual representation of the robot in a professional setting was also a novel challenge, as the clothes' dynamics heavily affected the motions that the robot could do, which was an evolving confined space compared to what kinematic constraints made. Continuous updates because of a large positive reaction from the community also signaled a clear need for natural language processing which also includes inferred motion outputs in the case that the robot is not simply a mobile kiosk. Documentation for non-roboticists presented a novel challenge, because the current lack of a similar look in robotics in both hardware and software meant either a need for a lengthy documentation or a frequent on-site visit by a developer.

Delivering a robot outside a controlled environment in a short time span where the robot had to perform 24 hours a day illuminated a lot of different areas that service robots will have to tackle in the future. Unfortunately, such information and insights are sparse in literature, despite their importance when trying to make a robot an actual product. This paper attempts to begin to fill this gap by providing insights to building and delivering service robots in the real world where human-robot interaction can repeatedly occur at an intimate level.


## ACKNOWLEDGMENT

The authors would like to acknowledge the assistance of various individuals who contributed their time, breadth of knowledge, and expertise to LARA's development: Hari Krishnan, Alexie Pogue, Colin Togashi, Matthew Williams, Xiaoguang Zhang, Chun Yin Chan, and Seulki Koo.


## REFERENCES


[1] Wikia. Audio-animatronics. [Online]. Available: http://disneyparks.wikia.com/wiki/Audio-Animatronics

[2] L. P. Release. (2017) Lg airport robots take over koreas largest airport. [Online]. Available: https://www.lg.com/sg/press-release/lg-airport-robots-take-over-koreas-largest-airport

[3] E. Ackerman, "Toyota's new human support robot gives disabled humans a hand (and an arm)," Sep 2012. [Online]. Available: https://spectrum.ieee.org/automaton/robotics/home-robots/toyotas-new-human-support-robot-gives-disabled-humans-a-hand-and-an-arm

[4] E. Guizzo, "How aldebaran robotics built its friendly humanoid robot, pepper," Dec 2014. [Online]. Available: https://spectrum.ieee.org/robotics/home-robots/how-aldebaran-robotics-built-its-friendly-humanoid-robot-pepper

[5] N. Statt, "Hilton and ibm built a watson-powered concierge robot," Mar 2016. [Online]. Available: https://www.theverge.com/2016/3/9/11180418/hilton-ibm-connie-robot-watson-hotel-concierge

[6] I. Ha, Y. Tamura, H. Asama, J. Han, and D. W. Hong, "Development of open humanoid platform darwin-op," in *SICE Annual Conference (SICE), 2011 Proceedings of.* IEEE, 2011, pp. 2178–2181.

[7] M. P. Xochicale. (2017) Lara - luskin automated robotic assistant". [Online]. Available: https://www.youtube.com/watch?v=f1jZo6TgdWs


[1]https://github.com/RoMeLaUCLA/LARA