PROJECT: Symbiosis

The development of an interactive swarm installation for Lumus Instruments.

Graduation Project Report
Creative Technology - University of Twente
Wouter Achterberg - s2151863

Supervisor
Edwin Dertien

Critical Observer
Angelika Mader

January, 2021

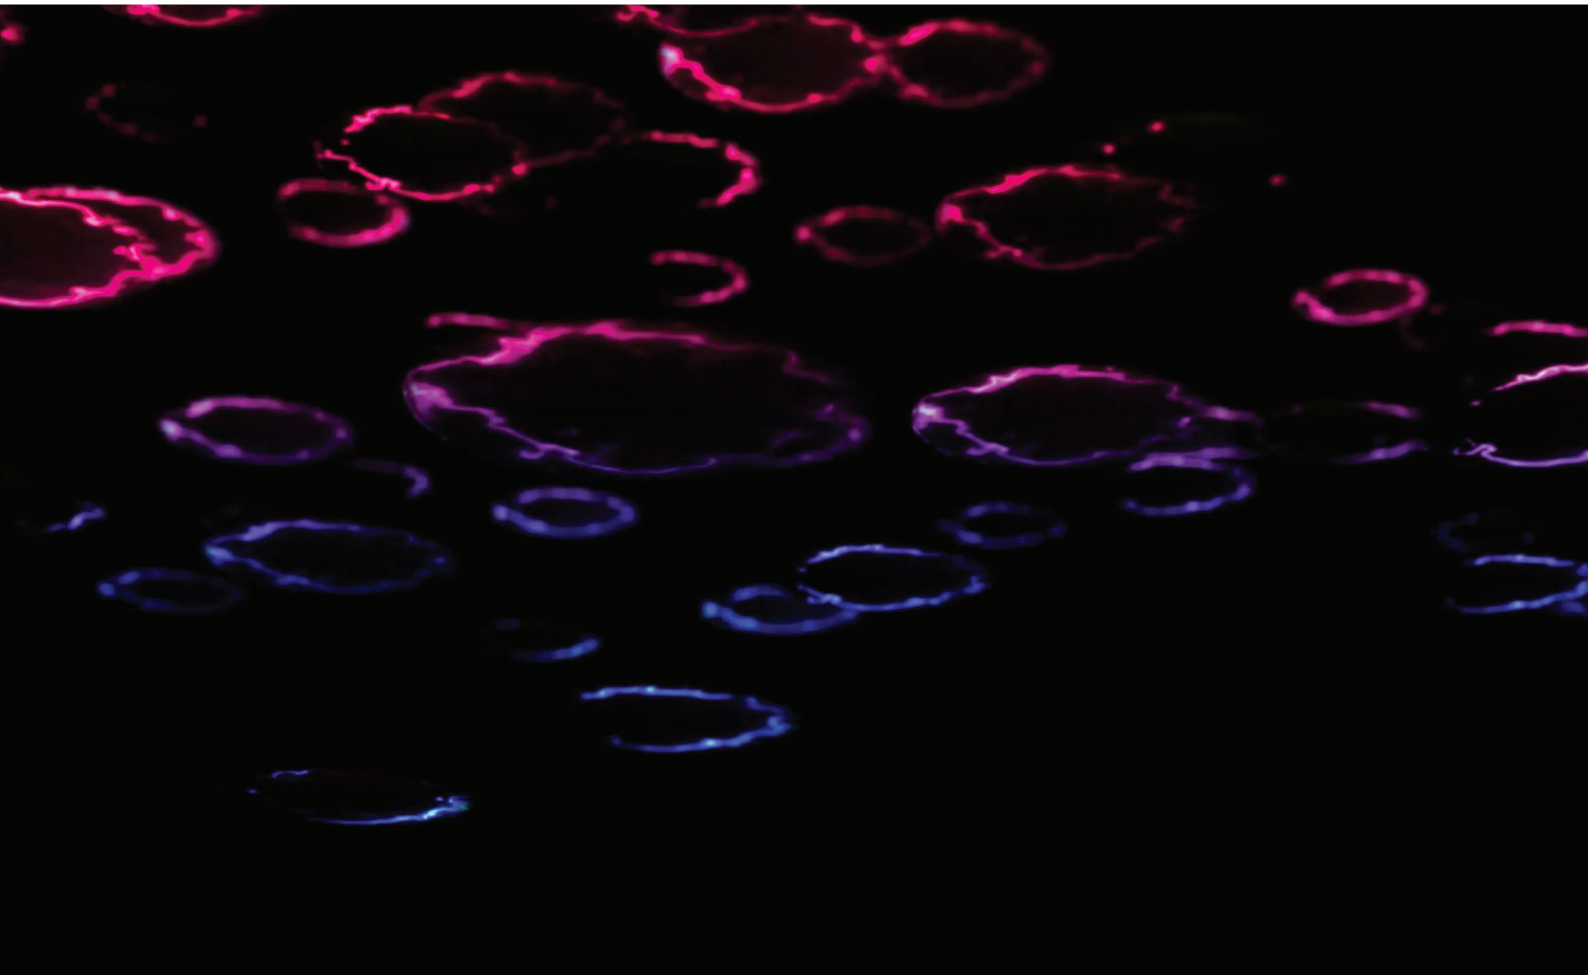

ABSTRACT

Swarm installation and dynamics allow for excellent opportunities in the future as in the military, healthcare, and entertainment. The implementation of swarm robotics in the entertainment world have essential aspects to consider; the human-swarm interaction and the inter-agent interaction, and the implementation of centralized or decentralized intelligence. This graduation project focuses on implementing decentralized inter-agent interaction while keeping the installation interesting entertainment-wise. Therefore, multiple iterations of the prototype have been developed using machine vision and deep learning. By thoroughly testing the prototype technically, recommendations for the use of neural networks, inter-agent algorithms, and the agents' physical appearance have been proposed to extrapolate towards a final installation.

swarm dynamics, decentralized intelligence, inter-agent interaction, human-swarm interaction, engagement and enticement

ACKNOWLEDGEMENTS

Thanks to **Timo**, **Timothy**, and **Julius** from Lumus Instruments for allowing me to work on one of their unique concepts and assisting me throughout the development process.

Thanks to **Edwin Dertien** for thoroughly assisting me with valuable feedback and thoughts throughout the whole process.

Thanks to **Angelika Mader** for providing feedback for my report.

CONTENTS

List of Figures

1 INTRODUCTION

1.1 Context

Lumus Instruments is a creative lighting studio that originated from TU Eindhoven industrial design engineers and is currently based in Amsterdam. Their vision is to design and develop products that inspire and enable sustainable art. Their designs are showcased on several venues, such as light-festivals, art sceneries, and audiovisual live performances.

Lumus Instruments has been brainstorming about a new concept to emphasize the harmony between nature, technology, and humankind. This concept is PROJECT: Symbiosis. Symbiosis aims to research the disconnection of humans from their environment and nature. Symbiosis will represent the environment they reside in with many individual creatures that will together form a field of creature: a swarm. The swarm's creatures will harmoniously live within the group and respond to internal and external inputs. Symbiosis will amplify the context and slowly grip on new emerging behavior stimulated by the context. This amplification of the context must be achieved by sensor data only, requiring the agents to be contextually aware. Also, Symbiosis is meant to be a calm, slowly emerging installation, but it must also be interesting enough to visit. Therefore, Symbiosis must partially reside in the background and be engaging and enticing enough to be visited.

1.2 Problem Statement

The challenge for Symbiosis will be to design a swarm installation with inter-agent communication without using a data network across the swarm's agents. Each agent will pick up external factors with its sensors (camera and microphone). These external factors can vary from dB levels to average colors in the frame or 'states' from neighboring agents. External factors can essentially be any sensor data input Lumus Instruments would like to add. These states will be discussed thoroughly in the latter of the report. These external factors combined make up the contextual awareness of the agent. As the agents are only interconnected with a power grid, they cannot exchange data and rely on their contextual awareness to construct audiovisual output behavior. The lack of data interconnection forces decentralized intelligence.

The installation will be in the background of the scenery/context and must be interesting enough to visit. Staying interesting in an installation is a difficult obstacle, defined by the novelty effect, decreasing visitors' interest in an installation quickly. Finding the balance between the calm, background technology, and the installation's engagement and enticement are crucial for designing Symbiosis.

1.3 Goal and Research Questions

The swarm installation goal is to amplify their environment and behave as a swarm with inter-agent communication. The creatures should use environmental factors to construct their behavior. The interaction with users should lean towards the calm, background, and ambient technology. The interaction with users and inter-agent interaction should stay interesting, avoiding the novelty effect and incorporating enticing aspects.

However, this graduation project's scope and goal will not be the entirety of the swarm installation. This report will primarily cover the inter-agent interaction and how agents communicate while keeping the swarm dynamics in mind. This will be explored by designing two swarm agents and modeling their behavior and expected behavior within a swarm context.

The challenges and the goal of the swarm installation together allow for the following research question:

- How to design a swarm installation with decentralized inter-agent interaction?

The main question allows several subquestions to be composed and answered during the graduation project.

- How to design contextual awareness of agents?
- How to incorporate engagement and enticement in an (artistic) installation?
- How to encalm technology?
- How to translate current human-computer interaction concepts into aspects of human-swarm interaction?

The project contains no constraints. There is no time frame in which the art installation needs to be finished. The art installation has got no rigid requirements and leaves plenty of space for personal and artistic interpretation.

1.4 Document Structure

The graduation report starts with the exploration phase. This part will consist of three main parts. The first part will consist of the initial brainstorming, stakeholders, and state of the art. The second part consists of background research and relevant research to understand essential and relevant concepts. The third part will consist of the requirements deducted from the exploration phase and a discussion with Lumus Instruments.

After the exploration phase, the ideation phase will be explored. The ideation phase starts with a detailed experience design. A more detailed overview of the creature's behavior will be discussed from the experience design, followed by evaluations of hardware and technical setups.

In the implementation phase, the discussed concepts from the ideation phase will be translated into physical prototypes. There will be several iterations of prototypes constructed, tested, and evaluated. This chapter concludes with a deconstruction of a minimum viable product.

In the testing phase, the minimum viable product will be tested on several aspects. These aspects include tests with response performance, lighting conditions, distance, and noise sources.

After the testing phase, the minimum viable product will be evaluated. The primary evaluations will cover requirements, stakeholders, and extrapolation to the final swarm installation.

In the concluding chapter, the results from the testing chapter and evaluation chapter will be compared to earlier chapters' literature. The main conclusion from the project will be discussed, including limitations and future work.

2 EXPLORATION

The exploration phase covers all the initial research that has been conducted. The exploration phase starts with an initial brainstorm with the client. This brainstorm reveals important aspects of the project. The brainstorm also brings a clearer image of the project stakeholders, which will be discussed right after. With the help of current, similar installations, state-of-the-art will be addressed.

The next part of the exploration phase will be conducted with scientific papers, case studies, and related research projects. Relevant domains, frameworks, and concepts will be discussed. These relevant concepts will be worked out in the deepening relevant research section, giving more detailed insights. Also, an in-depth reflection on the transition from HCI to HSI from the author will be included. The exploration phase will be concluded with requirements deducted from the exploration phase.

2.1 Initial Brainstorm

The start of the exploration phase is a result of a conversation with Lumus Instruments. This conversation aimed to get a general understanding of the desires and opportunities from both sides. During this conversation, four main areas got attention. Within these four subcategories, essential aspects or quick side-thoughts have been written down. These initial thoughts are the starting point for the exploration phase. They give broad guidelines and ideas for implementation opportunities. They have been bundled as a mindmap in appendix A.

One of the main components of the mindmap is sensing, which has been deducted to contextual awareness. How would the swarm agents perceive their context, and how would they distinguish neighbors from humans? The discussion covered which modalities the agents should have, such as hearing, seeing, and feeling. The establishment of the primary sensing components is of great importance.

Another important topic of the mindmap is behavior. The behavior of the swarm will be of great importance, and there are lots of possibilities. How will the agents react to certain situations, and will they do this individually, as a swarm, or somewhere in between as subswarms? Determining the fundament of their behavior is crucial. Which characteristics do the creatures get, and how will these be influenced? Relevant research fields are essential to explore further.

The modularity and materials are equally important but will mostly be handled by Lumus Instruments or a third party.

2.2 Stakeholders

PROJECT: Symbiosis will have several potential uses. The concept and project belong to Lumus Instruments, which operates in the entertainment area. The primary purpose of Symbiosis is to be deployed/displayed on light and art festivals. The stakeholders will therefore be visitors of the festival. Generally speaking, the age of the stakeholder will vary between ten years old and seventy years old. They will likely have an interest in technology and artistic, creative lighting installations. There is a high probability that they have attended more related festivals or shows, making it more challenging to grab and maintain their attention.

Next to the end-users, organizers of these festivals would need to be willing to rent Symbiosis to be displayed. Organizers are hard to pin down since they vary significantly in age, gender, and demographic data. However, the fact that they organize such festivals shows their affinity with the topic. This, as stated above, indicates that they have seen more installations and are curious for new, creative installations that will want the visitors to keep coming.

2.3 State of the Art

State of the art for this project focuses on the installations that are already out there. State of the art will give insights into how likewise structures are being set up and constructed. How are these installations novel, and which features or characteristics can be relevant for Symbiosis? Deconstructing some of the art installations from GOGBOT, Lumus Instruments themselves, and the NXT museum: Shifting Proximities.

The installations discussed in this section have been explicitly selected to cover specific aspects of Symbiosis. Playmodes by Cluster closely resembles the audiovisual show with spontaneous, oscillator driven data. The installation comes to life at night, which will also be the case for Symbiosis. The Wixel Cloud by Blendid is chosen because of the swarm-like approach. It could give insights into the multitude of objects approach, which is similar to the multitude of creatures in the Symbiosis swarm. SVNSCRNS by Strijbos focuses on the grouping of static objects while remaining intriguing for visitors to watch. Since Symbiosis will consist of stationary creatures, too, there could be interesting takeaways. The three final installations, of which one is of Lumus Instruments, were chosen to understand Lumus Instruments' vision better. Their installation perfectly resembles their style. Lumus Instruments recommended visiting the NXT installations to indicate the direction they want to go to.

2.3.1 Cluster - Playmodes

Cluster[1] is an installation designed and created by Playmodes. It is an audiovisual show, which is oscillator-driven. This means that there is no intelligence or spontaneous behavior involved since the data is being generated by an oscillator. The light show will therefore be pre-generated since the only difference between each show is the randomness of the oscillator data.

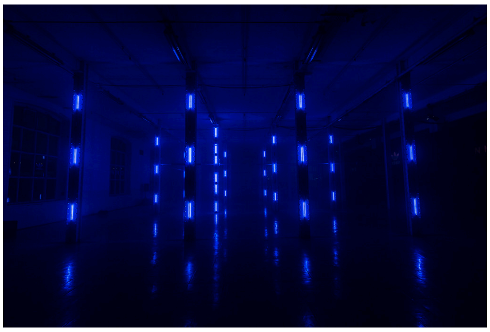

(a) Playmodes Active

(b) Playmodes Inactive

Figure 2.1: Playmodes Different States

2.3.2 Wixel Cloud - BLENDID

The Wixel Cloud [2] consists of 75 Wixels (wireless pixels) that creates a spatial resemblance of the environment they reside in. The project aims to have 3d representations of compositions. With the help of Blender, Python, and the openFrameworks library, the 3d composites are projected onto the individual wixels. There is no intelligence involved in this art installation since the compositions are being displayed on the wixels and the individual wixels are only passively used. Nevertheless, this installation shows organic behavior and allows for smooth motion.

[1] https://www.playmodes.com/home/cluster/
[2] http://blendid.nl/index6258.html

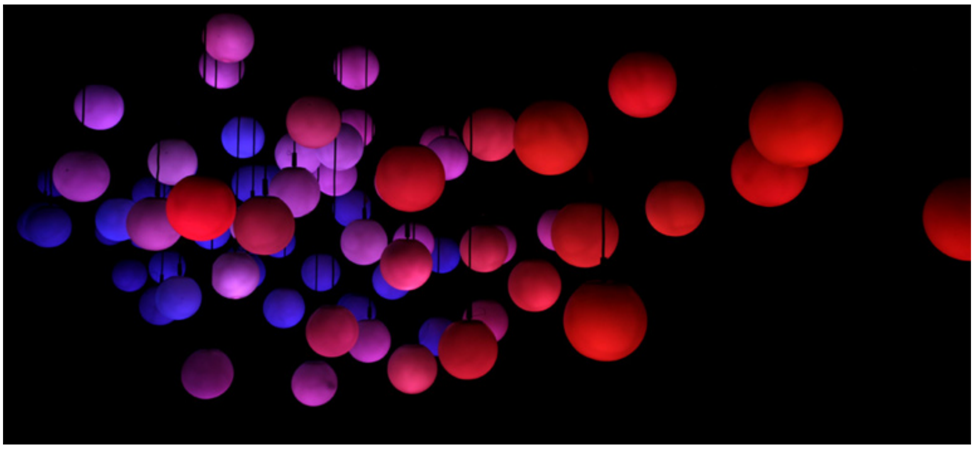

Figure 2.2: WixelCloud - Blendid

2.3.3 SVNSCRNS - Joris Strijbos

SVNSCRNS [3] is an art installation on GOGBOT 2020, which consists of seven rotating screens that each display an image. The idea behind the installations is to let the audience realize the content they are watching. It allows being displayed during live performances or premade compositions. There are also speakers present, which together with the screens can be controlled from a central computer. Light, sound, and motion then join forces to create a multi sensational experience. Custom-built software allows the user to determine the imagery and sound, making it a centralized experience. There is no spontaneous behavior, neither does one of the screens/sounds develop itself.

Figure 2.3: SVNSCRNS - Joris Strijbos

[3] https://www.jorisstrijbos.nl/work/svnscrns

2.3.4 Contratrium - Lumus Instruments

Contrarium [4] is the latest addition to Lumus Instrument's line-up of light installations. Contrarium is an audiovisual live experience and uses symmetry as one of the core aspects. The audio output is split into different bands, which will produce a lighting output. However, this output can only be generated with the lighting controller. This does not necessarily make the installations less interesting to watch, but does not include an intelligent interaction form. Lots of the interaction is pre-programmed or determined on the spot by the light controller.

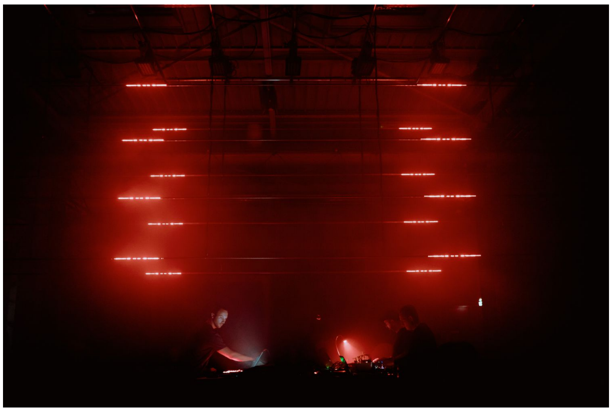

Figure 2.4: Contratrium - Lumus Instruments

2.3.5 NXT Museum - Shifting Proximities

Shifting Proximities [5] is a new media art exhibition at the NXT museum. Upon the recommendation of Lumus Instruments, I visited Shifting Proximities. Shifting Proximities consists of new media installations, focusing on the human experience and interaction in the face of social and technological change. In appendix B, the reflection upon the full exhibition can be read.

Econtinuum [6] by Thijs Biersteker is an organically designed form of Symbiosis between two trees' roots and the installation visitors. The similarities between PROJECT: Symbiosis and Econtinuum may not be apparent initially, but there are some underlying similarities worth mentioning. The interaction between the installation and the users looks like calm, background technology, and uses sensors to achieve this interaction, which is PROJECT: Symbiosis's goal. Fluid gas sensors and temperature sensors change the symbiosis speed between the displayed trees' roots and alter the projection on the wall.

---

[4] https://lumus-instruments.com/project-detail/5f3aa688afbcb716d0692a79

[5] https://nxtmuseum.com/nl/event/shifting-proximities/

[6] https://nxtmuseum.com/nl/artist/econtinuum/

(a) Distortions in Time - Marshmallow, Laserfeast

(b) Econtinuum - Thijs Biersteker

Figure 2.5: NXT: Shifting Proximities - Distortions in Time & Econtinuum

Distortions in Spacetime [7] is an artistic installation on the topic of black holes. The black hole visualizations are chaotic but emerging and give an organic, esthetic atmosphere to the installation. The use of strong sound effects builds on the visual fundamentals and increases the impact of the installation. PROJECT: Symbiosis's goal is also to incorporate esthetic visualization and emerging sound experiences. The execution of Distortions in Spacetime was well done, which creates the opportunity for PROJECT: Symbiosis to incorporate these elements.

2.4 Background

2.4.1 Research Fields

There are several ways to approach the concepts of swarms and agents. The way agents within a swarm interact and communicate is one of the critical aspects of swarms. Next, the communication and interaction between individual agents and users is another key aspect. It is crucial to get a grasp at the current understanding of the various ways swarms are dissected. The relevant research fields give brief insights into multi-agent systems, artificial intelligent forms, and contextual awareness.

The state of the art installations described in the previous section do not utilize artificial intelligence, except for Econtinuum. For Lumus Instruments, the visuals must be engaging, but the technology that drives the installation must be interesting, hence the decentralized intelli-

[7] https://nxtmuseum.com/nl/artist/distortions-in-spacetime/

gence. Lumus Instruments strives to resemble an organism that can see and hear. That is why the primary sensors will be a camera and a microphone, which resemble seeing and hearing, respectively. The agents in the swarm will focus on predicting the state of neighboring agents. As agents will express themselves using dynamic lighting, these lighting states' snapshots can be made and analyzed. These predictions are based on the live feed of the agents' camera, which is the only visual input an agent has. When analyzing visual data and making predictions based on visual data, machine vision, and neural network classification are interesting topics.

Distributed Artificial Intelligence

Distributed artificial intelligence [16] can tackle complex problems and require autonomous learning processing nodes (the agents). The systems are robust and elastic. The power of distributed systems is that they do not require the data to be centralized, but rather have it distributed over the autonomous processing nodes. The main challenges for DAI are the communication and interaction of agents, the coherency of agents, and the synthesis of results among agents.

Multi-Agent Systems

Multi-agent systems [9] are a form of distributed artificial intelligence. A multi-agent system consists of multiple intelligent agents living in an environment. In our case, the agents are active, with simple goals as directing lights and responding to neighbors. Whether the environment is virtual, discrete or continuous has yet to be determined. Agents within a multi-agent system have requirements such as autonomy, local views, and decentralization. Agents are partially independent, have no global view and no agent is controlling. The main concept is that these systems can be self-directing, and self-organizing, while the individual actions of the agents are fairly simple.

Agent-Based Models

Agent-based models [8] aim to solve and describe complex phenomena. It looks at the interactions of individual agents and subgroups and their effects on the total environment/group. There are general aspects of agent-based models; agent granularity, decision-making heuristics, learning rules or adaptive processes, an interaction topology, and an environment. An

[8] https://en.wikipedia.org/wiki/Agent-based_model

agent-based model can result in interesting behavior. Often, agents' location and responsive behavior are encoded in a central algorithm, which takes away the intelligence from the individual agent. When the model is inductive, the behavior encaptures the unexpected behavior, which is what can be regarded as organic. Inductive reasoning allows for organity since the agents use some form of evidence in the form of observations and truths from other agents, but synthesize their own conclusions from this evidence making it an uncertain conclusion.

Agent-based models would be a great opportunity to generate organic patterns since they are widely used to resemble natural phenomena such as vegetation ecology, landscape diversity, and plant-animal interactions to name a few. The difference between agent-based models and multi-agent systems is that ABM looks into the collective behavior of the agents (so there is no need to be intelligent).

Cyber-Physical Systems

In cyber-physical systems [1], hardware and software components are combined and controlled by a computer-based algorithm. CPS is a widely used term for lots of systems. CPS is strongly related to the Internet of Things, but focuses more on the parts of the embedded system and does not necessarily need to be connected to the internet. However, for this GP, IoT would suit the job more than CPS, since the internet is able to provide valuable information for the swarm, such as time and location.

Neural Network Classification

Neural networks can cluster and classify input data and map the input data on output labels [8]. Especially the supervised learning variant can perform well on classifying. Supervised learning uses pre-labeled data to extract features from data to map unseen data to the correct output label. Neural networks use automatic feature selection, whereas traditional machine learning algorithms use human-selected features. Neural networks are composed of several layers, existing of nodes. Data traverses through these node layers and will be mapped to an output layer.

Deep Reinforcement Learning

Deep reinforcement learning [10] is part of machine learning and therefore belongs to artificial intelligence. Deep reinforcement learning allows machines or agents to learn from their actions and the results thereof. An agent can be penalized or rewarded for taking specific steps, forcing them to understand typical behavior related to the task. Machines deployed for deep reinforcement learning have an end goal. Choices bringing the agent closer to this end goal are reinforced, whereas choices resulting in moving away from the end goal are penalized. The downside of DRL is that there needs to be specified a particular purpose or end goal, which cannot always be achieved beforehand. Most applications of DRL therefore find their purpose in games.

Engagement and Enticement

Engagement and enticement [2] for complex systems and art installations are of great importance. Installations often suffer from the novelty effect, which takes away the newish impact and decreases the interest users have in the installation. Engagement and enticement ensure that the installation will be exciting and make people want to interact with the installation. Engagement and enticement are essential for all installations, but more so for artistic installations. The primary purpose users visit these installations is to be entertained, without any practical side-uses, which enlarges the pressure of creating inviting, enticing installations.

Background Technology

Background technology, also referred to as calm technology [4] or ambient technology, is an essential aspect of Symbiosis. Background technology ensures that the installation's technology is subtle and not too intense. Incorporating this calm form of technology also allows the technology to organically live independently without needing too much attention from the (human) user.

Context Awareness

Context awareness [7] refers to how mobile devices are capable of sensing their environment with the help of sensors. Especially regarding the core aspect of Symbiosis, which is to amplify its context, context awareness is essential. The swarm agents should proactively respond to external and internal input and reflect these inputs through light and sound.

Humand-Computer Interaction and Human-Swarm Interaction

Human-Computer Interaction (HCI) focuses on the interface between humans and computers. It focuses on designing user-oriented interfaces and how we could beneficially design computer interface to improve usability [19]. Many models and theories have been developed since the popularization of HCI, which was around 1980. Interesting concepts and ideas for this project revolve around bonding between humans and computers, the sense of credibility towards the computer, and the need for human autonomy and competence in design. As HCI has served as motivation for design choices, deconstructing important concepts and thoughts from this research field can motivate design-related choices in this project.

Human-Swarm Interaction (HSI) is an emerging research field, with lots of new possibilities arising along with it. Since the swarm interaction research field is relatively new, and swarms have not found their specific use yet, it is challenging to establish concepts and regular practices when developing an HCI related product. Human-Swarm Interaction is a broad research field, such as HCI, and can cover many aspects even slightly related to the interaction between humans and swarms. Whereas with HCI, a user interacts with a single computer, with HSI, the user interacts with many agents (computers). Would these computers be controlled centrally, or would each agent in the swarm be capable of making decisions themselves, creating decentralized intelligence?

In section 2.5.4, a literature review will extrapolate concepts from HCI and transform them into the field of HSI. The overlap between HCI and HSI allows concepts within HSI to be introduced.

2.4.2 Development Ecosystem

When considering all the environments that can serve this installation, there are thousands of different possibilities. All the light installations Lumus Instruments have built over the years are developed in the same ecosystem. The modular and small form factor they work with has led them to use microcontroller-based systems. Lumus Instruments installations require a lot of computing power and dynamic, easy control over the lights. Over the years, they have had the best experience with the Teensy and FastLED library combination.

This project's goal is to stick to the Teensy and FastLED ecosystem. The project's smart sensing aspect requires hardware outside of this ecosystem to be compatible and easy integrability. Also, Lumus Instruments will be continuing the project themselves after this graduation project, making it easier for them to develop within a familiar ecosystem.

2.5 Relevant Research

Following from the relevant research fields in the previous section is the deepening relevant research. The interesting development frameworks and tools will be further researched in the ideation phase. The deepening relevant research utilizes and deconstructs concepts established in the relevant research fields section. The contextual awareness, engagement and enticement, and calm technology are important aspects that will be discussed. The deepening part will go over what these concepts would mean in the context of Symbiosis. Lastly, a reflection on the human-computer interaction and human-swarm interaction will be provided.

2.5.1 How to design contextual awareness for agents?

Context awareness [7] is an utterly important aspect of Symbiosis. The agents must be aware of their context and adjust their behavior accordingly. There are no data lines between the agents for the exchange of information. There are several ways to collect data for analyzing the environment the agents reside in.

There are three main components of context awareness; the ability to see, hear, and feel. The ability to see is of significant influence on the context we perceive. To implement a camera module in each agent seems like a logical choice to make. The data output from camera modules can provide object detection, movement, RGB levels, and light levels. Another aspect is the ability to hear. The decibel levels of the context can be metered with a microphone. This can provide us information about the context's loudness, as the averaged decibel levels will be lower in a remote place than in a highly visited gallery. These dB levels can reflect, e.g., the intensity of the light and sound output of the agents/swarm. With a multitude of infrared sensors, individual agents can perceive motion close to their bodies. Four IR sensors, two on each axis, can already provide sufficient context. Is something approaching the agent, or is there perhaps a constant distance to neighboring agents?

In addition to these three main components, several other sensor inputs either support the main elements or get other, less common data inputs. One of these would be a light sensor, such as an LDR, digital light sensor, or sunlight sensor. This sensor can support the camera module in determining day time and detect local (sun)light levels. Another possibility would be to have a temperature and humidity meter, which can add another dimension to the contextual awareness and reflect this in colder/warmer light representations. Gas sensors (O2) combined with the temperature and humidity sensors can provide extra depth to the contextual understanding. They can also help determine more polluted areas such as next to the train station instead of in the woods.

Haptic feedback incorporated in machines seems to evoke human emotions and relations. The physical buzzer, combined with a touch sensor, can mimic communication between the agents and the human touch. The touch sensors can detect contact from humans and can, in any desirable way, 'communicate' back with haptic feedback.

2.5.2 How to incorporate engagement and enticement in installations?

When developing new technology, maintaining interest in the specific technology got high priorities. The novelty effects, engagement, and enticements are essential to consider when creating installations for the public.

The novelty effect [13] is a common term in the human performance context. The novelty effect addresses the phenomena of a slight improvement of use when encountering and interacting with technology for the first time. This can be triggered by many different elements of technology; new features, other interactions, or a completely new design. However, the novelty effect is an illusion, and the slight improvement of use decreases quickly after the first encounter.

There are three different activity spaces when it comes to engagement [15]. The first level is the peripheral awareness level. Users in this activity space are aware of the installation but do not know much about it. The second level is focal awareness. People are talking about it, gesturing to it, and seeing other people use it. The third level is direct interaction, wherein people directly interact with the device. An unaware activity space can also be added before the first level and is when the people are not aware of the installation's presence.

An important aspect to consider is in which activity space the installations' creator wants the user/visitor. Specific enticing triggers can influence the interaction with the user and the engagement level the user is in. As discussed by [2], a particular threshold of participation is influenced by external factors and enticement. A crucial external influence that increases the threshold is social embarrassment. Brignull [2] adds that the visitor should be able to transition between participant and onlooker seamlessly and comfortably. People enjoy the simplicity of walking to an installation and using it the same way they saw the previous participant using it.

Another influence on the participation threshold is the location of the installation [18]. The engagement in busy places like train stations or metros has lower retention rates. However, areas that are too remote will also require low engagement since the initial threshold to get there is too high. The honeypot effect also greatly influences the engagement of visitors. The honeypot effect occurs when bystanders watch other people participate and interact with the installation [26]. It is a social learning influence and allows the audience to observe and participate in the installation more easily. One other important aspect is the dropout of participation. Participants should be able to drop out of the interaction without serious repercussions [26].

Regarding engagement and enticement, there are several aspects to consider when designing installations for public display. At which engagement level does the creator want the user? After considering the engagement levels, several enticement triggers and external aspects influence the participatory audience state. The installation's location, social embarrassment, the honeypot effect, and dropout repercussions are the most important aspects influencing the participation threshold.

Complex Adaptive Systems

Complex Adaptive Systems are systems that are not fully chaotic but also not linear (MacDonald). These systems have a loosely defined framework and are therefore very flexible. Most of the agents within such a system are locally aware of, mostly, a few simple rules. CADs are often used to describe ecosystems. Since there is no central control, the agents' behavior within the systems is harder to predict. The outcomes are emergent, involved, and responsive. CADs are capable of providing a non-repeating storyline. The CADs are of great use to keep engaging the audience since the novelty effect is actively avoided, and the agents keep responding differently to users.

2.5.3 How to encalm technology?

Calm technology focuses not on the 'in-your-face' interaction with technology but focuses on the user's peripheral. The main goal is to supply the user with information if necessary subtly and to stay in the background at all other times [25]. There have been identified several core principles to establish a framework for calm technology:

1. Technology should require the smallest amount of attention.
2. Technology should inform and create calm.Technology should inform and create calm.
3. Technology should mainly focus on the periphery.
4. Technology should amplify the best of technology and the best of humans.
5. Technology can communicate but does not speak.
6. Technology should work even when it fails.
7. The right amount of technology is the minimum needed to solve the problem.
8. Technology should respect social norms.

To encalm technology, the user's periphery is crucial [24]. The calm technology should easily move from the center of the periphery to the back. Video calls on your phone are frequently at the periphery center, whereas larger screens as your desktop allow for more comfortable video calls. We can experience facial expressions and know who we are talking to and who is not in the conference. A result of calm technology is to make us feel comfortable and familiar. When our periphery is stimulated just right, we feel excited about what we have experienced, are currently experiencing, and what we are going to experience. This connection to the context is an essential aspect of calm technology, referred to as locatedness.

The Symbiosis project will include numerous calm technology aspects that are relevant. The technology requires the smallest amount of technology. The creatures within the swarm do not even need attention from humans at all. They can self-provide with inter-agent communication and have endless different behavioral outcomes without a human visitor's intervention. The creatures also focus on the periphery of human visitors. They are capable of passively living amongst the visitors, creating calm ambiances. Even when visitors are walking past the installations, the creatures can passively pick up relevant contextual aspects and use these to construct new behavior.

2.5.4 Building a bridge between HCI and HSI

Technology has inevitably made its way into everyone's daily life. Ever since we created technology, we have developed solutions, opportunities, and frameworks to use technologies to its best extent. Human-Computer Interaction has been a dominant research discipline that considers how humans can most effectively and intuitively interact with computers. With the expanded capabilities, better resources, but most profoundly, the higher demands of humans, we have started to evolve technology in the field of swarm robotics. Swarms can be useful in several areas, such as firefighting, combing an area, or entertaining people. For my graduation project, swarm installations will be used as an art installation, which should interact with the audience deeply. However, the interaction between humans and swarms is hugely different from the interaction between humans and computers, while swarms are computers. With HSI as an emerging research discipline, many aspects have yet to be determined, including possible frameworks and deployment opportunities. HCI has been around for much longer and a lot of design decisions are based on HCI concepts. Therefore, with the known concepts of HCI, this paper will give insight into *how we can translate concepts from HCI to the field of HSI?*

The first part of the paper will deconstruct the most dominant HCI concepts and HSI's critical aspects from early research. In the second part of this paper, HCI concepts will be adapted and translated to practically use and enhance essential HSI aspects.

Identifying Concepts in HCI

There are three main concepts in HCI. The first essential concept within HCI is that people can bond with the computer. Szalma [22] states that human factors in machines emphasize the relationship between human characteristics and machine characteristics. Such an approach treats the human component of the system as the psychological process inside the human. The application of human characteristics can contribute to better implementation of the design principles. These human characteristics will influence the way humans interact and look towards machines positively. Epley [6] adds that anthropomorphism is another aspect of HCI that enables to create a sense of efficacy to increase the apparent competence of interaction between computers and humans to improve technological agents' usefulness.

Anthropomorphism can facilitate creating social bonds to increase the social connection between the technological agent and human. As has been found by Epley [6], anthropomorphism evokes more social connection than agents without human features. This feeling of social connection causes the human to relate to something with feelings instead of an object. This feeling of connectedness has also been agreed upon by Sproull [21]. The differences in responses to text interfaces and face interfaces were remarkable. Both men and women were more aroused when talking to the face interface and presented themselves in a more positive light talking to the face interface. It also shows how the thoughts of an actual person you are talking to elicited social behavior, which corresponds with Epley's [6] findings. Although the participants were told it was a cued emotion, meaning it was not human, they felt more connected and at ease with the human interface.

The second important concept in HCI is the sense of credibility and confidence towards the computer. Credibility is a medium being reliable in its message and sources according to Burgoon [3]. The influence may seem subtle. However, in reality, whenever a recipient fails to comprehend information, the information will never be used in their argumentation in any context. The study also reflects that humans assess and judge computers and their credibility in the same way they do with humans. It is also stated that the participants rated the credibility of the computer on par with human subjects. They believed that the arguments from the computer were better informed than their own opinions. Jiang [12] adds to that three communication elements that are psychological determinants of belief change. These are the source, message, and receiver. One of the critical factors is discrepancy, which is the distance between the source's position and the receiver. The lower the discrepancy, the greater the acceptance of the receiver. That means that when the source and receiver share common beliefs, they would more likely accept the sender's proposed information since the discrepancy is low. The receiver's self-confidence is also of great importance since this influences the likeliness of accepting the answer. Also, the perceived reliability of the source is of great importance. This perceived reliability factor from Jiang [12] closely relates to the credibility aspect proposed by Burgoon [3]. When the receiver judges the sender as credible or reliable, the acceptance will be greater.

The third concept is the need for human autonomy, competence, and involvement. Autonomy and competence can be facilitated through technology in several ways as stated by Szalma [23]. One of the ways is with the intuitiveness of the controls of the system. This closely relates to traditional HCI aspects, where the computer interface and how this interacts with humans is central. The ease of use of an interface positively affects intrinsic and extrinsic motivation. The experience of presence in the interface is also being related to an increase in competence. The rationale of activity or interface influences the autonomous motivation. The psychological aspects of human involvement are based on psychological needs. These psychological needs are in line with Jiang's [12] psychological communication elements. One of the psychological needs is the experience of pleasure and will be increased if autonomy and competence are experienced. The competence component is also influenced by the credibility and confidence aspects proposed by Burgoon [3].

Identifying aspects in HSI

To translate the concepts of human-computer interaction towards HSI, essential aspects need to be distinguished within the field of HSI to get a tangible grip on the research, opportunities, and possibilities. Human-swarm interaction is a new research discipline and requires thorough research to say something useful about these concepts and aspects. The first important aspect within HSI is the inter-agent interaction. A swarm of agents will have specific behavior programmed and will react according to that behavior. How one agent responds and interacts with another can be defined in several ways. Olfati-Saber [17] states that within multi-agent systems two interaction types can be determined: lead-by-attraction and lead-by-repulsion. They correspond to traditional leadership and predator-prey relationships, respectively. Goodrich [11] adds to that another inter-agent relationship: lead-by-orientation. In that way, repulsion and attraction models are incorporated, and class-agents and types are played with. For each inter-agent interaction model, different workloads from human control are expected. The lead-by-repulsion model drives the agents apart, demanding a higher human controller workload.

The second aspect of HSI is communication and control between agents and humans. The way agents organize and react depends on the behavior they inherited and adapted from other agents within the swarm. Kolling [14] also states how the communication between the agent and the human controller is one of the most challenging swarms aspects. A human controller must have all relevant information available to control the swarm if unexpected behavior occurs since the automated agents do not have control over interaction outside the algorithm that is programmed in them. An essential aspect of communication is the type of agents that are within the swarm. These types determine how the agents respond to human and external inputs. Goodrich [11] adds four different agent types within the human-aware and the human-blind category. Within human-aware agents, there are special agents and stakeholder agents. Special agents are influenced by human input only, whereas stakeholder agents are influenced by human input and other swarm agents. Among human-blind, there are type aware and type blind agents. Type aware agents are influenced by other agents and influenced differently per type. Type blind agents are influenced by other agents only and make no distinction between the different agents. This is only one way of interaction and control between humans and agents in swarms, and there is no one definitive way, yet.

The third aspect of HSI to consider is the (level of) autonomy of the swarm. Sheridan [20] configured a 10-point scale of autonomy within human-computer machines. Ten means full autonomy of the computer, whereas one means that humans got full control. Swarm robotics is a seven on this scale according to Sheridan, based on his experience with swarm robotics. Kolling [14] states that flexible levels of autonomy could be beneficial. That means that the swarm or human operator can take more control at a discrete moment in time. Dorais [5] states how the level of autonomy is of great importance for human-centered autonomous agents. He explicitly focuses on labeling autonomy levels and the communication of autonomy levels between the system and the swarm. The human operator must be aware of the system's current state to make useful adjustments. Essential aspects of the autonomy levels need to be considered when deploying and designing swarm robotics. What tasks can be executed by the swarm, and what tasks can be executed by humans? Setting the autonomy levels of the swarm and determining who controls this level once set are important aspects.

Discussion

The goal of this review was to identify important HCI concepts that could be relevant for swarms, identify essential HSI aspects, and to build the bridge between HCI and HSI. The research above suggests three concepts important from within HCI and three essential aspects of the HSI. Within HCI, the bonding with computers, the sense of credibility towards computers, and the need for human autonomy, involvement and competence are essential concepts. For HSI, the crucial aspects are inter-agent communication, communication, and control between human and agent/swam, and the swarm's autonomy level. The first translating aspect is the bonding between humans and computers. This concept within HCI can be directly translated towards the HSI domain. The human that is being interacted with must be part of the system. When humans interact with swarms, it will also be essential to establish connectedness between humans and swarms, or agents. The distant feeling humans have towards disruptive technologies can be limited by establishing this social bond. The bonding concept from HCI translated best to the communication and control between the human and swarm. The inter-agent communication aspect and the swarm's autonomy require no social bonding since the agents have no emotional feeling and response to bond to. Within the communication and control between human and swarm aspects, the human aspect is still involved, and bonding between them has beneficial effects.

The second translating concepts are credibility and confidence towards a computer and critical importance in the HSI domain. Depending on the application, the communication between swarm and humans can adapt this concept the best. The effect swarms have on humans and how human control/interaction with agents are positively affected by the sense of credibility and confidence. Incorporating design principles to enhance the feeling of credibility and confidence would certainly directly apply to HSI. The credibility also affects the inter-agent communication since the agents have individual credibility states. Once an agent is confident of action, it can trigger an action that can trigger another action, creating a chain within the swarm. This case's credibility gets out of the human-swarm context but is still relevant for the inter-agent credibility. The credibility is highly intertwined with the autonomy level of the swarm. One could imagine how the mutual credibility among agents influences the autonomy levels and, therefore, the swarm dynamics.

The third translating concept is the need for human autonomy, competence, and involvement. This concept translates to the interaction between the communication and control between swarm and humans mainly. As mentioned in the research, autonomy and competence can be facilitated through intuitiveness. The interaction's intuitiveness can positively influence intrinsic and extrinsic motivation within the communication between agents and swarm. These concepts do not translate to the swarm's inter-agent communication and autonomy because those aspects of HSI are not related to human influence and intervention. With an increased human autonomy within the system and interaction, the system's state is more comfortable transferable to the human operator/interaction.

When considering HCI concepts for translation to HSI, we can identify one main red line.

The concepts of HCI translate best to the HSI aspects with the communication and control between humans and swarms. The swarm's inter-agent communication and autonomy level are specific to HSI because they relate to the swarm aspect and not so much to humans' interaction. This shows that most HCI concepts do not apply to the fundamental swarm component central in HSI. Therefore, future research should dive deeper into the human influence on inter-agent communication and how agents can effectively communicate. Next to that, determinants for autonomy levels should be discussed, and a possible framework should be proposed. In relation to my personal GP, how one human interacts with many agents in a swarm should be researched thoroughly. This is another form of intelligence than the standard one to one interaction commonly found in daily life.

2.6 Requirements

The project's requirements originate from the results of the literature and a brief discussion with Lumus Instruments. The conceptual requirements focus on contextual awareness, combining the sensor inputs and the lights and sound output. An important aspect is the resonance and resemblence of the environment, closely related to the engagements and enticement research outcomes.

The requirements described below are requirements strictly for the graduation project prototype, which focuses on the inter-agent interaction. There is an overlap between the graduation prototype requirements and the final installation, but this will be further discussed in section x.

The physical and budget requirements are purely from Lumus Instruments. Lumus Instruments determine the system and budget requirements because the power grid supply is essential to the technology's novelty. It refers to all the agents in the swarm being powered by the same power grid. The budget is pre-determined by Lumus Instruments.

2.6.1 Conceptual

1. The creatures should communicate using actuators and sensors only.
2. The creatures express themselves using dynamic LED lighting.
3. Using sensor input, the creatures should 'be aware' of neighbouring creatures and characteristics, and decide whether to respond with complementary behavior, resulting in inter-agent interaction.

2.6.2 Physical

1. A power grid is the only thing connecting the creatures. No physical data lines should be used in between creatures.

2.6.3 Budget

1. A single creature should be priced in the range of 100 - 600 EUR. 100 EUR for creatures that are strong in numbers / swarm behaviour. 600 EUR for more complex creatures in terms of learning capabilities.

3 IDEATION

The ideation phase of the project focuses on concrete ideas and implementations following from the exploration phase. The ideation starts with an experience design. This experience design has carefully included elements from the exploration phase and has been constructed with Lumus Instrument's view in mind. Following the experience design, creature behavior will be deconstructed in detail. The ideation phase will conclude with evaluations of hardware and technical setups.

3.1 Experience Design

The experience design is a crucial determinant for the technical implementation of a possible prototype. The experience design will cover aspects of the experience users will get when 'visiting' the installation.

The basic setup for the installation will consist of a multitude of creatures. - somewhere in the region of fifty to two-hundred. A single unit will be this circular shape of organic, black acryl and positioned on a stand. The height of the stand will differ, as well as the creatures' separation, as seen in figure 3.1 on page 32. This separation and height difference will create a 3D field effect of units. This field of creatures will be around twenty meters wide and at least the same length, perhaps even longer. Not only will the creatures be able to be placed on a stand on the ground, but they will also be capable of being fixed to the wall or the ceiling.

The experience design will focus on the installation's final experience. This graduation project's prototype is the inter-agent communication and awareness part of the total experience design described in this section.

The installation will be placed at light festivals. It is common to walk around on a large exhibition at most light festivals, where all the different installations will be placed. This could either be in a single room or the open space and outdoors. Symbiosis will be a large installation, which will immediately grab the visitors' attention as soon as they lay their eyes upon it. As mentioned beforehand, there will be a 3D field of creatures creating the swarm. The visitors

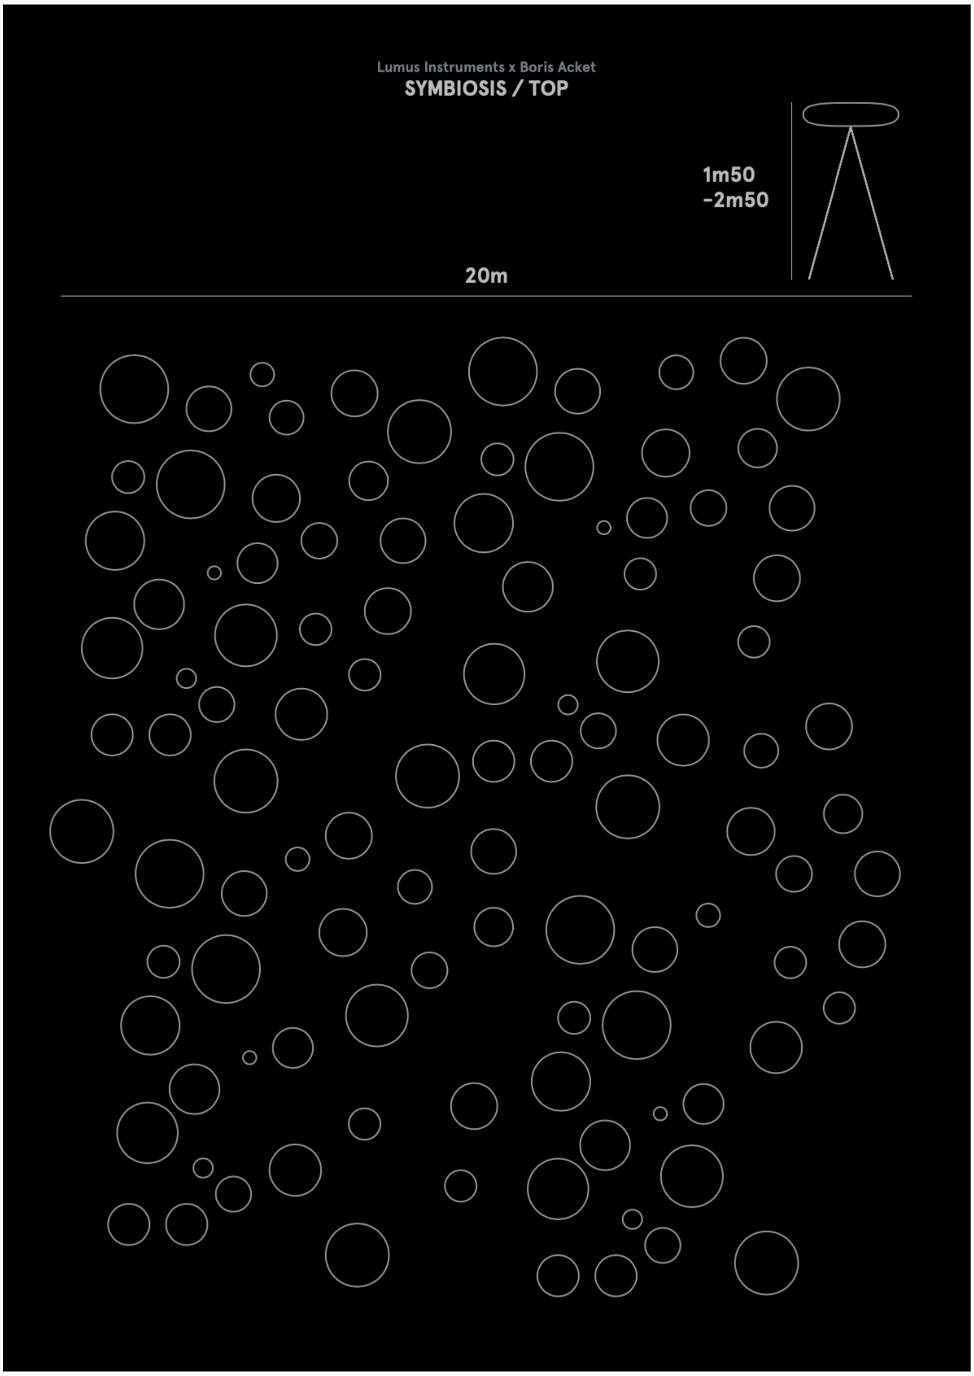

Figure 3.1: Symbiosis - Conceptual Construction

will be overwhelmed by the large installation and curious to see what it is and what it does. When they walk through or past the installation, they can look wherever they want but seem to be captivated by the organic appearance they create. The swarm seems to represent the environment - they function as an amplification of the context.

There is more to the agents than just representing and amplifying the environment they reside in. When one strolls past the installation, a creature, or more of them will be aware of your presence as a human. It will recognize what colors reflect you (clothing) and your personal space (personal attributes). The creature can also track your coordinates. As you stand still in front of one of the cameras, the creature will be fully focusing on you - after all, that specific creature's context is influenced predominantly by you at that point.

However, this creature observing what is happening around it is only a single link in the whole swarm-chain. The creature will output light and audio as a response to its context. When a visitor wearing red is occupying one or multiple creatures' vision, this creature will, e.g., incorporate red more intensely in their behavior and use red as a dominant color in their output. Depending on the context, the creature will process the input differently in its behavior. Are the color coordinates rapidly moving over the screen, causing anxiety or anger? Then the creature will amplify this by utilizing the lights more intensely.

Each agent in the swarm is identical in the sense that they have equal hardware. They can sense the same and produce/output lights and audio in the same manner as the other creatures. What differentiates the creatures in the swarm is the position in the hive and the exact context. Aside from sensing visitors, they will also be aware of neighboring creatures. When visitors walk through Symbiosis, there is a high probability that multiple creatures will sense this and output their interpretation of the influence this visitor has on the creature's context. This creature's output will trigger neighboring creatures because the neighboring creatures' context is affected by the creature that initially got affected by the visitor. This reactance chain can be started by the (unaware) visitor(s), giving them strong but subtle interaction opportunities.

Aside from visually experiencing Symbiosis, there will be an immersive audial experience. Each creature is equipped with a microphone capable of picking up environmental sounds. Visitors can vaguely rediscover the environmental ambiance in the voices of the creatures. The direct sound input seems warped, distorted, and filtered - creating an organic atmosphere. If there are peak amplitude discovered in the environment when you, e.g., scream or honk, the creature will adjust, digest the sound, and process it in their vocabulary to merge it with their contemporary ambiance.

When there is no peak interaction, and the creatures reside in a quieter environment, they still represent their environment. At night, the swarm will not sense as much as during day time. The creatures will construct their behavior from what they have sensed previously and use this to adapt to the environment. At night it will become a luminescent, harmonious swarm to witness. When you are watching it from a distance, it becomes this sea of elegantly breathing organisms. If they reside in more living areas, they will be more active during nighttime. If they live in sunnier environments, their colors will become brighter. If they experience unstable weather with lots of changes, their behavior will be more unpredictable.

Biological-inspired behavior will be an essential aspect of Symbiosis. As with birds flocking, several natural rules establish the respective behavior. For flocking, these are separation, alignment, and cohesion. The ways agents in such a swarm behave are forced upon by nature - increasing survival chances. Symbiosis will adapt to a form of biologically-inspired behavior, giving the visitors the feeling of watching a living, adjusting installation.

Symbiosis will also be capable of several other deployment options. The first, as discussed above, is the field-like approach. Visitors can walk through this field of creatures, where they surround them. However, with the modular setup of Symbiosis, it will also be capable of deployment in other shapes to create different experiences. It could be utilized in a tilted, circular shape, but also as two straight lines at the same height. This setup deviates from the regular, perhaps noisier field-setup, but allows for interesting behavior. How will such a swarm adapt to the unusual clinical setup? Neighbors can only be found on a single axis, forcing signals to pass linearly.

3.1.1 Relation to Exploration

Section 2.5 gives a thorough overview of the experience design. However, this is the experience design, as discussed with Lumus Instruments, and does not explicitly contain the relevant elements' exploration phase. This section will briefly cover how the broad exploration will be incorporated into the envisioned experience/installation.

The contextual awareness discussed in section 2.5.1 will profoundly find its purpose in using a camera and microphone. The camera and microphone will function as the seeing and hearing aspects of the creature, which covers two of the three essential aspects of contextual awareness. Each creature in the swarm will have a constant live-feed, providing the creature with audiovisual data of its surroundings. At later stages, the contextual awareness can be extended by adding other sensors to support the hearing and seeing. These sensors can, for example, include humidity sensors, light sensors, or fluid gas sensors.

Factors discussed in section 2.5.2 influence engagement and enticement as the novelty effect, the intended activity space, participation threshold, and the installation's location. One of the primary aspects of the installation is that creatures construct their audiovisual behavior using sensors. The data from these sensors will likely always fluctuate, making the behavior from the creatures not static. The creatures' ever-changing behavior helps counter the novelty effect, as the visitors will never know exactly what to expect. Also, the intended activity space is subtle, as the swarm must live its own life without relying on human intervention. The subtle activity space does not require the visitors to interact with the installation actively. However, users will always subtly be interacting with the installation when being near, making the participatory threshold extremely low. The installation location will also be great since it will be displayed at light festivals where people are genuinely interested in the installations.

The low participation threshold goes hand in hand with the calm aspects of technology discussed in section 2.5.3. The envisioned installation complies with critical aspects of calm technology. The swarm requires the smallest amount of attention since it will not be dependent on human input can 'live' entirely on its own. The technolog y also focuses on the periphery, making the interaction between installation and visitor passive rather than active. Visitors only need to wander through the installation to be picked up by the creators. The creatures can communicate, especially with other creatures, without speaking. It senses using a camera and microphone outputs this audiovisually. The audio output will also be subtle and nowhere near regular speaking.

3.2 Creature Behavior

One of the key technical aspects of the creatures/swarm is contextual awareness. Contextual awareness can be divided into peak behavior and adaptational behavior. The peak behavior covers the harsh sensor inputs and the deducted action from these peaks. The swarm's adaptation refers to the gradually shifting behavior related to the environment the swarm is in.

3.2.1 Peak Interaction

The peak interaction of the creatures will be essential. There are several critical factors tied to the peak interaction to consider. The first one is the sensor inputs and data to work with. The creatures will likely have at least a camera and a microphone to pick up the environment's information. These sensors can be triggered by human interaction but also by neighboring creatures. This could indicate no direct communication protocol for inter-agent communication, but they solely rely on their senses to construct behavior.

A possibility is to have installations settle within their context and absorb data for a couple of days straight to develop and train a neural network model. Once the model is sufficiently trained, it can be deployed and used in the installation's context. The environment will construct the behavior of the creatures. How the creatures will evolve individually and as a swarm will depend on environmental triggers and can create, for example, anxiety, anger, happiness, or excitement. One of the possibilities is to use the training week for supervised classification or clustering using the k-nearest neighbor algorithm.

When a creature senses a specific trigger, it can react accordingly. When the camera senses much blue and the decibel levels are below a particular threshold, it can extract behavior, which can be detected by the neighboring creatures and replicated. This reaction implies explicitly for peak interaction. When a threshold is abruptly exceeded, the creatures will perhaps display specific behavior based on previous experiences. This behavior is quickly picked up by neighboring creatures, combined with the sensory data of that particular creature (microphone also picks up loud dB level, e.g.) and sets of a chain reaction.

3.2.2 Adaptation

The adaptation of the creatures and swarm is closely related to the peak behavior of the swarm. Whereas the peak interaction focuses on the direct interaction with others, the adaptation will be more towards developing the swarm's behavior over a more extended period. It is vital that the swarm's behavior reflects the environment it resides in, but needs time to adapt to this environment. The sensor inputs construct the adaptation. With the adaptation occurring over a more extended period, the sensors that would typically not fluctuate significantly in a shorter period, but do so over a more extended period, will be of great use in adaptation. This would be humidity sensors, pressure sensors, and temperature sensors. The use of these sensors for the long-term (two/three weeks) adaptation also allows organism-like behavior.

One of the possibilities to let the swarm adapt to its context is to utilize sensor stamps. Multiple sensors are used for the creatures—the peak interaction sensors as the camera and microphone and the adaptational sensors as humidity and temperature sensors. Generating an array of sensor data stamps every 30 minutes (or so) creates an average over the course of days, which can be represented in the swarm's lighting and audio behavior.

When the swarm moves to a new environment, the sensors will get different inputs. The behavior will gradually adapt to the new environment. Was there much pollution in the previous city, and was it a cold city? If the installation came over from a colder city, the swarm would perhaps have more blue colors included in its behavior and react slower. When people visit the installation right after it has been moved to a warmer city, it will still include some characteristics (such as more blue for colder areas) in its output since it does not immediately adapt and reflect the new environment it is in.

3.3 Technical Prototype

The experience design aligns with several requirements established in section 2.6. The conceptual requirements allow for more concrete technical requirements to be deducted. The design experience covered in the previous section creates a blueprint for the requirements of the technical prototype.

The primary modality of the prototype will be seeing. The creature will sense another creature and its emotional state and characteristics. The creature senses this information with a smart camera module.

3.3.1 Smart Cameras

The creature's most essential feature will be seeing to develop contextual awareness. Intelligent behavior will be crucial, so microPython cameras or microcontrollers will be considered because the high-level Python allows for more straightforward AI implementation. Microphones will also be considered, although the hearing modality will be considered beneficial and not mandatory for the prototype.

OpenMV Cam H7

The OpenMV Cam H7 (A) in figure 3.2 is a machine vision camera that operates on a small, low-power board. The form factor allows the camera to be subtly integrated into any design. One can program the camera to its liking with microPython and manipulate the inputs and outputs in any desirable way. The board comes with a built-in interface library, which allows easy connection to other microcontrollers such as the Arduino. Also, TensorFlow Lite is easy to integrate. With an 80MBs bus, live video and pictures can be streamed to other devices. The micro sd card can be used to provide machine vision assets to the program. There are 1MB SRAM and 2MB of flash memory available. Consumes 170mA @ 3.3V. The price is $65.

OpenMV Cam H7 Plus

The OpenMV Cam H7 Plus (B) in 3.2 is almost identical to the standard version. It runs on another processor, with 32MBs SDRAM + 1MB of SRAM and 32 MB of external flash + 2 MB of internal flash. Consumes 240mA @ 3.3V. The price is $80.

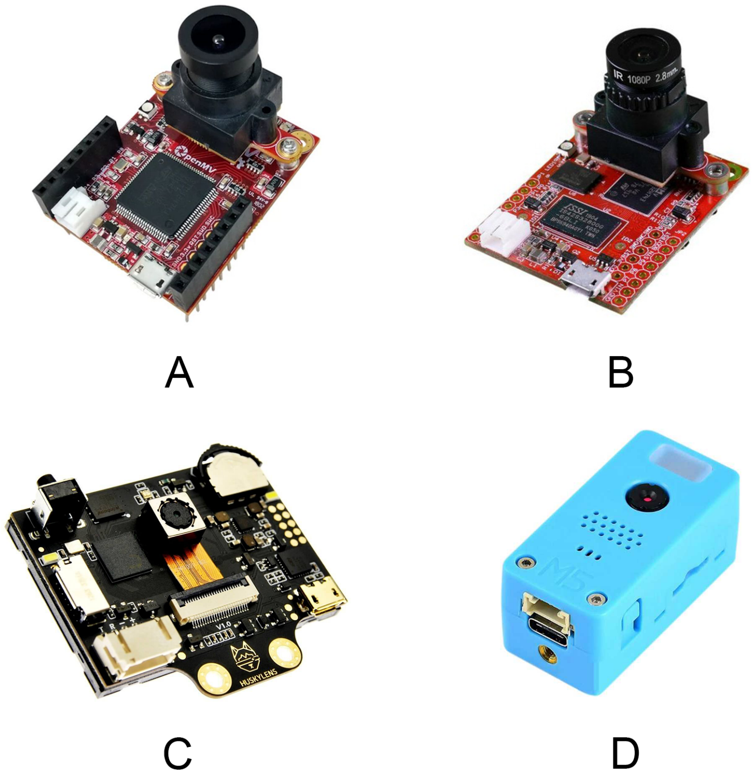

Figure 3.2: Smart Cameras Overview

HuskyLens DFrobot

The HuskyLens (C) in 3.2 is also based on the Kendryte K210 SoC. It also features the basic machine vision applications as face recognition and object tracking (color, location, size). The HuskyLens has a 2” IPS screen, which allows for easy parameter configuration. For the HuskyLens, there is a built-in library within the Arduino IDE. This allows the Arduino to control the input/output of the HuskyLens. It consumes 230mA @ 5,0V, and the price is €50.

M5StickV K210 AI Camera (Without Wifi)

The K210 by M5StickV (D) in 3.2 is also a small, low-power camera module powered by the Kendryte K210 SoC. It is focused on neural network calculations and supports the detection of coordinates, size, and type of target. The K210 has a dedicated neural network processing (KPU) unit. It has 8MB of SRAM and 16MB flash memory. It supports microPython and has a field-programmable IO array, allowing it to be programmed and configured using another microcontroller. It has no built-in interface library but offers UART / I2C support. The camera module has a built-in screen and overall consumes 500 mA @ 5V. The price is $29,90.

3.3.2 Microcontrollers

Teensy 4.0

The Teensy 4.0 (A) in 3.3 is already used by Lumus Instruments, combined with Open Sound Control (OSC) over ethernet and an Arduino. The Teensy has a clock speed of 600MHzm, which makes it suitable for rapid audio and lighting applications. It is also able to be used with Arduino using the Teensyduino library. The Teensy has an extensive audio library, which is way better than the Arduino UNO's standard audio library. The Teensy 4.0 comes in at €24.

Aduino UNO

The Arduino UNO (B) in 3.3 is a suitable option for the main microcontroller. There UNO is one of the most well-known microcontrollers and has broad connectivity. Many third-party modules can connect to the Arduino, and there are tons of external libraries to be used with the Arduino. The clock speed is way lower than the Teensy at 16MHz, making it more challenging to operate with rapid lighting and audio applications. The price is €20.

Arduino Portenta H7

The Arduino Portentia H7 (C) in 3.3 is part of Arduino's Pro line-up. The portenta has two parallel cores, allowing scripts to be uploaded and run simultaneously, interacting with each other. The processors run at 480 and 240 MHz, respectively. Arduino sketches can be uploaded, as well as microPython and TensorFlow Lite. High-level programming combined with a dedicated graphics engine ensures a more accessible machine learning application, with tons of ready to use libraries. The Portenta H7 costs €109.

Arduino NANO 33 BLE SENSE

The Nano 33 BLE (D) in 3.3 is the tiny AI-enabled board from Arduino. The board runs a 32 bit 64 MHz processor, with 1MB program memory and 256 KB of SRAM. The board has several onboard sensors, including inertial sensors, humidity, and temperature sensors and microphones. The data from these sensors can be used to train models using TinyML or TensorFlow Lite. It costs €29.

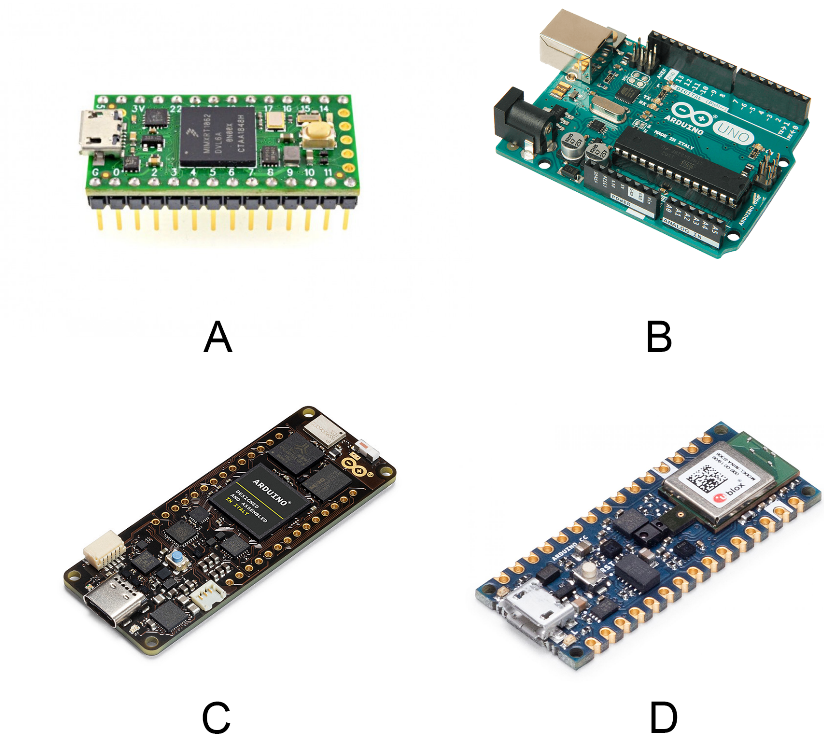

Figure 3.3: Microcontrollers Overview

3.3.3 Microphones

The microphones in this project are just as crucial for contextual awareness as the camera modules. Whereas the camera needs to be smart, the microphone simply needs to be excellent at the single purpose it has: picking up sound.

Portenta Vision Shield

The Portenta Vision Shield, which is also discussed in the camera section, also includes two onboard microphones. These are the MP34DT05 microphones, which is a compact omnidirectional microphone. It costs $1.30 per unit when bought separately but comes included on the Vision Shield.

Elektret Microphone MAX9814

The Elektret microphone has a bit larger form-factor than regular microphones. The microphone is soldered on the MAX9814 amplifier, which is designed for unpredictable audio. It has an automatic gain control, allowing the louder sounds to become quieter and the softer sounds to become louder. It functions as a compressor and has got an adjustable gain and attack/release ratio. It detects sounds from 20 - 20k Hz. The price of the microphone-amplifier combination is €9.

Adafruit Silicon MEMS Microphone Breakout - SPW2430

This silicon MEMS microphone is a tiny microphone that does not require an additional amplifier. The amplifier is built into the breakout board. It can detect sound in the 100 - 10k Hz area. The sound region it detects is substantially smaller than the Elektret microphone, which is useful to consider. Do the creatures need to hear frequencies above 10K Hz or below 100 Hz? This microphone comes in at €5,95.

Grove - Analog Microphone (MEMS)

The Grove analog microphone uses acoustic technologies to detect sound. The MEMS microphone offers extremely low noise and can provide 20dB of gain. It has a small form factor but finds its use in high-end audio applications as the ReSpeaker 2.0. The Grove Analog microphone costs €8,65.

3.3.4 Hardware Setups Evaluation

With the current knowledge of the available hardware components and requirements, a table consisting of several hardware setups are judged on several factors, making the choice of a hardware design easier. The hardware setup's most important factors will be the speed and RAM / ROM combination. The high speeds are essential for fast processing of data and controlling the LEDs, whereas the RAM / ROM combination is important to deploy a neural network on the controller. A hardware setup with a Teensy at its core will likely be chosen since this adapts well to the current ecosystem Lumus Instruments works in. That indicates that the vision processing will likely occur at the camera side.

This project focuses on the visual part since an external company will cover the audio part. That is the primary reason microphones are not in this project's scope, so the previous section on microphones will not be evaluated in the hardware setups. Also, at later stages of the prototype, the prototype's hearing modality will not be part of the prototype.

	Costs	Size	Power	User-friendly	Speed	ROM	RAM
AiY Vision [1]	-	-	-	+/-	+	+	+
Coral Dev Board + Coral Cam [2]	–	-	-	+/-	++	++	++
Nvidia Jetson Nano [3] + Pi Camera Module V2 [4]	–	+/-	+/-	+/-	++	++	++
OpenMV H7 [5] + Teensy 3.2 [6]	++	+	+	++	+/-	-	-
OpenMV H7 Plus [7]+ Teensy 3.2 [8]	+	+	+	++	+/-	+	+
ESP EYE DevKit [9]	++	++	++	-	-	-	-

Table 3.1: Hardware Setups Evaluations

4 IMPLEMENTATION

The implementation phase will translate the theoretical knowledge acquired in previous sections into physical prototype iterations. The prototyping part will continuously be evaluated and reconsidered to get to a minimal viable product, which is what the implementation phase will be concluded with.

4.1 Prototyping

The implementation phase of the project covers several aspects of the development of a sophisticated prototype. The implementation phase starts with the exploration and iterations of the prototyping phase, beginning with the design's basic functionalities, evaluating the prototype iterations, and adjusting them towards a final prototype.

4.1.1 Face Recognition and Environmental Awareness

The prototyping phase starts with the implementation of the basic features required. These are the machine learning aspects of the OpenMV camera and the sending of data to a microcontroller, storing them in variables.

Materials

Component	Amount
Arduino Uno	1
OpenMV M7	1
Breadboard	1
Jumper Wire	3
USB-B - USB-A Cable	1
micro-USB - USB-A Cable	1

Hardware Design

The basic implementation's hardware design consists of the Arduino Uno and OpenMV M7 connected to the computer via USB, and the OpenMV M7 connected to the Arduino Uno with the Rx and Tx pins. The jumper wires connecting the OpenMV M7 to the Arduino Uno were long enough to freely point the lens of the OpenMV M7 into any desirable direction.

Software Design

The software consists of the OpenMV part and the Arduino part. The OpenMV M7 is responsible for capturing sensor data, processing it, and sending it to the Arduino. The Arduino is responsible for receiving the data via the serial monitor and storing the serial data in the appropriate variables.

The OpenMV M7 runs a microPython script, which uses the pre-trained person detection network. This network continuously analyzes the frames of the camera and is capable of detecting faces. Once a face is in the frame, the boolean isPerson turns to 1. Next to that, basic contextual awareness is implemented in the form of the LABS color scheme. The lens's histogram object is called, and the appropriate L, A, and B mean are stored in variables. These four variables are sent as bytes over a UART connection to the Arduino Uno. A screenshot of this implementation can be found in figure 4.1.

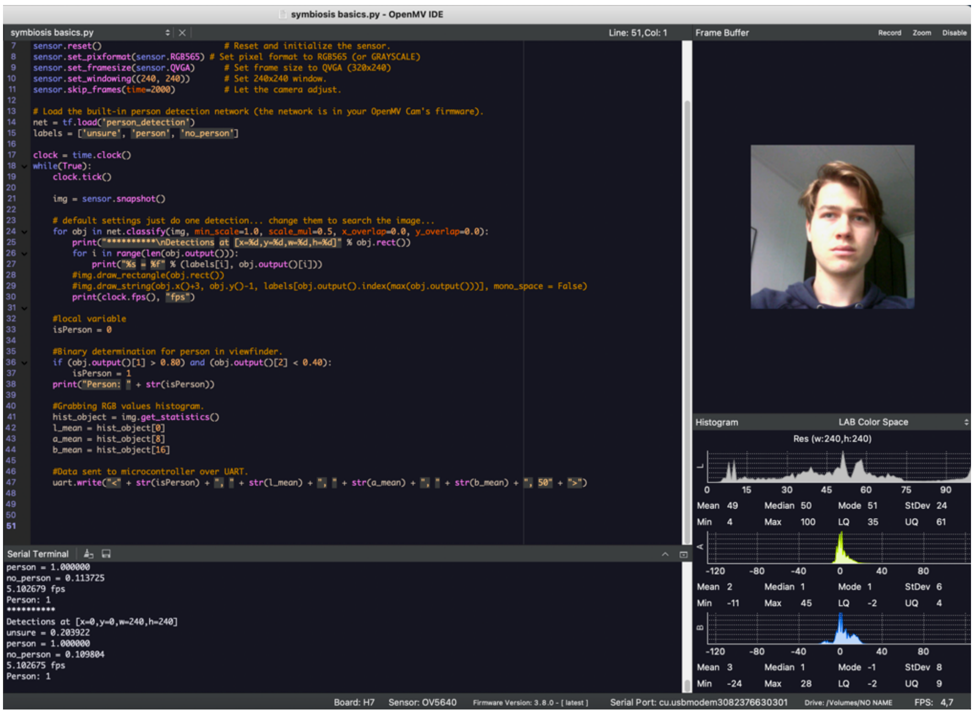

Figure 4.1: OpenMV Face Recognition

The Arduino Uno is listening to the serial port and recognizes start and end markers to sort the bytes coming in. A buffer of the incoming data is created and analyzed with the comma as a delimiter. These chunks of bytes are stored in tokens and converted to the appropriate variable and type.

Evaluation

This basic implementation of communication has proven to work conveniently in this prototype. The data OpenMV M7 sends to the Arduino has a minimal delay, and the Arduino can quickly process and store the information in appropriate variables. However, the network is only capable of detecting a face, which also needs proper lighting.

4.1.2 Creature Recognition Neural Network

This prototype focuses on the recognition of other creatures and their respective states. The technical challenge for this prototype is that there is no data grid connecting the individual creatures. This led to the implementation of smart camera modules, to implement deep learning algorithms for state recognition. An intelligent camera module aims to process live video and run these frame by frame through the trained model, which can then label the creature's corresponding state.

Materials

Component	Amount
Teensy 3.2	1
OpenMV H7 Plus	1
WS2812B 5050 RGB LED Ring	1
Breadboard	1
Jumper Wire	3
micro-USB - USB-A Cable	2

Hardware Design

The basic design setup for creature recognition was a circular led strip, placed behind black, diffusing acrylic, which can be seen in figure 4.2b. The led strip is driven by a Teensy 3.2, which is powered via USB of the computer. The acrylic plate, the led strip, and the Teensy with breadboard are attached to a wooden surface, keeping all the individual components in place, as seen in figure 4.2a. Black cardboard is used as a frame, blocking direct light from the outside.

The OpenMV M7 was set up on a cardboard stand with the lens facing the led ring. The OpenMV M7 was also powered via USB by the same computer. The OpenMV M7 captures training data for the deep learning model, while the led ring outputs specific, predefined behavior.

(a) Internals Led Ring Prototype

(b) Organic Behavior Prototype

Figure 4.2: Neural Network Training Prototype

Software Design

The prototype consists of three main software parts: the Teensy 3.2 code, the OpenMV M7 code, and the convolutional neural network structure.

The Teensy runs in the Arduino IDE with the TeensyDuino package, allowing the Teensy to be programmed precisely like a regular Arduino (Uno). In discussion with Lumus Instruments, four types of behavior were designed that are reflected in the lighting behavior of the led ring. Within TeensyDuino, these five types of behavior are created, resulting in the following behavioral types: blank, organic, strobe, random, and breathing. Teensy controls the led ring with the FastLED library. Most of the behavior types are simple programs, except the organic kind, which is more complicated due to the 2D Perlin noise implementation.

The OpenMV code in this part of the prototype is mainly to capture training data for the neural network. There is no implementation of the neural network yet, only gathering training data of the neural network. Since the final neural network will likely be deployed on an OpenMV cam, the network's accuracy would increase if the training data and input data are of equal dimensions. The OpenMV software was a microPython script that automatically captured two thousand screenshots of the sensor and automatically stored them in the appropriate folder on the inserted SD card. This script creates five folders on the SD card, with two thousand images of each respective behavioral state.

The neural network's goal is to recognize the behavioral states of neighboring creatures with only machine vision. The training data from the five respective folders is inserted into the Edge Impulse application via their CLI. T The input data is split into 80% training data and 20% test data. The training data images are squashed from 320x240 (QVGA) into 96x96 images and fed into a feature extractor, which extracts raw features and parameters. This results in a three-dimensional graph of the training data, which shows how well the images can be separated based on TensorFlow's RGB components. These features then function as the input layer of the neural network, consisting of six layers: an input layer, a 2D convolutional layer with 32 neurons, a 2D convolutional layer with 16 neurons, a flatten layer, a dropout layer, and the output layer.

Evaluation

This prototype's evaluation is based on the trained network. The statistics of the confusion matrix look promising, with a 99% accuracy. The network also works excellent on unseen test data with equal accuracy. The trained network is a TensorFlow Lite model, with a size of 74 KB. However, the TensorFlow model will be deployed on the internal memory of the OpenMV cam and will have specific peak RAM usage. The OpenMV M7, the camera module used for this setup, was incapable of utilizing the TensorFlow model. Therefore, this setup cannot test live situations and requires a more powerful camera module.

4.1.3 Creature Characteristics

The previous iterations of the prototype gave insights into the advantages and disadvantages of this implementation. The current network is trained on static images, making the network analyze the images frame by frame.

This approach's difficulty is that behavior, which consists of multiple frames, is now generalized to a single frame. So far, this results in several forms of label flickering. These flickering label artifacts can be tackled with the help of a rolling average, which creates a label output buffer and sums x iterations of the output layer before giving a final label. This buffer mimics the classifications as if they are videos instead of single frames.

The current behavioral states are also not thought out well since the previous prototype iterations only required the prototype to recognize states. There was no intention or focus on creating behavioral states that would be essential for the final product.

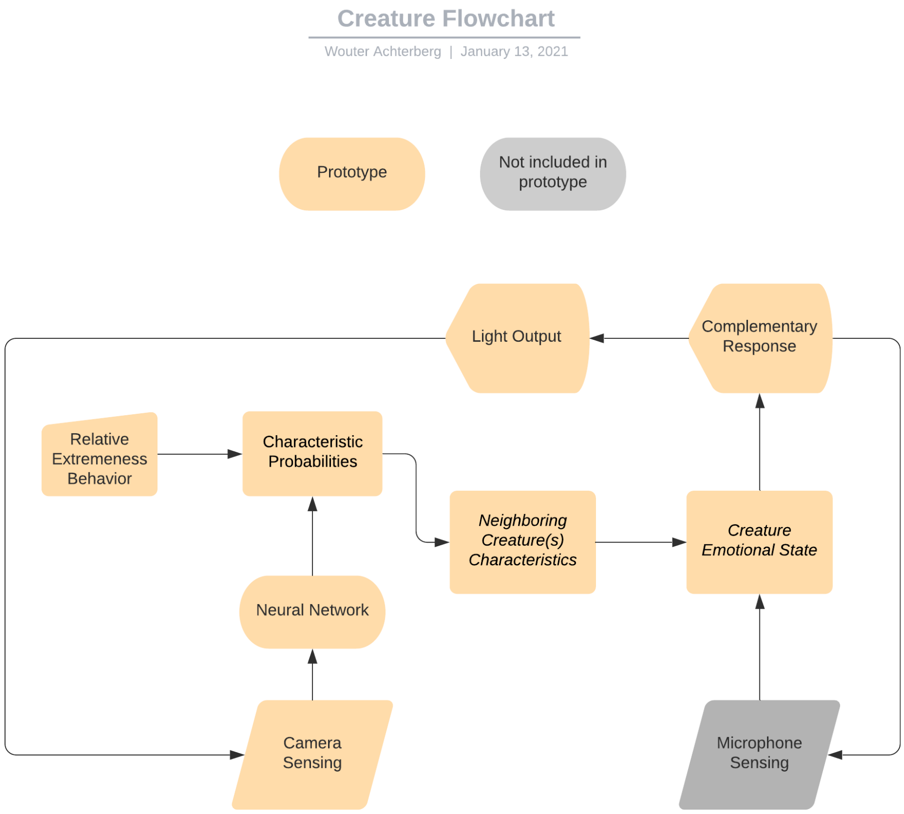

Figure 4.3: Creature Flowchart

Moving towards a more sophisticated prototype, a flowchart has been constructed to give a clear overview of the technical structure. This flowchart can be seen in figure 4.3. The creature consists of two main parts: the creature's characteristics and the creature's emotional state. The previous prototypes focused on the creature characteristics parts, which involve the probabilities, the neural network, the camera sensing, and the light output. These characteristics are from neighboring creatures and sensed with a camera and a custom trained neural network. These probabilities of characteristics influence the creature's emotional state, which responds with complementary audio(visual) output.

The next iteration of the prototype will be more directed towards a more polished implementation. As shown in the flowchart, the character probabilities are based on the relative extremeness of characteristics. The defined characteristics are smooth, random, chaotic, organic, graphical, and symmetrical.

Materials

Component	Amount
Teensy 3.2	1
OpenMV H7 Plus	1
WS2812B 5050 RGB LED Ring	1
Breadboard	1
Jumper Wire	3
micro-USB - USB-A Cable	2
Black Acrylic (30x30)	1

Hardware Design

The hardware design is identical to the previous iteration.

Software Design

This iteration of the prototype has some additions to the current software, with different lighting behavior corresponding to behavioral states. These will be smooth, random, chaotic, organic, graphical, and symmetrical. A new neural network has been trained with these labels, passing probabilities to the Teensy as parameters for the behavior.

Evaluation

This iteration of the prototype performs reasonably well. This prototype's main objective is to assess how the neural network would pick up the defined lighting characteristics. These probabilities are picked up well and sent to the Teensy as behavioral parameters. The probabilities are stable when the test data is similar to the training data. When the camera is moved around and thus moving away from the training images, more label flickering occurs. The next prototype iteration will focus on the interaction between two such prototypes and use the lighting characteristics to construct their behavior new behavior, developing constant interaction.

4.2 Minimum Viable Product

Currently, the prototypes have included a single creature. This creature only consists of predetermined lighting states. This creature is being monitored by the OpenM H7 Plus, which utilizes the trained deep learning model to obtain lighting characteristics. These values have only been read out in the OpenMV IDE to understand the predictions' accuracy in real life.

The last prototype will be the minimum viable product, which will be the prototype used to test and analyze previous chapters' set requirements. The primary purpose of the minimum viable product is to incorporate and combine previous iterations of prototypes. Another creature (creature B) will be built with equal hardware and software components as the previous prototype. Both creatures will have their camera with the trained deep learning model and will communicate the probabilities via UART to the Teensy.

4.2.1 Sensor and Actuator

There are two main components in the minimum viable product, each with its essential functions. The OpenMV cam is responsible for detecting the characteristics probabilities and passing these values over UART to the Teensy. It functions as the eyes of a single creature. The Teensy is the heart of a creature, which receives the characteristic probabilities. It is also responsible for controlling the led ring, which is the body of the creature. The controlling of the led ring will be done with behavioral functions. These behavioral functions rely on incoming lighting characteristics.

4.2.2 Behavior Construction Mechanism

A crucial aspect of this prototype iteration is the combination of previous iterations. Specifically, the software design of the behavioral construction of the lighting behavior using the lighting characteristics obtained by the OpenMV H7. The prototype's current iterations have mainly revolved around behavioral functions to create test data for the neural network. There was no software setup that made the creature function as an organism. The minimal viable product's software will be structured differently. The creature has several (emotional) states, of which he will always be in one and only one. Switching between states occurs by the shifting, incoming characteristic probabilities. This will function just like hormones work in human bodies; when we sense someone is angry, we will most often try to keep calm or get angry at them too. When the creature senses high chaotic probabilities, it will probably calm down or get chaotic itself too.

Character - State Relation

To get started with the mechanics' development, it is important to get an overview of which state relates to which characteristics and how strong the association is between the character and state. Together with Lumus Instruments, an overview has been created on how, intuitively, these relationships exist, which has been visualized in figure 4.4. The thicker the connection between the state and characteristic, the more intense the characteristic influences the state. The colors of representing the different characteristics have been chosen arbitrarily. There is no additional meaning to these colors and they are not specifically linked to certain behaviors. Each creature will always be in only one state. The characteristics are determining this state. The characteristics are probabilities that are being passed from the OpenMV H7 to the Teensy. Figure 4.4 shows how these characteristics intuitively correspond to which state, so there is no definitive or correct link between characteristic and state. This flowchart may be handy when programming the algorithm for the state determination of the creature.

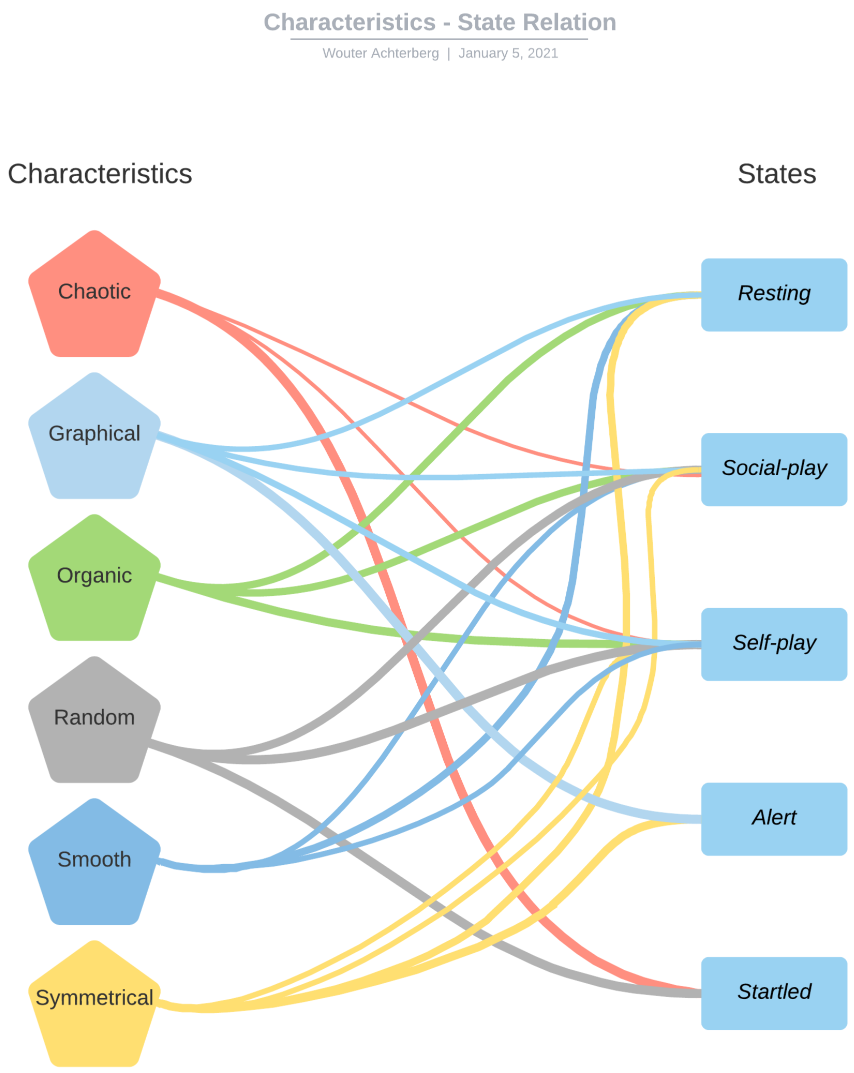

Figure 4.4: Character - State Relation

Figure 4.5 gives another overview of each state with its respective influential characteristics. This figure gives an overview of the backbone of the mechanism. It is also a great starting point for determining which parameters will affect which functions the most. For consistency, each state will be 1.00 maximum, and this will be divided among the state's influential characteristics.

Character - State Relation Table

Wouter Achterberg | January 5, 2021

Resting
Smooth (0.30)
Organic (0.30)
Symmetrical (0.20)
Graphical (0.20)

Social-play
Random (0.30)
Organic (0.30)
Smooth (0.15)
Graphical (0.10)
Symmetrical (0.10)
Chaotic (0.05)

Self-play
Random (0.30)
Organic (0.30)
Smooth (0.15)
Graphical (0.10)
Symmetrical (0.10)
Chaotic (0.05)

Alert
Graphical (0.50)
Symmetrical (0.50)

Startled
Chaotic (0.60)
Random (0.40)

Figure 4.5: Character - State Relation Table

Figure 4.5 shows how each state will logically consist of which characteristic. However, the most crucial factor is how to translate the incoming data to complementary behavior. How will a creature react if it senses much chaos? What will it do if it senses random and smooth characteristics?

Data Flow

The mechanism will function with predefined, hard states. The creature will either be in state A or state B. This indicates that a creature will not be able to, for example, be resting and alert at the same time. These states are equal to the states defined in figure x. Within each state, a specific effect generator will be active. This effect generator will behave according to the state. For some of the states, this will be relatively straightforward - the 'startled 'state will have an effect generator that will strobe with a given speed and intensity.

The characteristics that the OpenMV H7 passes to the Teensy are raw and estimated values. Upon entering the Teensy serially, they should be cleaned up and prepared to be used within the

Teensy code. This process will be referred to as smoothing. Once the characteristic values are smooth, they will be utilized for determining which state the creature will be in. If the creature determines lots of chaos, the chaotic value will increment. The creature will switch to the chaotic state, in which the effect generator of the chaotic state will be utilized to output complementary behavior. This process has been visualized in figure 4.6.

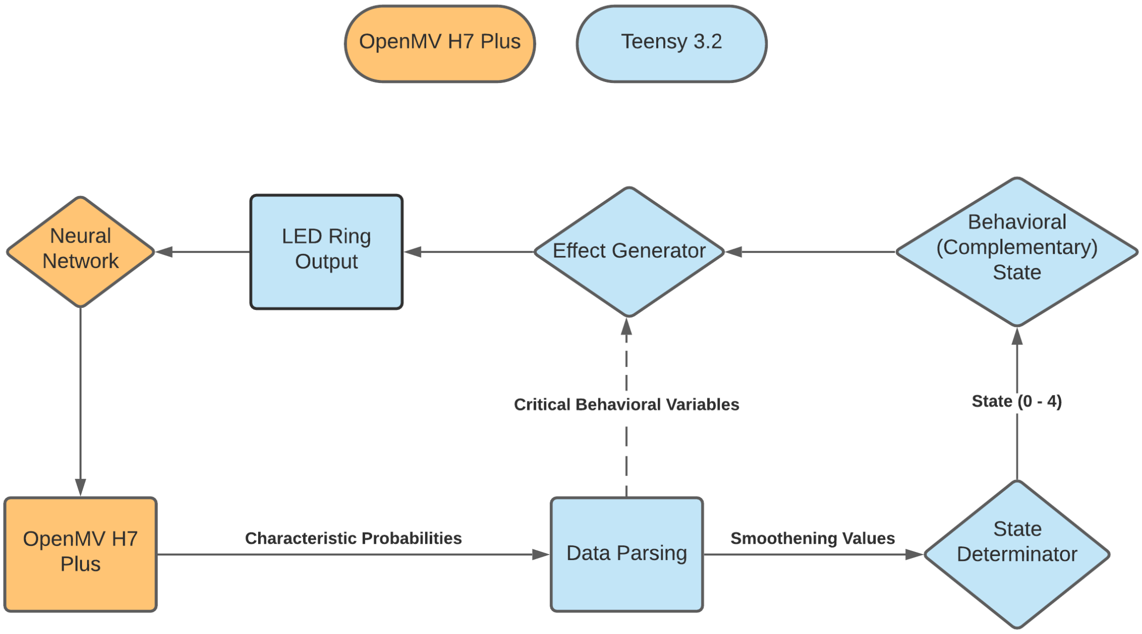

Figure 4.6: Creature Data Flow

Within each state's effect generator, several values can be used to alter the effect generator's behavior and output. These values will be referred to as critical behavioral variables. When considering the startled state, the critical behavioral values can be the intensity of the LEDs and the speed at which the LEDs flicker. Each effect generator will be tailor-made for the state it serves, and the critical behavioral values will be determined along the way.

One key concept to discuss is the utilization and implementation of the characteristic values. Whereas the characteristic values and corresponding states discussed in figure 4.5 will exclusively be used to determine a creature's state, it does not necessarily induce that other characteristic values can not be used to alter the effect generator's output. The random and chaos values determine the startled state, but the startled state's effect generator can also use other characteristics to construct effects.

Component	Amount
Teensy 3.2	2
OpenMV H7 Plus	2
WS2812B 5050 RGB LED Ring	2
5V Power Supply	1
Thingiverse H7 Stand	2
Breadboard	2
Jumper Wire	16
Black Acrylic (35cmx35cm)	1
Black Acrylic (25cmx25cm)	1
Cardboard	1

4.2.3 Materials

4.2.4 Hardware Design

The minimum viable product's hardware design will be more sophisticated and complicated than the previous iterations, primarily since several aspects of previous iterations have to be combined.

The Teensy and the OpenMV H7 will not be powered by USB anymore but will use a custom 5V power supply. From this power supply, all the electronic components will be powered. This includes two Teensies, two OpenMV H7's, and the two 60 LED rings. The power supply is fixed on the inside of creature A, which has an adapter connector to it. The other side of this adapter is fixed on the inside of creature B, allowing the power to easily be disconnected. The power adapter between creature A and B is long enough to separate the two adequately for testing. The distance between the two creatures will be roughly fifty centimeters.

Both creatures have an opening at the side of their case, allowing the power cable, as well as the combined data cable for the OpenMV H7 to the Teensy, to fit through. The OpenMV H7 will be placed next to the creature it belongs to, using a shielded cable that includes the Vin, ground, Rx, and Tx required to send data to the Teensy over UART.

In figure 4.7, the creatures that are used for the minimum viable product are shown. Creature A has been used for prototyping in the previous iterations, whereas creature B has been created solely for the minimum viable product. Creature A has been altered by implementing the power adapter, to which creature B can connect to power its electronics.

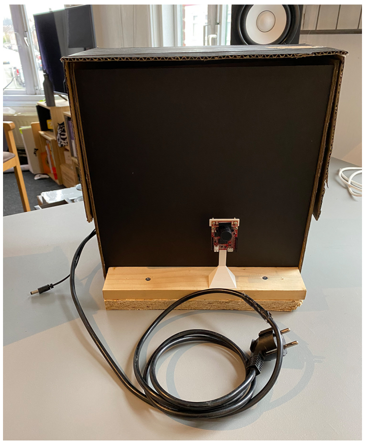

(a) Creature A

(b) Creature B

Figure 4.7: Creatures A and B

(a) Creature A Side Setup

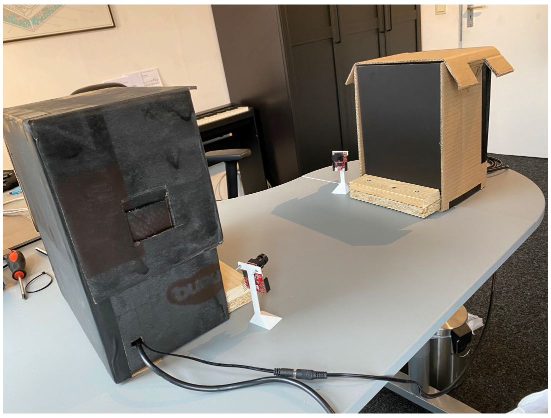

(b) Creature B Side Setup

Figure 4.8: Minimum Viable Product Setup

In figure 5.6, the setup for testing is shown. The distance between the creatures and the power connection between them closely resembles the final product. The camera integration and connection are not yet on that level. This will be further discussed in the testing and evaluation section of the minimum viable product.

4.2.5 Software Design

Teensy Code

The software design combines previous iterations of the prototypes with the primary adjustment in the creature's behavior. The OpenMV H7 - Teensy communication over UART from the previous iteration is implemented in this iteration. Also, the trained deep learning model from the previous iteration has been utilized in this iteration.

The Teensy code's data flow has been altered to mimic a more realistic and organic approach to behavior. The OpenMV probabilities are smoothened in the Teensy code. Within the Teensy code, multiple states have been defined within a switch case construction. Each case, which represents a different state, calls the state's respective function. Case 0 is resting, in which the resting function is called. Within the resting function, an effect generator is responsible for producing the lighting output.

A global, smoothened variable activates each state. This variable is a product of the difference between the probability it receives from the OpenMV H7 and the current probability for that particular characteristic. The thresholds within the smoothening function can be adjusted to, for example, adjust the time it takes to switch from one state to the other. When choosing the states, there are conditions to determine which state to enter. These conditions are subject to change, which will need finalizing in the prototype's fine-tuning and testing.

The code for controlling the creatures can be found on Github [1]. This applies to the TensorFlow code, the OpenMV code and the Teensy code.

OpenMV Code

The OpenMV H7 Plus operates on a C/C++ basis allowing the hardware to be way quicker than conventional programming languages as Python. However, the cam's top-level can be accessed with MicroPython scripts while utilizing the full I/O. The script for Symbiosis is therefore written in Python.

The first part of the code imports the necessary modules and the trained neural network as a TensorFlow Lite (tflite) model with its respective classifier labels. An infinite loop is created in the scripts, which runs as fast as the hardware allows the loop to run, allowing the script to pass as many values to the Teensy as it physically can. The tflite model's output contains probabilities linked to the model's labels, which are all the characteristics of the creature's light (chaotic, smooth, organic, symmetrical, graphical, and random). Each probability is assigned to a variable within the script, which is written as a string over UART to the Teensy.

TensorFlow Code

The network that the OpenMV H7 Plus uses has been created with the help of Edge Impulse. Edge Impulse is an application that specializes in machine learning deployment on microcon-

[1] https://github.com/wouterachterberg/symbiosis_gp

trollers. OpenMV and Edge Impulse work together closely, allowing the implementation of the software to be smooth.

Within the Edge Impulse application, all the relevant data has been imported and labeled. The data has been split into training data and test data. The training pool is used to train the model, whereas the training assesses the trained model's accuracy. The training pool consists of 2,461 images, and the test pool consists of 539 images, giving an estimated split of 80/20.

All the images are resized using Lanczos' downsampling algorithm from 240x240 to 96x96. The features are then extracted from the 96x96 RGB images. This results in the feature explorer of figure 4.9. The input for training the neural network layer will be the 27,648 features created in the previous step.

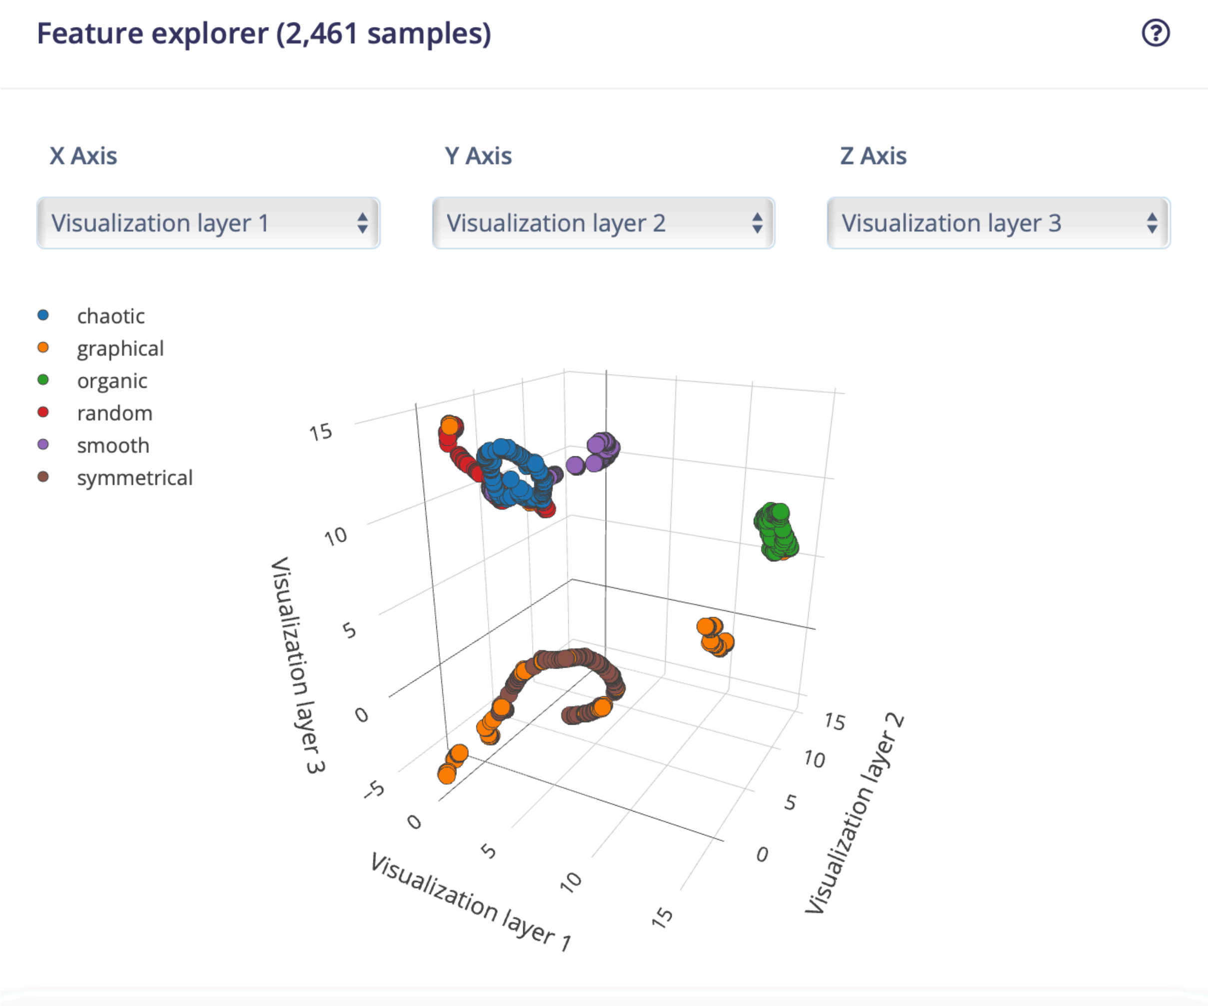

Figure 4.9: Feature Explorer Neural Network

The final step is to train the neural network. Ten epochs are responsible for training the neural network, with a learning rate of 0.0005 and a minimum confidence rating of 0.60. The images are three-dimensional inputs since they are 96x96, with each pixel have three color channels, hence RGB. This results in 96x96x3 per pixel.

For the training of the model, the widely using 2D convolutional layer will be used. This 2D layer will have a moving window, referred to as a filter, that mover over each image three times, once for each color channel. The filter scans the entire image and calculates the dot product between the image's pixel value and the filter weights. After the first 2D convolutional layer, the features from the filter are downsampled. This decreases the dimensions of the features, making them more efficient and lighter in use. The combination of a 2D convolutional layer followed by a pooling layer is repeated with fewer filters to focus on more specific features of the images.

The output of the 2D / pooling layer combinations is flattened. This means that spatial, multi-dimensional outputs are reduced to a simple vector. After flattening the layer, a dropout function is added to prevent overfitting of the model. The dropout shuts down random neurons at each training cycle to decrease the dependency on a single feature within the network. Once an image passes this neural network, it will be assigned one of the six outputs (chaotic, smooth, organic, symmetrical, graphical, and random).

	CHAOTIC	GRAPHICAL	ORGANIC	RANDOM	SMOOTH	SYMMETRICAL
CHAOTIC	100%	0%	0%	0%	0%	0%
GRAPHICAL	0%	96.7%	0%	0%	0%	3.3%
ORGANIC	0%	0%	100%	0%	0%	0%
RANDOM	1.1%	0%	0%	98.9%	0%	0%
SMOOTH	0%	0%	0%	0%	100%	0%
SYMMETRICAL	0%	9.9%	0%	0%	0%	90.1%
F1 SCORE	0.99	0.95	1.00	0.99	1.00	0.93

Figure 4.10: Confusion Matrix Neural Network

The trained model has been evaluated based on the test data. The confusion matrix in figure 4.10 shows how the model performs. On this exact data, the models perform very well.

5 TESTING

The project's testing phase will take the minimum viable product and conduct tests to validate the prototype. The tests will be split into two parts; the performance part and the external conditions part. Since the minimum viable product is a part of the final installation, most tests will cover the technical implementation.

5.1 Performance

The performance of the minimum viable product will be tested on accuracy and speed, primarily. The speed en accuracy of the neural network and the final creature's algorithm will be tested.

5.1.1 Neural Network Classifier

Speed

The most critical aspect of evaluating the trained neural network's speed is by considering its inferencing time. The inferencing time indicates how long it takes for the network to make a prediction based on input data. It is essentially the travel time from the input layer of the network to the network's output layer. As the complexity and density of a neural network increases, the inferencing time does so too. Primarily for real-life applications, inferencing time is essential. The lower this time, the quicker the application can react. In figure 5.1, the calculated inference time for the trained neural network can be found.

The inference time for the network is at 665ms, which is a little more than half a second. For real-time applications, the goal is to keep the inferencing time under 1000ms. The inferencing time is heavily influenced by the clock speed of the device the network is running on. The higher the clock speed of the device, the lower the inferencing time will be.

Also, the peak RAM usage and overall ROM usage for the network have been calculated. These come in at 362,5K, and 87.8K, respectively. These numbers are well within the boundaries of the OpenMV H7 Plus' capabilities.

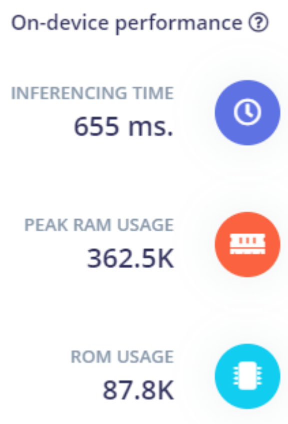

Figure 5.1: Neural Network Classifier Performance

Within the OpenMV environment, the neural network's probabilities are sent to the Teensy over UART at roughly 5Hz. 5Hz is more than enough for the current application, as the creature does not need to change states that fast.

Accuracy

The accuracy of the model can be measured at two stages. It can be measured after the training data and after the test data. The accuracy after the training data will likely be higher since the model used the training data to train itself and has seen the data before. In figure 5.2, one can see the confusion matrix for the training data.


Last training performance

ACCURACY 97.8%

LOSS 0.08

CLASSES 6

Confusion matrix

	CHAOTIC	GRAPHICAL	ORGANIC	RANDOM	SMOOTH	SYMMETRICAL
CHAOTIC	100%	0%	0%	0%	0%	0%
GRAPHICAL	0%	96.7%	0%	0%	0%	3.3%
ORGANIC	0%	0%	100%	0%	0%	0%
RANDOM	1.1%	0%	0%	98.9%	0%	0%
SMOOTH	0%	0%	0%	0%	100%	0%
SYMMETRICAL	0%	9.9%	0%	0%	0%	90.1%
F1 SCORE	0.99	0.95	1.00	0.99	1.00	0.93

Figure 5.2: Neural Network Training Data Accuracy

The accuracy of the model is 97.8%. The inaccuracy of the model is primarily in the graphical and symmetrical output. This inaccuracy indicates that the graphical and symmetrical data's differences were hard to pick up using the model's input features. In figure 5.3, one can find the accuracy of the model after being tested on the data. This means the model has never seen this data before.

The accuracy of the model on the test data is 95.73%. As expected, this accuracy is lower

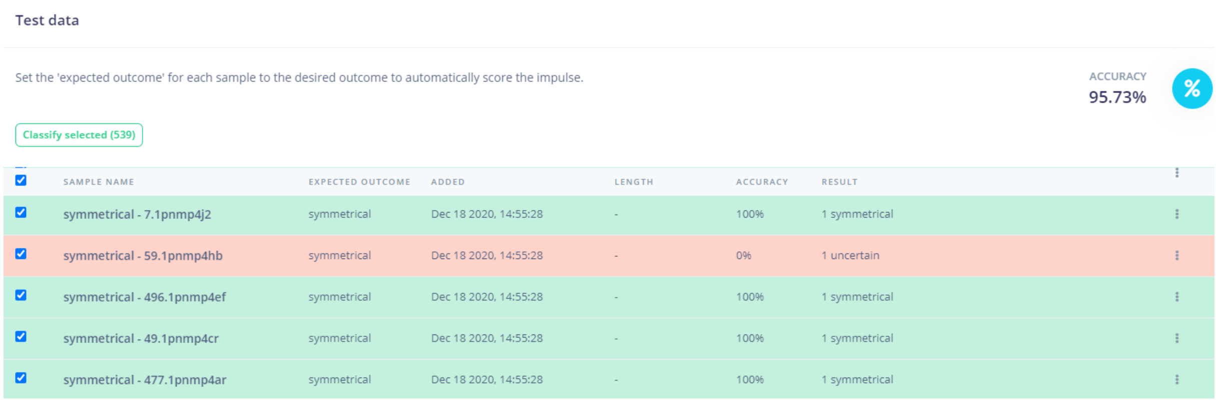

Figure 5.3: Neural Network Test Data Accuracy

than the training data accuracy but is still high. On the test data, the inaccuracies come from the graphical/symmetrical mix-up.

5.1.2 Creature

Speed

The speed of the creature can be measured in several ways. The most obvious way is to determine how quickly the creature responds to the probabilities it gets from the camera. When the creature is in state A, how long will it take before it gets to state B? Another way of measuring the speed of the creature is the frames per second the LEDs run on. Is the creature capable of running the LEDs at a high rate while parsing and interpreting the probabilities of the OpenMV H7?

For the testing of the state recognition speed, a response timeline graph is created. This graph visualizes how long it takes before the creature switches to a different state when interrupted by a human hand. It can be found in figure 5.4. The time has been recorded with a timer on a phone between the first image taken and the last image taken. In 5.5 the probabilities from the OpenMV H7 Plus have been added. The third values is the organic probability, whereas the value is the chaotic probability.

When analyzing the time response graph, the creature starts with regular, smooth behavior. In t1, when a hand is disturbing the creature's view, the output starts to get more chaotic since the network's probabilities are not trained to recognize any (human) disturbance. At t4 the creature is in a different state. At t7, when the hand is removed, the creature gets back to the smooth state in roughly 0.5 seconds, with t8 being the state from t0 again.

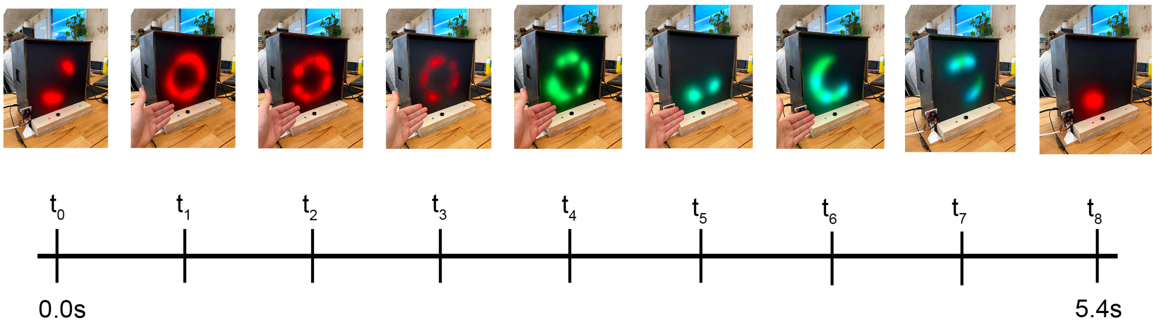

Figure 5.4: Switching States Time Response Graph

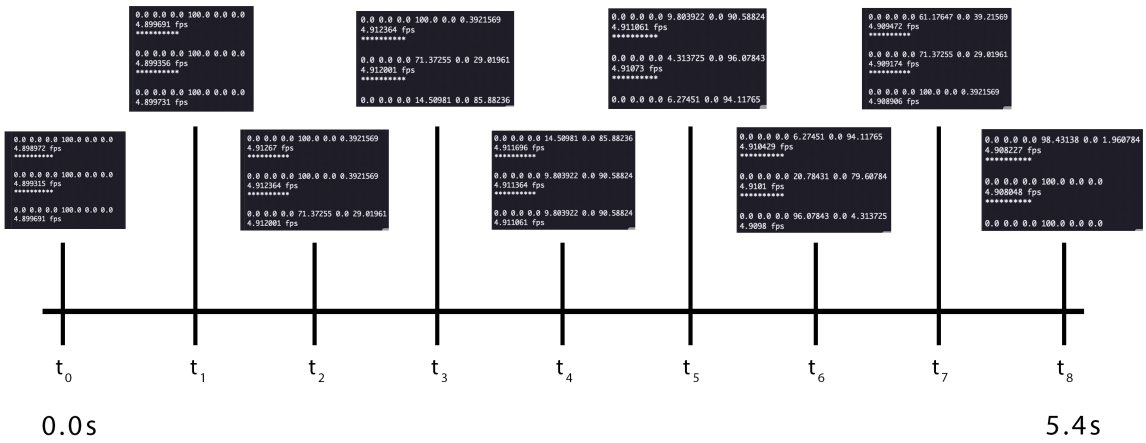

Figure 5.5: Characteristic Probabilities Values over Time

5.2 External Conditions

The creature's neural network is at its core. It determines how, when, and why it will detect specific characteristics of the other creature. Whereas the previous section showed us how the network performed well on training data and test data, the real-life application will be different. The test and training data for the creatures were generated under perfect conditions. In real-life, these conditions are not perfect anymore, and this section dissects how the creature performs influenced by external conditions.

5.2.1 Lighting

The creature's training data is different from the live data feed of the creature the majority of the time. In figure 5.6a, one can see a sample of the training data. The creature is perfectly centered, and there is no visible background. The blacks are relatively black, and the colored

LEDs are clear. These conditions change when testing. In figure 5.6b, a screenshot of the live data feed from the OpenMV H7 can be found.

(a) Organic State Training Sample

(b) Sample Live Feed OpenMV H7 Plus

Figure 5.6: Organic Sample & Live Feed Sample

The blacks in figure 5.6b are way less defined and tend to become more white/grey. This causes the contrast between the acrylic plate and LEDs to decrease. As a result, the neural network can have difficulties with label the output correctly. The primary issue found was that there were a lot of false positives for the chaotic state. In figure 5.7, a sample of the chaotic characteristic training data can be found. The contrast between the acrylic and LEDs is less than in figure 5.6a. The increase of external light causes the contrast between the acryl and LEDs to decrease. It explains why, with increased external lighting, more false positives of the chaotic state occurred.

Figure 5.7: Chaotic State Training Sample

5.2.2 Distance

The creatures use a camera as their sight. The neural network is trained on images that were shot from a fixed distance to the creature. Since all the training data is equal to the images found in figure 5.8, it is crucial to obtain such images as live input for the creatures. The closer the real-time feed comes to the training data, the more accurate the neural network will perform. When increasing the distance between the creatures, the frame of the camera sensors will change gradually.

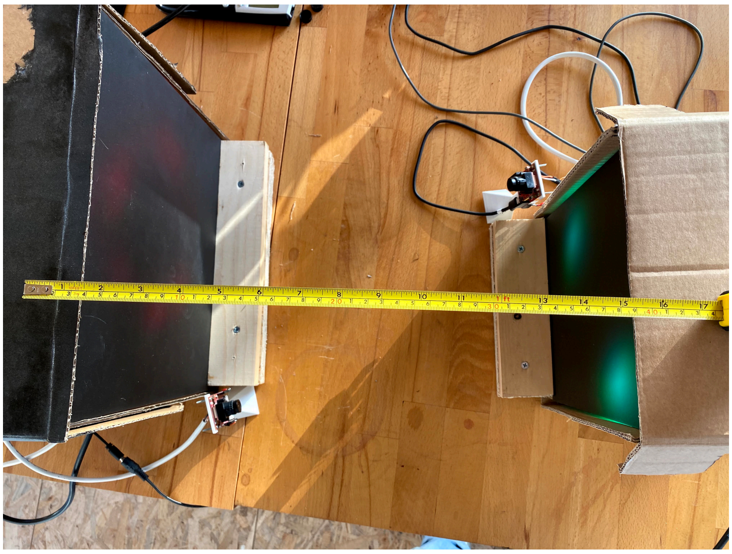

Figure 5.8: Creature Threshold Distance

In figure x, the distance at which the creatures could still identify each other state without too much label flickering is measured. This distance was roughly 40 centimeters. Closer was not an issue since the creature's acrylic takes up most of the camera sensor. When the creature's cameras exceeded 40 centimeters distance, the creature's acrylic was not the frame's primary focus, and the predicted output labels started flickering. In figure 5.9, the frame at the threshold distance of 40 centimeters can be found.

5.2.3 Noise

The prototype required lots of data connection within the creature self and a single power connection between the two creatures. Within each creature, the power adapter supplied the Vin and Ground. The Teensy 3.2, the OpenMV H7 Plus, and the LED ring were powered from this power supply. Creature B had all the relevant parts soldered onto a breadboard, whereas creature A worked with a breadboard and jumper wires.

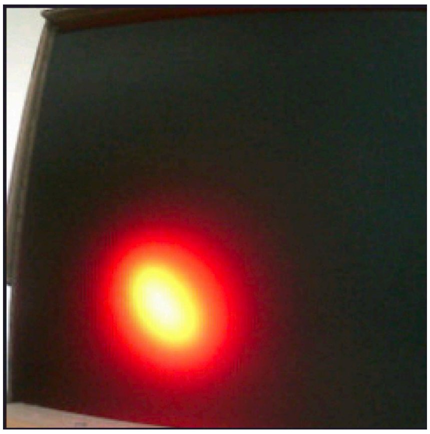

Figure 5.9: Creature Threshold Frame

The primary source of noise was within creature A, where all the components were connected with jumper wires. The most noticeable noise source was between the Teensy and the OpenMV H7 Plus UART connection. When the UART connection is not stable, random bytes will appear on the receiving side (Teensy), making the incoming data useless. Although the noise data was filtered out due to a specific start and end marker sent by the OpenMV cam, the sent data frequency decreases.

Creature B was not as prone to noise as Creature A, likely due to the soldering of the components.

6 EVALUATION

This section focuses on the evaluation of the project. The requirements will be evaluated first, after which a stakeholder evaluation will follow with Lumus Instruments. This section concludes with extrapolation to the final, envisioned installation.

6.1 Requirements Evaluation

The requirements that have been set together with Lumus Instruments in section 2.6 will be evaluated in this section. Each requirement will be discussed briefly. It will start with the conceptual requirements, after which the physical and budget requirements follow.

6.1.1 Conceptual

The creatures should communicate using actuators and sensors only.

Each creature has an OpenMV H7 Plus as sensor and WS2812B 5050 RGB LED Ring as actuator. This is the only communication method between the creatures since there are no other data connections between creatures.

The creatures express themselves using dynamic LED lighting.

Each creature expresses itself using the WS2812B 5050 RGB LED Ring. The exact expression is controlled by the algorithm on the Teensy and is based on live data and chaos theory, making the expression responsive and dynamic.

Using sensor input, the creatures should 'be aware' of neighboring creatures and characteristics, and decide whether to respond with complementary behavior, resulting in inter-agent interaction.

The creature is aware of the dynamic behavior of neighboring creatures and responds accordingly with lighting. This reaction starts a chain of endless interaction between two creatures, constantly obtaining new neighboring creatures' probabilities. The model does not include a cascading classifier, but this will be further discussed in section x.

6.1.2 Physical

A power grid is the only thing connecting the creatures. No physical data lines should be used in between creatures.

The only connection between the two creatures is a power line. So there is no data line connecting the creatures.

6.1.3 Budget

A single creature should be priced in the range of 100 600 EUR. 100 EUR for creatures that are strong in numbers / swarm behaviour. 600 EUR for more complex creatures in terms of learning capabilities.

A single creature consists of an OpenMV H7 Plus, which costs 80 EUR, and the Teensy 3.2 costs 20 EUR. All the other materials are subject to change, along with the costs of it. However, these additional costs will likely not exceed 50 EUR. The total amount for a single creature comes down to roughly 150 EUR. This total amount purely consists of shelf components. Lumus Instruments should still take into account additional costs as labor and custom parts as PCBs and the power grid.

6.2 Performance Evaluation

The overall performance of the creature is as expected. The neural network classifier is capable of calculating all the relevant probabilities five times per second. The neural network's inferencing time comes in at 655ms, with an accuracy of 95.73% on test data. The inferencing time combined with the accuracy makes it suitable for live-feed applications.

The creature states can be recognized by other creatures quickly, shifting from one state to another state within 0.5 seconds. The switch within 0.5 seconds happens when the characteristic probability is 100 on one state and suddenly switches to a different arbitrary state with a probability of 100.

The neural network is trained under daylight conditions inside a room. The state recognition of the neural network would perform optimally under identical lighting conditions. The further away the lighting conditions deviate from the original lighting conditions, the more difficulties the neural network will have with recognizing creature states. However, these inaccuracies are more drastic under brighter conditions than under darker conditions. This is a result of the decrease in contrast under brighter conditions.

The distance between the creatures at which the neural network performs best is around 40 centimeters. From the previous section on testing, it became apparent that the neural network has trouble recognizing creature states beyond this threshold. Moving the creatures closer, to around 30 centimeters, works fine. The implementation of a cascade classifier could increase the working distance between creatures.

Considering these evaluations, scaling the size of the swarm is achievable with equal hardware. The next section will discuss what is required to upscale the current minimum viable product towards the envisioned final installation.

6.3 Extrapolation to Final Product

Within the report, there has been a clear distinction between the final installation and part of the final installation that has been developed as a minimum viable product. Chapter 5 tests this minimum viable product, and from these results, recommendations for the extrapolation of this prototype will be given. This section will be split into three primary parts; the neural network, the physical appearance, and the creature's algorithm.

6.3.1 Neural Network

The current neural network performs reasonably well. It is capable of calculating probabilities for six output labels. However, there are some limitations to this neural network that need to be improved when developing the installation.

One of the most critical changes to the neural network would be to include a cascading classifier. Such a cascading classifier is used widely in computer vision applications. A cascading classifier scans the frame and detects the object it is trained for. When a dedicated cascading classifier for a Symbiosis creature is trained, it can determine if a creature is in the frame. However, it can also detect at which coordinates in the frame the creature is.

The benefit of using a cascading classifier is that the creature can determine another creature's location in its frame on exact coordinates, even though there are several other objects in the frame. This allows the creature to be more precise regarding contextual awareness. The creature can decide whether to respond to another creature or not and intelligently adapt its behavior accordingly. The creatures of the minimum viable product are not capable of determining whether there is another creature in the frame but only determine the probabilities of the given frame, whether there is a creature or not.

When the cascading classifier is used parallel to the current neural network, the accuracy of detecting the characteristics of neighboring creatures will increase. Once a creature detects another creature, it will focus on that part of the frame only to utilize the characteristic detecting neural network. This will exclude objects in the frame that are not relevant to determine characteristic probabilities.

With or without implementing a cascading classifier, the current neural network would need to be retrained. The network focuses on single frames of a live feed to determine probabilities. When the creatures' final states are determined, the neural network needs to be trained on the respective characteristics. This retraining of the network needs to focus on two aspects, specifically: depth and diversity. Depth refers to the expansion of training data. Each characteristic is currently based on 500 images, which should be increased to allow the neural network to extract more exact features in the images. Diversity refers to the incorporation of more diverse images. Currently, only images from a single point of view are included. When retraining the neural network, several distances from the creature should be included. Also, different angles from the creature should be included. However, this also relies on the potential implementation of the cascading classifier. If the cascading classifier is not implemented, more different angles will be needed for the neural network to perform well.

6.3.2 Physical Appearance

The physical appearance of the creatures will change. The shape of the creatures will be organic, entirely from translucent, black material. Acrylic is a possibility but will be hard to shape organically since bending and curving are difficult without a mold. Also, the creatures would be in more significant numbers, somewhere around 150. The creatures' lighting output will need to visible from 360 degrees and not only when facing it directly. Also, the creatures need stands to be able to position them at different heights.

The internals of the creature would also change physically. Designing a printed circuit board for each creature will be very useful. This ensures stable connections between the creature's hardware components and allows more comfortable construction, deconstruction, and potential debugging. The camera placement would change and be integrated into the creature instead of next to the creature. A possibility is to cut a subtle hole in the board through which the camera lens fits precisely. Lumus Instruments could also place it directly on top, for more straightforward implementation, designing the creature more animal-like with eyes. However, an essential aspect to consider is that the camera will probably need to rotate. With the first implementation, the creature's whole body would need to rotate, whereas the second implementation also allows the camera to rotate solely.

6.3.3 Creature Algorithm

The algorithm at the heart of the creature responsible for the incoming data and lighting output will need changes, too. The first change would be to design more refined behavior and allow the creature to crossfade between states. The more refined behavior allows more complex interactions, and perhaps more understandable behavior for human. The crossfading between states is more soothing to watch and smoothes the creature's behavior instead of jumping between states. Lumus Instruments can also experiment with color-behavior relations. Each behavior can correspond to a specific behavior, making the creature more sophisticated. A final addition to the behavior would be to investigate the use of critical behavioral variables within the effect generator. This will make the effect generator more dynamic by making it dependent on characteristics semi-related to the state.

6.3.4 Artist Impression

This section covers an artistic impression of Symbiosis. The hexagonal shape of the creature is based on the MAKERFACTORY MF-6324861. This is a hexagonal neopixel-panel containing 37 individually adressable LEDs. Figure x shows a rough sketch of the layout. The acrylic glass of the minimum viable product can be cut into six rectangular side panels and two hexagonal front and back panels.

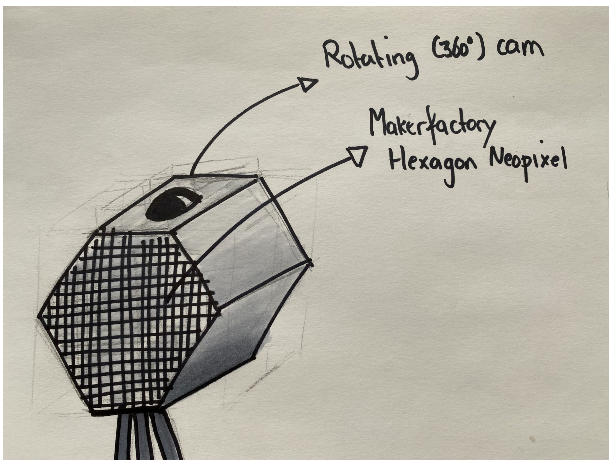

Figure 6.1: Rough Sketch Symbiosis

Within these panels, two MAKERFACTORY neopixel-panels are used for the back and front of the creature. The other six sides can be covered in rectangular neopixel-panels as the NeoPixel 8x8 grid. The side panels can also be covered with individual LEDs or left empty.

On top of the creature, a dome with the OpenMV H7 Plus cam on the Pan & Tilt shield will be installed. This allows the cam to look around and search for other creatures. When retraining the neural network for this specific application, the recognition time will be within 1.0 second and the accuracy will be 95% or higher. This works for usage in live scenarios. The render in figure 6.2, tripods have been included and the creatures have been set at different heights.

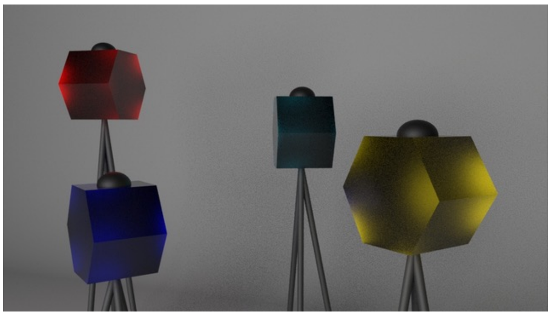

(a) Creature Group Render Side View

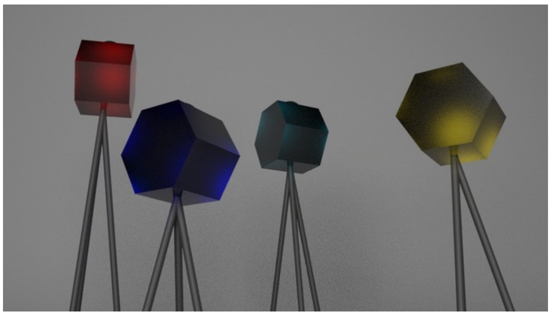

(b) Creature Group Render Bottom View

Figure 6.2: Symbiosis Group Render

7 CONCLUSION

7.1 Research Conclusions

After ten weeks of implementation, testing, and evaluation, a clear direction has been set for the Symbiosis project of Lumus Instruments. It was vital for their artistic swarm installation to understand how to design swarm dynamics with decentralized inter-agent interaction. The research question has translated to the minimum viable product, which resembled decentralized inter-agent interaction.

That question was tackled with the help of several sub-questions, which have been answered with literature research. The literature researched provided essential information for answering the sub-questions. The first sub-question is focused on designing contextual awareness for a single creature in the swarm. This has been achieved by implementing a smart camera and deploying a neural network on the camera, which has been trained to recognize neighboring creatures' lighting characteristics.

The engagement and enticement in artistic installations were covered in the project's next section. The novelty effect, participatory threshold, and installation location significantly influence artistic installations' engagement and enticement. The novelty effect has been tackled by designing ever-changing dynamic behavior, dependent on the characteristic probabilities. In contrast, the participatory threshold is naturally low, with visitors of the installation passively interacting with the installation when walking past it. Also, given the project's origin, the installation will be displayed at appropriate locations, where attention for the installation will be focused and extensive.

The third sub-question focuses on encalming the technology and installation. Several aspects of calm technology have been implemented in the minimum viable product. The installation requires the least amount of attention and is capable of generating appealing content without human intervention. The installation's primary focus is on the periphery and does not need to be actively engaged with to interact. The installation also does not speak but communicates to the outside world using dynamic lighting behavior.

The minimum viable product incorporates the decentralized inter-agent interaction. The minimum viable product has been tested on technicalities. The results from the tests, such as the neural network performance, speed, accuracy, and influence of external conditions, have led to numerical insights. These insights have shaped the recommendations for extrapolation to the envisioned end product.

7.2 Future Work

Section 6.3 already covers recommended adjustments to the minimum viable product to get closer to the envisioned end product. The first recommendation for future work would be to implement the suggestions made in section 6.3. Extensive testing of these new components would be the next logical step. After extensive testing, new insights will be retrieved.

There are also other thoughts to consider when further developing the swarm installation. The durability of the minimum viable product is minimal. The parts that have been used will not last long, especially in outdoor situations. Improving the overall build quality would be wise, including the use of PCBs internally.

The interaction between humans and the entire swarm is also an essential topic to explore. This project focused on the swarm's contextual awareness and inter-agent interaction, but humans' involvement is critical for the swarm, too. When extrapolating to a prototype with over one hundred creatures allows for rigid testing with human subjects' involvement.

The inter-agent communication and recognition also benefit from additional testing in the final setup. Furthermore, how the inter-agent communication would translate to more sophisticated swarm dynamics. The current inter-agent interaction focuses on the interaction between two different creatures, but this could vastly change when a creature can simultaneously identify multiple creatures. This would open up the possibilities to advanced swarm dynamics and allow for multi-agent interaction.

The artist's impression shows a possibility regarding the implementation, but the final installation also needs to explore its limitations and boundaries. This primarily focuses on the recognition boundaries and behavior in a finalized setup.

REFERENCES

[1] Radhakisan Baheti and Helen Gill. "Cyber-physical systems". In: *The impact of control technology* 12.1 (2011), pp. 161–166.

[2] Harry Brignull and Yvonne Rogers. "Enticing People to Interact with Large Public Displays in Public Spaces." In: Jan. 2003.

[3] Judee Burgoon et al. "Interactivity in human-computer interaction: A study of credibility, understanding, and influence". In: *Computers in Human Behavior* 16 (Nov. 2000), pp. 553–574. DOI: `10.1016/S0747-5632(00)00029-7`.

[4] Amber Case. *Calm technology: principles and patterns for non-intrusive design*. " O'Reilly Media, Inc.", 2015.

[5] Gregory Dorais and David Kortenkamp. "Designing Human-Centered Autonomous Agents". In: Aug. 2000, pp. 321–324. DOI: `10.1007/3-540-45408-X_32`.

[6] Nicholas Epley, Adam Waytz, and John Cacioppo. "On Seeing Human: A Three-Factor Theory of Anthropomorphism". In: *Psychological review* 114 (Nov. 2007), pp. 864–86. DOI: `10.1037/0033-295X.114.4.864`.

[7] Seblewongel Esseynew et al. "Context-aware Technologies and Systems". In: June 2016, pp. 120–132. DOI: `10.20544/AIIT2016.16`.

[8] Raphael Féraud and Fabrice Clérot. "A methodology to explain neural network classification". In: *Neural networks* 15.2 (2002), pp. 237–246.

[9] Jacques Ferber and Gerhard Weiss. *Multi-agent systems: an introduction to distributed artificial intelligence*. Vol. 1. Addison-Wesley Reading, 1999.

[10] Vincent François-Lavet et al. "An introduction to deep reinforcement learning". In: *arXiv preprint arXiv:1811.12560* (2018).

[11] Michael Goodrich, S. Kerman, and S.-Y Jung. "On leadership and influence in human-swarm interaction". In: *AAAI Fall Symposium - Technical Report* (Jan. 2012), pp. 32–37.

[12] J.J Jiang, Gary Klein, and Richard Vedder. “Persuasive expert systems: The influence of confidence and discrepancy”. In: *Computers in Human Behavior* 16 (Mar. 2000), pp. 99–109. DOI: 10.1016/S0747-5632(00)00014-5.

[13] Michael Koch et al. “The novelty effect in large display deployments–Experiences and lessons-learned for evaluating prototypes”. In: *Proceedings of 16th European Conference on Computer-Supported Cooperative Work-Exploratory Papers*. European Society for Socially Embedded Technologies (EUSSET). 2018.

[14] A. Kolling et al. “Human Interaction With Robot Swarms: A Survey”. In: *IEEE Transactions on Human-Machine Systems* 46.1 (2016), pp. 9–26. DOI: 10.1109/THMS.2015.2480801.

[15] Hannu Kukka et al. “What makes you click: Exploring visual signals to entice interaction on public displays”. In: Apr. 2013, pp. 1699–1708. ISBN: 9781450318990. DOI: 10.1145/2466110.2466225.

[16] Greg MP O’Hare and Nicholas R Jennings. *Foundations of distributed artificial intelligence*. Vol. 9. John Wiley & Sons, 1996.

[17] Reza Olfati-Saber. “Flocking for Multi-Agent Dynamic Systems: Algorithms and Theory”. In: *Automatic Control, IEEE Transactions on* 51 (Apr. 2006), pp. 401–420. DOI: 10.1109/TAC.2005.864190.

[18] Gonzalo Parra et al. “Quantifying the interaction stages of a public display campaign in the wild”. In: Oct. 2014. DOI: 10.1145/2639189.2639216.

[19] Jenny Preece et al. *Human-Computer Interaction*. GBR: Addison-Wesley Longman Ltd., 1994. ISBN: 0201627698.

[20] Thomas Sheridan, W. Verplank, and T. Brooks. “Human and Computer Control of Undersea Teleoperators”. In: (Dec. 1978).

[21] Lee Sproull et al. “When the Interface Is a Face”. In: *Human-computer Interaction* 11 (June 1996), pp. 97–124. DOI: 10.1207/s15327051hci1102_1.

[22] James Szalma. “Individual differences in human-technology interaction: Incorporating variation in human characteristics into human factors and ergonomics research and design”. In: *Theoretical Issues in Ergonomics Science* 10 (Sept. 2009). DOI: 10.1080/14639220902893613.

[23] James Szalma. “On the Application of Motivation Theory to Human Factors/Ergonomics: Motivational Design Principles for Human–Technology Interaction”. In: *Human Factors* (2014), pp. 1453–1471. DOI: 10.1177/0018720814553471.

[24] Mark Weiser and John Seely Brown. “Designing calm technology”. In: *PowerGrid Journal* 1.1 (1996), pp. 75–85.

[25] Mark Weiser and John Seely Brown. “The Coming Age of Calm Technology”. In: (1996).

[26] Niels Wouters et al. “Uncovering the Honeypot Effect: How Audiences Engage with Public Interactive Systems”. In: June 2016. DOI: 10.1145/2901790.2901796.

A INITIAL BRAINSTORM MIND-MAP

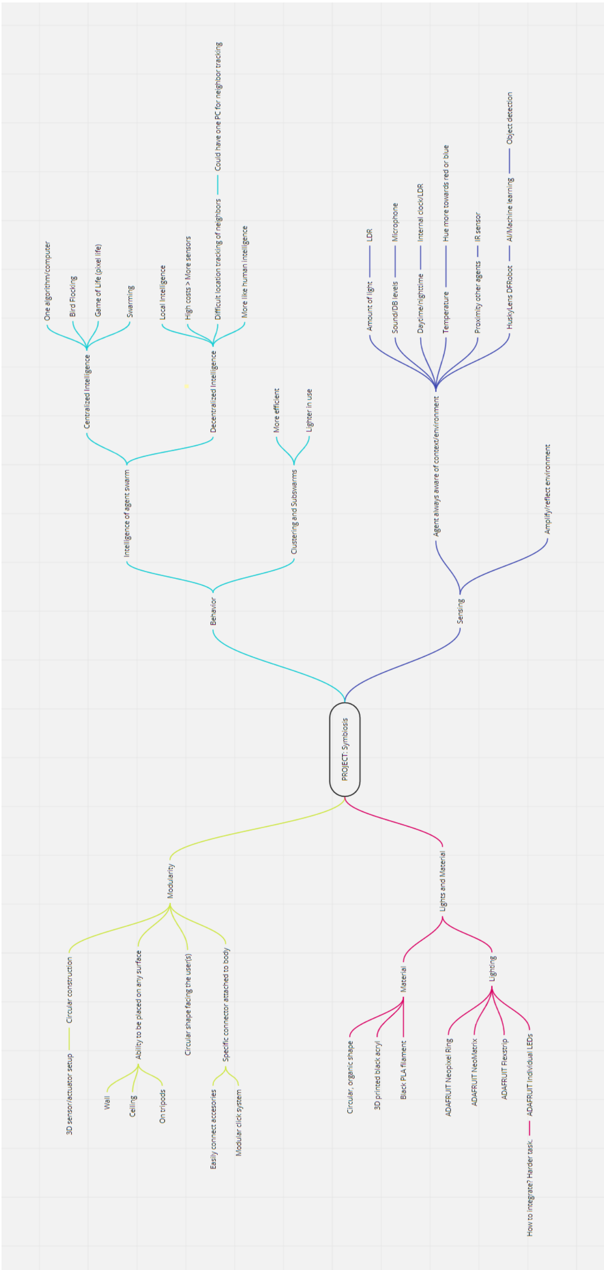

Figure A.1: Initial Mindmap

B NXT: SHIFTING PROXIMITIES REFLECTION

Shifting Proximities is a new media art exhibition at the NXT museum. Upon the recommendation of Lumus Instruments, I visited Shifting Proximities. Shifting Proximities consists of new media installations, focusing on the human experience and interaction in the face of social and technological change. I want to cover seven installations and explain my thoughts on that installation in this brief reflection.

The exhibition starts with Connected, by Roelof Knol. The idea behind the installation is how you share your personal space with strangers in the same area. A simple camera, something like a Kinect, I reckon, is tracking people from above. An algorithm then creates shapes connecting the different people in the room. The geometries change as you move through space, breaking and chaining you to other people. I enjoyed the concept since it was aesthetically pleasing. There was room for interpretation, and you were free to do what you wanted. The tracking of the shapes worked pretty well, which is extremely important with such installations. If the technology does not perform its 'simple' task, the message it conveys becomes useless.

The next installation is Topologies #1, by United Visual Artists. It a long hall filled with ambient smoke, and several planes of light are being projected. I did not understand what the installations wanted the audience to think about, so I started to gaze into the hall and watch the planes elegantly move through space. The whole show took around twelve minutes, in which most of the movement was relatively slow. It would have been more interesting if the artists played more with speed to create a more intense experience. I also believe that sound design plays a vital role in installations, which I did not experience with Topologies #1.

Habitat by Heleen Blanken was an interesting experience. When I entered the room, a couple of people watched the large screen on which spacy visuals were displayed. There was flat water in front of the visuals, which beautifully reflected the ground's visuals, which I appreciated. The visuals on the screen were like you were flying through a mountain and caves landscape. There was also a rock in front of the screen, and it was not before the security showed us there were IR sensors in the rock that we found out you could interact with the visuals. I would consider that down-side of the interaction because it should be more apparent to the audience.

The interaction itself was very subtle, in which the user could move his/her hand closer to the IR sensor and control the movement through the visuals. I enjoyed the abstract visuals but think that the interaction could have been incorporated better in the installation.

Distortion in Space by Laser Feast was one of the more exciting installations of the exhibition. It was quite an extensive experience. It started with clear visuals and an explanation of black holes. After these visuals, we entered a room covered with glass from the inside. One wall was a screen and gave chaotic, abstract animations about the black hole. There was once again a Kinect-like sensor that allowed subtle interactions. These interactions were delightful because it was a little uncertain of what action caused a reaction. This uncertainty made the interaction mysterious and exciting. One of the significant aspects of this installation was the sound. These were sound effects that suited the visuals very well, which enhances the experience significantly.

Econtinuum by Thijs Biersteker was beautifully, organically designed. Its purpose was to bring the visitors back to the roots of the ecosystem. It looked almost like a form of Symbiosis. The roots of two trees were interconnected, and LEDs were implemented in these roots. The LEDs lit up in sequence to mimic a feeding/sharing mechanism. There were also projections on the walls, which were rather abstract. The important part is how they stated that artificial intelligence was incorporated into the installation. There were sensors located in the room that were able to capture specific gas concentrations and temperature. These data inputs controlled the visuals in the room. I wonder how the AI component is incorporated because I could not immediately grasp how it was implemented. I did enjoy the subtle interaction the installations had with the audience, even though we could not influence the interaction. I once again am convinced that more substantial audio effects or sound gives such an installation a boost. Although the interaction was subtle, there could also have been a more active component. Most visitors were intrigued by the interconnected roots but lost interest quickly because the novelty was gone.

Biometric Mirror by Lucy Mcrae and Niels Wouters was an artificial intelligence installation that would make assumptions about your physical features. It was witty to see how machines would estimate human characteristics as beauty and age, but it was not very interesting. This type of artificial intelligence is well known, and computer vision has been around for quite some time. The personal interaction was pleasant since you could get a one to one interaction with the installations. It made it more personal, which was enjoyable. Other than that, there was no real novelty to this installation.

Dimensional Sampling by Yuxi Cao (James) was an enormous screen accompanied by loud music. It focuses on the way QR codes and scanning has become part of daily life. The installation was a repeating audiovisual show, making it repetitive. On the other hand, the music was very enjoyable and allowed you to blend in with the visual and connect with it more in-depth. Other than that, there was no interaction with the installation, making it very universal. I could understand how this would be less intriguing than other, more personal installations.