Exploring the TAME Framework: Insights from AI, Biomedical Engineering and Robotics

Bachelor's Project Thesis
Jiske de Vries
Supervisors: Marleen Schippers & Fred Keijzer
Word count: 15.550

Abstract: This study examines Levin's TAME framework to determine if it provides useful guidelines for understanding and developing systems in AI, biomedical engineering, and robotics. Experts from these fields were interviewed to explore how the framework applies to artificial systems. Levin argues that the divide between artificial and natural systems is outdated, suggesting hybrid systems can bridge these two worlds. He emphasises goal-directed behaviour as a key component of cognitive systems. He proposes to design systems with agents at multiple levels that can scale up to form a larger agent.

1. INTRODUCTION

Contrary to popular belief, worms cut in half do not turn into two new worms. Some species might grow a new tail, but generally, you end up with two dead worms. However, our current technology makes it possible to go beyond this natural limit. Scientists can manipulate worms to grow two heads or no head at all (Emmons-Bell et al., 2019). But there are more examples of entirely new body types. Computers can design organisms to perform tasks using simulations. After the simulation, one design wins and is assembled in real life with biological tissues. The resulting life form is called a Xenobot. Xenobots can walk, swim, push pellets, carry payloads, and work together in a swarm. They can survive for weeks without food and heal their wounds. Intriguingly, they defy categorisation, being neither a traditional robot nor a known species of animal (Kriegman et al., 2020).

If we assume that cognition is embodied, we will have to deal with new types of minds. Embodied cognition means that your body shapes and constrains your cognition. So, when scientists create new bodies, they create new minds. According to Michael Levin, one of the inventors of the Xenobot, philosophy of mind has to catch up with the new technological developments because it is too focused on full-grown adult brains. At the same time, these developments can teach us more about the relationship between body and mind (Levin, 2022). For these reasons, Levin wrote the paper *Technological Approach to Mind Everywhere: An Experimentally-Grounded Framework for Understanding Diverse Bodies and Minds (2022)*, wherein he developed the philosophical framework Technological Approach to Mind Everywhere, or TAME. The framework adopts a practical, constructive engineering perspective to map how to predict and control new types of minds. Why is it called TAME? Levin holds that ‘engineering is not just making technology, it is bringing an approach, one that is fundamentally about truth over inherent biases and limitations, the task of recognizing and ethically relating to beings around us’ (Frequently Asked Questions, n.d.). His approach is technological: building new minds or modifying existing ones. Besides, he ties his perspective on the mind to basal cognition, an approach to cognition that starts from the simplest and smallest organisms that exist and scale up from there (Lyon et al., 2021). Levin uses the words cognition and mind not only for humans and other animals but also for cells or plants. This explains the last part, ‘minds everywhere’; agents can exist in different spaces and scales. The cells, tissues and organs in your body can all be considered agents, which entails that your body is built up of many different agents that can work together or compete. This raises the question how smaller agents scale up (or down) to a larger whole that can walk, eat and talk. Levin proposes that bioelectrical patterns, the electrical signals that our cells and tissues use to communicate and function, underlie scaling up and down in agents. Going further into the bioelectrical patterns is beyond the scope of this thesis, but further information can be found on Levin’s own website.

The TAME framework works very well for biological agents, but how does it work for artificial agents? Levin states that the framework facilitates experimental approaches for detecting, understanding, and functionally interacting with both natural and artificial intelligence. In the paper *Living Things Are Not (20th Century) Machines: Updating Mechanism Metaphors in Light of the Modern Science of Behavior*, he even writes that the stark conceptual distinctions between artificial and life are no longer viable or productive (Bongard & Levin, 2021). This argument rests strongly on hybrid systems, which have both living and artificial parts, and so erase these boundaries. Levin believes this leads to a ‘powerful unification’ (2022, p.3) and sharing tools across disciplines and systems. This brings me to the following research question: ‘Does Levin’s TAME framework provide sound guidelines for future developments and understanding in AI, robotics and biomedical engineering design?’ In order to answer the research question, I interviewed AI researchers, roboticists and biomedical engineers about their predictions and systems. This way, I can make the framework more concrete and determine how the concepts in TAME are perceived. Additionally, I used secondary literature to clarify TAME’s philosophical principles, specifically functionalism and embodiment. This study is relevant to scientific literature because the framework has not been applied yet to current systems in AI, robotics and engineering. Applying the framework to these systems can provide practical feedback and insights into the strengths and limitations of the framework.

2. THEORETICAL BACKGROUND

What is TAME all about, then? In this chapter, I provide an overview of TAME by outlining the key features of his framework. Then, I explain the possible tension in his framework. Before discussing these features, I have to make a critical remark. The terms agency, cognition, mind, intelligence and self all have different meanings, especially in the philosophy of mind. Levin acknowledges this issue and holds that the set should be taken together, rather than focusing on the precise demarcation of each term (Table 1). The framework is still in its infancy, so the concepts are meant to be useful rather than uniquely correct. Since his main focus is cognition, I will focus on this concept, which Levin loosely defined as 'all the activities undertaken by a self' (2022, p.42).

Table 1: Levin's terminology

Term	**Levin's description**
Agency	Set of properties related to decision-making and adaptive action
Cognition	The functional computations between perception and action
Consciousness	The first person phenomenal experience of any self.
Intelligence	The functional ability to solve problems in various spaces.
Mind	Dynamic aspect of self
Self	A coherent system emerging within a set of integrated parts.

2.1 Key features of TAME

Levin's TAME framework includes four key features. He argues that goal-directed behaviour is the key invariant for comparing different complex systems. These systems can be placed on a persuadability scale, which shows how they can best be manipulated. Levin recognizes a multi-competency architecture in biological agents and suggests this as a design principle. He is committed to gradualism, which he uses to support his other ideas.

2.1.1 Goal-directed behaviour

TAME takes goal-directed behaviour as the critical component of all cognitive systems. This implies we make no distinction between systems based on their material substrate, which ranges from artificial to natural materials, or their origin, which can be designed or evolved (Seifert et al., 2024). Taking goal-directed behaviour as the key invariant rests on a functionalist perspective on cognition and 'dethroning evolution' (Levin, 2022, p.8). A functionalist perspective on cognition entails that you define cognitive processes by their function rather than their internal constitution. So, when systems show goal-directed behaviour, this is enough for cognition. Thus, imagine that you and a zombie go to get chocolate ice cream; there is no distinction between your 'true' preference and the zombie's preference for chocolate. Even though the zombie has no brains (its internal constitution), you show the same goal-directed behaviour. Next, Levin argues that evolution is nothing magical, so we should not withhold cognition from machines just because they are designed. He argues that we can mimic evolutionary processes in the lab, an example of this is the Xenobot. This creature is designed with an evolutionary algorithm. The algorithm combines all kinds of biological tissues to perform a function and then selects the best one. Besides, he describes evolution as a hill-climbing search algorithm that results in 'selection among random tweaks', and we do not know the consequences of these small tweaks in the long term. Levin argues that 'if this short-sighted process can give rise to true minds … then so can a rational engineer approach' (2022, p.8) How sophisticated the goals of a system are depends on multiple factors: how much memory and forward planning does the system have, and how well does it adapt to obstacles and different starting points?

2.1.2 Persuadability scale

TAME focuses on how we can predict and control complex systems. Which practical and conceptual tools can we use to control and predict systems? Levin believes these tools should be shared among different disciplines, so he uses a scale to map systems based on the way they can be persuaded to change their behaviour. Hence the name of this scale: the persuadability scale. This scale is continuous but has four clear intermediate steps at the moment: hardware modification only, modification of the setpoint, training by punishments and rewards, and communication through reason. The scale is inspired by Dennet's intentional stance, which works with three different levels of abstraction. The idea behind both scales is that you choose the correct scale of abstraction for each system. We can make more precise predictions at a low abstraction level. However, we can ignore irrelevant details and make broader, more generalised predictions when we shift to a more abstract perspective. For example, when you throw a ball, you can precisely predict the ball's behaviour. However, determining whether your dog will chase the ball is more accessible based on his ‘beliefs’ than calculating his body mass and amount of neurons. The difference between Dennett's intentional stance and Levin's persuadability scale is that Levin's scale provides more specific guidelines for using engineering techniques to predict and control behaviour. Hence, we switch from ‘physical stance’ to ‘hardware modification only’ in Levin. Levin argues that we should map these systems based on experiments, not on preconceptions about the systems. For example, to determine the place of a monkey, we first have to do experiments and try out different methods to predict and control its behaviour. Our estimate of where we place a monkey on this scale also depends on us, the designers of the experiments. Are we smart enough to see how smart animals are? If we position the monkey at the punishment and reward step, this means this tool is currently the best way for us to control the monkey. At the same time, this result also depends on our limitations and capabilities in designing and interpreting experiments with monkeys. This scale builds on the idea that goal-directed behaviour is the key invariant across different cognitive systems because we map all types of agents on the same scale. He makes no distinctions based on the origin story, designed or evolved, or the material, organic or inorganic. So, artificial and biological systems are placed on the same scale.

2.1.3 Multiscale competency architecture

TAME argues that cognitive agents are made up of smaller agents. These agents also have their own goals, for which they cooperate or compete with each other on different levels. For example, in a school of fish, fish cooperate and compete for food simultaneously. Here, the school of fish and the individual fish are considered agents that solve problems. For example, in a school of fish, each fish helps avoid predators by copying its neighbours' movements. This creates a confusing, swirling motion that makes it hard for predators to target any single fish. By sticking close and moving together, they protect the whole group. This shows the multi-scale competency architecture of the school of fish. Now, we can also apply this architecture to our body, where smaller agents, such as cells, compete and cooperate for nutrients with each other. At the same time, the cells scale up to one larger body and work together towards a common goal, such as food or shelter. How agents scale up or scale down can be explained by bioelectricity. When cells connect via gap junctions, they scale up, which has several results. They get more information, they have more power to deal with this information, their memories melt together, and they can not fool each other anymore (Levin, 2022). The cells can still achieve their own goals, but they are also part of a bigger whole with its own goals. This implies that ‘other minds’ not only exist outside but also inside us; cells compete for nutrients in the body but also work together to gather food. This is an embodied cognition claim because the body plays a fundamental role in cognition. I will explain this point further in the next section, but for now, it shows us that not only your brain solves problems; other parts of your body help, too. Next to bioelectricity, stress is also a component of scaling up and down. When something goes wrong, you see a predator, your essay is terrible, or your leg hurts, you feel stressed. This spreads throughout your body and causes the rest of your body to feel stressed as well. Each part starts to work to reduce its stress, which results in working towards the same goal. The type of goals agents are concerned with, or in psychological terms, stress about, can tell you how sophisticated the agent is (and how sophisticated you are in observing).

The type of goals agents are concerned with, or in psychological terms, stress about, can tell you how sophisticated the agent is (and how sophisticated you are in observing). Is he or she concerned with global warming (very sophisticated) or getting enough ice cream? The sophistication of the system's goals depends on how much memory and planning the system has and how much information the system can digest.

2.1.4 Gradualism

TAME is committed to gradualism: evolutionary changes are slow and happen in small steps. This principle also applies to cognition; it would be weird if your parents had no cognition and you could suddenly do puzzles. Levin uses two gradualist arguments to argue for a fully non-binary approach to cognition in different areas. The arguments imply that we have no excellent criteria to distinguish between systems with and without cognition. His first argument is that the journey from 'just physics' to cognition is continuous because we start with just a few cells and grow to full adults. His second argument is that we cannot make a distinction between the true cognition in biological systems and the fake cognition in machines. He believes that in hybrid systems, where a living system works together with artificial parts, there is no clear demarcation between the 'true' and 'fake' cognition, because the parts work together as one (2022).

2.1.5 Unification

These features show how Levin makes functionalist and embodied claims about the nature of cognition. On the one hand, goal-directed behaviour is the critical component, not the material part. On the other hand, multiscale competency or scaling up homeostatic functions is only seen in biological agents (Bongard & Levin, 2021). This creates tension because functionalism suggests that the material substrate does not matter, while embodied cognition holds that the material substrate constitutes cognition. Before discussing this possible tension further, I explain why this is relevant.

Levin wants to bring together different disciplines to share tools that help predict and control complex systems. For instance, in morphogenesis, tadpoles turn into frogs without any single cell knowing the final shape. Using tools from fields like computer science and psychology can improve our understanding and ability to work with these complex processes. We can use machine learning to make predictions about the process or use behavioural terms to understand the process better. He also mentions this practical purpose of the framework; it is not 'just philosophy' (2022, p.16). Instead, he proposes a cycle between engineering and philosophy where you 'philosophise, engineer, and then turn that crank again and again as you modify both aspects to work together better and facilitate discoveries and a more meaningful experience' (n.d.). Here, engineering means both building something and engineering yourself by changing your perspective or commitments. Merging different approaches, such as philosophy and engineering, or disciplines like biology and artificial intelligence, can help us look at problems with fresh eyes. At the same time, it also takes time to learn how a framework works, what terms mean and which tools you use. Levin emphasises that attributing agency should not depend on philosophical preconceptions about origin or material (2022). However, our role as observers means that our (often unspoken) philosophical assumptions shape what we observe. For instance, believing that the human brain is necessary for cognition confines research to human subjects. Understanding and refining the philosophical foundations of TAME improves the framework because this makes it easier to find possible problems and adapt to them.

2.2 Embodiment and functionalism, two peas in one pot?

Levin describes his approach to cognition as both functionalist and embodied. In this section I discuss why Levin thinks they are compatible. To do this, I look at his arguments in favour of his commitments. As mentioned before, the tension between these two commitments lies in the constraints of the body. According to embodied cognition, the body shapes and constrains your cognition. At the same time, functionalism claims cognition depends not on the internal constitution but only on the function it performs. This means cognitive processes can be abstracted, which makes them more generalizable. So, both theories lead to a different approach to designing systems and measuring their performance. A purely functionalist approach might miss meaningful physical interactions and sensory feedback, making solutions less effective. On the other hand, while more comprehensive, an embodied approach could make systems less adaptable in different environments and overly complex to design. Since there are many strands in both theories, we have to look at Levin's arguments to determine his position.

2.2.1 Functionalism

Starting with functionalism, Levin holds that there is no privileged material substrate for cognition. The only aspect that matters is the function of a cognitive process. So, when you solve a puzzle, the most critical aspect is the result: knowing which puzzle piece to put where. Levin also describes cognition as 'the functional computations between perception and action' (2022, p.2). This implies that cognitive processes can be abstracted; once you have the input, you can make a step-by-step plan to generate an outcome.

He makes three arguments in favour of this claim. First, he argues that we have no criteria to determine that biological systems have 'true' motivation. Maybe you have 'true' motivation to learn for a test, but do we know whether fish or single cells are truly motivated? There is no criteria to determine this, which suggests that we can not distinguish between biological systems that have 'true' motivation for a goal and artificial systems that are faking this motivation (Levin, 2022). Second, he points out that hybrid systems function as one integrated being and the living tissue interacts closely with the artificial parts. The living tissue, which has subjectivity or 'true' understanding bleeds over into the artificial system. This means we cannot draw a bright line between artificial and living systems based on the presence or absence of subjective experience, as the boundaries between them become increasingly blurred (Bongard & Levin, 2021).

Both these arguments support the idea that cognition is non-binary, but they do not prove functionalism. To see why, we look at people with and without beards. These are two clear categories, still there is a continuum of beardedness: clean-shaven, light stubble, heavy stubble, short beard, medium beard, long beard. So, even though there is a continuum between the two extremes, this does not mean there is no significant difference between them. The same goes for a hybrid system. When living systems use or take up artificial components, this makes the demarcation between the two systems more difficult. However, this vagueness does not mean the two categories are not useful or that all the properties of living systems exist to the same degree in artificial systems. However, Levin could reply that at least in the case of hybrid systems there is no bright line between the two. The integration of biological and artificial components in hybrid systems demonstrate that cognitive function can be distributed across different substrates. This rests on the principle of plasticity, which we investigate in the next section.

Finally, Levin has another argument in favour of functionalism: basal cognition. This approach starts from the simplest and smallest organisms and points out that these organisms show behaviours that suggest a basic form of cognition, such as sensing, perception, memory and decision-making. For example, bacteria can move toward nutrients and away from toxins, which suggests they can sense and perceive (Lyon et al., 2021). This supports the idea that cognition can be realised in different substrates, which is also known as multiple realisability. This does not mean that cognition can be realised in every type of substrate, but it does suggest that we should not exclude cognition from systems that do not have a nervous system, such as bacteria. Moreover, basal cognition has a biological basis, namely the simplest organisms. So, we should not take it as a direct argument for functionalism.

2.2.2 Embodiment

Levin presents TAME as an approach to strongly embodied cognition (2022). The claim that cognition is embodied can be read in two ways (Foglia & Wilson, 2013). On the one hand, the body constraints cognition. For example, humans can only hear sounds within a specific range, and we can not hear the high-pitched calls of bats or the low rumbles of elephants. This limits our ability to understand and respond to these sounds, constraining our cognition. On the other hand, the body also shapes cognition. An example of this is the hand gestures we make when we talk. Hand gestures are a good indicator of vocabulary development in children (Rowe et al., 2008), suggesting they are not just by-products of speech but actively shape cognitive processes. However, in both cases, cognition is dependent on your internal architecture. Now that we understand embodied cognition better, we can look at the strong and weak versions. The strong states that cognition heavily depends on the body. It emphasises that cognition is distributed across the brain, body and environment and we can not separate the mind from the body. The weak version agrees the body and environment influence cognition, but holds we can still understand cognition through the internal mental processes. Levin maintains he is a proponent of strongly embodied cognition, but what is his idea exactly?

Levin's concept of embodied cognition stems from the multi-scale agency architecture, which is evident from the following statement: 'There is no truly monadic, indivisible yet cognitive being: all known minds reside in physical systems composed of components of various complexity and active behaviour' (2022, p.2). This multiscale competency architecture results from evolution and allows organisms to solve problems at different scales. This structure is essential because it is very energy efficient. To illustrate this, we look at the heart at the cellular and organ levels. Heart cells helping the heart: Heart cells' reliable signalling and contractions ensure the heart works as an efficient pump. Heart helping heart cells: The heart maintains a healthy environment for heart cells, providing them with oxygen and nutrients. Cognition arises when agents at different levels work together. This bodily structure allows complex behaviour and shows how our bodily processes shape and constitute cognition. Cells solve problems on their own level, which simplifies the problems for the organism as a whole. At the same time, the compound structure of cognition also implies that cognition adapts to changes in the bodily structure: minds are plastic. Returning to the plasticity principle, cognition can mould to changes in the body, such as artificial replacements of body parts. Here we clearly see a strong version of embodied cognition, since cognition heavily depends on the body and we can not separate mind and body.

Levin believes that artificial systems can have this type of embodiment too. This is because he holds that plasticity is essential for embodiment and that we can build this into robots. Levin expects that future machines will have a multi-scale competency architecture to ensure plasticity. He also refers to current research by Bongard et al. (2006) and Kwiatkowski and Lipson (2019) on robots that learn to understand and adapt to changes in their own bodies, which shows that 'morphological change is occurring alongside mental changes' (Bongard & Levin, 2021, p.6). Moreover, he states that: 'Embodiment is critical for intelligence, but it doesn't mean necessarily embodiment and motion in 3D space: embodiment can occur in many different problem spaces. In the end, the relevant factor is … whether the system itself believes it has a body and a perception-action loop in a space it models.' (n.d.) This indicates a very loose conception of embodiment that is compatible with functionalism. Emphasising the system's belief in having a body and a perception-action loop highlights the functional aspect of embodiment rather than its specific physical realisation. It shows that what matters is how the system functions and interacts with its environment, not the specific physical details.

2.2.3 Hybrid systems

Still, hybrid systems could make a case for functionalism and strongly embodied cognition. This study specifically looks at designing smart implants and prostheses. Once they are implanted or worn by a human, the device and human form a hybrid system. These hybrid systems are real-life examples of known cognitive agents with artificial parts, showing that cognition can be functionally defined and, thus, modelled and embodied simultaneously; the artificial parts are taken up in the homeostatic loop, showing that cognition is distributed across the artificial parts and limited by their capabilities.

This would be very similar to the idea of extended cognition, yet another variation of embodied cognition, which holds that the tools you use to solve problems are part of your cognition. The classic argument in favour of this idea also rests on a functionalist approach to cognition. Imagine two people, Otto and Inga, who both want to go to a museum. They both look up the address for the museum, and Inga remembers it while Otto writes it in his notebook. When Inga goes to the museum, she uses her memory of the address to navigate. Otto grabs his notebook, which functions like a memory for him, so we should take it as an extension of his cognition. Extended cognition is close to embodied cognition, but the same criticism remains. This extended approach also focuses on a material neutral approach to cognition, while embodied cognition focuses on the specific execution or implementation of cognitive processes. To illustrate this point better, I use an argument made by Godfrey Smith against functionalism.

He argues that 'perfection functional duplicates' does not make any sense. A perfect simulation or realisation, but in different hardware, is impossible. Functional similarity is always a matter of degree; even slight differences in the details can matter. For example, slight variations in the timing of brain processes can result in functional differences. These details are often ignored in functionalist discussions. As replacements are done, differences accumulate, and the small changes in functionality also result in changes in experience and behaviour (2023). So, even though Otto uses his notebook the same way Inga uses her memory, this small change in the physical implementation results in small changes in behaviour and experience. Otto might need to grab his notebook a few times and switch between walking and reading.

Levin seems primarily functionalist and envisions that the multi-scale competency architecture can be realised in artificial and natural systems. His argument rests on the smooth integration between artificial and natural parts seen in hybrid systems. So, further investigation into how hybrid systems work is required. These systems can provide insights into balancing functionalist and embodied principles in design and how well the systems integrate with the user. Next to this the framework is still very young and has not been applied yet except by Levin himself. Thus, in trying to answer the research question ‘Does Levin's TAME framework provide sound guidelines for future developments and understanding in AI, robotics and biomedical engineering design?’ I consider three subquestions:

1. Where do current systems fall on the persuadability scale, and what is their cognitive light cone?
2. Do hybrid systems form a bridge between artificial and living systems?
3. If and to what extent is there a multi-scale competency architecture in place?

By addressing these subquestions, I aim to assess both the practical and theoretical contributions of the TAME framework. The first subquestion is practical, focusing on implementing TAME for current systems. The second subquestion examines Levin's claim that all systems can be placed on the same scale, which suggests there is no real divide between artificial and living systems because they work closely together. The final subquestion evaluates the key design principle of TAME: building systems with a multi-scale competency architecture. In the next chapter, I outline the methods to answer the central research question and the three subquestions.

3. METHODOLOGY

3.1 Research design

In trying to answer the research question, this study used a qualitative approach, since this makes mapping and analysing easier. The researchers can explain their systems and ideas in simple terms, while research papers are often more difficult to understand for outsiders. By conducting the interviews I can gather insights about the functioning of hybrid systems, how well they work, challenges in designing them and their interaction with humans. During the interviews with theoretical researchers who were still very early in their project, too early to have actual systems to discuss, I used their insights about challenges in hardware and the divide between natural and artificial systems to reflect on themes in TAME.

3.2 Participant Selection

The participants for this study were selected from two primary locations:

1. Groningen AI Department and Nijmegen AI department: This group consisted of experts in AI, providing foundational background information relevant to the cognitive framework.
2. University Medical Center Groningen (UMCG): This group included biomedical engineers and roboticists whose insights on hybrid systems serve as the primary data for this research.

A non-probability sampling method, specifically convenience sampling, was used due to the ease of access to the experts in these institutions. This practical approach allowed for the efficient collection of relevant data from individuals with expertise and experience in their fields. In total, 12 interviews were conducted, of which three were not used for the report because the research topics did not align with themes in TAME (Table 2).

Table 2: Overview interviews

Interviewee	Field / Discipline	Duration (minutes)
Ajay Kottapalli	Nanoscience & Nanotechnology, Materials Science, Characterization & Testing, Engineering, Electrical & Electronic	22
Andreas Milias Argeitis	Biochemistry & Molecular Biology, Mathematics, Interdisciplinary Applications	20
Bert Kappen	Physics	39
Elisabetta Chicca	Engineering,Neurosciences, Physics	57
Erika Covi's team	Electrical & Electronic Engineering	44
Gloria Araiza Illan	Robotics, Biomedical Engineering, Interdisciplinary Applications	51
Herbert Jaeger	Computer Science: Theory & Methods, Artificial Intelligence, Hardware & Architecture	42
Pim Haselager	Philosophy, Psychology, Artificial Intelligence	35
Raffaella Carloni	Robotics, Electronic Engineering	48

3.3 Data collection

3.3.1 Semi-Structured Interviews

Semi-structured interviews were chosen for this study due to their flexibility. This enables the questions to be adapted as necessary based on the flow of conversation and each respondent's specific expertise. Interviewees' statements were only summarised during the interview to correct possible misinterpretations; otherwise, they were not interrupted to ensure my response did not influence them (Brett & Wheeler, 2021). The language proficiency of the interviewees determined the language of the interviews. If the interviewees spoke Dutch, the interviews were conducted in Dutch. Otherwise, the interviews were conducted in English, a common language for both the interviewer and the interviewee. All interviews were conducted in person in the interviewees' offices. Before each interview, oral permission was obtained from the participants to record the conversation. The interviews were conducted from April until June 2024.

3.3.2 Interview Structure

The interviews can be divided into two categories, reflecting the shift in focus from general AI concepts to specific aspects of hybrid systems. The first set of questions was developed based on the key features in TAME, and the researchers were asked how they perceived these ideas or saw that they were present in their system.The second set of questions was more focused on hybrid systems and also asked specific questions about the interaction with the body. The first set of questions was tested out with a researcher, and based on his feedback, the questions were made more straightforward, e.g., only ask one question at a time and not use ambiguous words.

Category 1: AI Experts (Groningen AI Department):

- Introduction: These questions aimed to establish context and rapport with the interviewee, including questions about their current projects and how they became involved.
- Intelligence: These questions investigated the interviewees' definitions and perspectives on artificial intelligence and natural intelligence.
- TAME: A summary of TAME was given to establish context for the questions; the questions delved into the goals, experiments, and challenges related to determining the intelligence of AI systems and the conceptual understanding of agents within these systems.
- Ending: These questions aimed to gather final thoughts and advice from the interviewees, such as what they felt was the most crucial topic discussed and any tips for the researcher.

Category 2: Biomedical Engineers and Roboticists (UMCG):

- Introduction: These questions aimed to establish context and rapport with the interviewee, including questions about their current projects and how they became involved.
- Design, capabilities, and interaction with the body: This set of questions was more specialised, focusing on the design and functioning of hybrid systems, the materials used, challenges in modelling, and the systems' goals and decision-making capabilities. It also included questions on how these systems interact with the human body, integrating biological and artificial parts and conceptualising cognition within these systems.
- Ending: These questions aimed to reflect on the interview, asking what topics the interviewees found most interesting and if there was anything they felt should have been discussed.

The difference between the two sets of questions reflects the shift from a broader exploration of AI concepts and intelligence in category 1 to a more focused investigation of hybrid systems and their interaction with the human body in category 2. The complete list of interview questions is included in the appendix.

3.3.3 Adaptation of Interview Questions

During the first interviews, it became clear that interviewing people who build systems could give more insight into the topic. Since Levin's framework is also based on an engineering perspective, this suits the framework better. However, at this point I had already scheduled, and partly conducted, all the interviews. I decided to revise my interview question for the upcoming interviews and to conduct two additional interviews with professionals working directly with AI and biomedical systems. Since some researchers mainly did theoretical work, I adjusted the questions to delve deeper into their research relevant to a theme in TAME. I also decided to leave out the summary of TAME because discussing and applying the framework requires more time than is available in the interviews. This meant that the data became more interpretive from my side, but it did make the interview more fluent.

3.3.4 Secondary Sources

In addition to the primary data collected through interviews, secondary data was used to provide a theoretical foundation and context for the research. This secondary data comprised some of Levin's academic papers and articles about philosophy of mind to determine the theoretical aspects of his framework. Besides, in preparation for the interviews, I read the articles of the interviewees. This helped me to get a better understanding of their research and adopt the questions when necessary.

3.4 Data Analysis

The analysis of the data involved several steps. First, I listened to the interview recordings and transcribed them. Next, I scanned the transcriptions and categorised the data into different themes. The main themes were the system's functionality and modelling, mapping, and multi-scale competency architecture. These themes are discussed individually for each system to contextualise the results within the specific workings of each system.

Examining the system's functionality and modelling helped me to understand the system in the first place. Additionally, this also gave an insight into the current limitations of the system and how the artificial and living parts stick and work together. This involved understanding which materials are used and how the device is attached or implanted in the user. Additionally, I looked at how body parts or functions are modelled, giving insight into the balance between functional and embodied principles. I also considered what aspects are not modelled, indicating limitations or gaps in the current understanding and integration with the body.

In the mapping theme, I identified how to change the system's behaviour to map it onto the persuadability scale. I also examined what type of goal the system pursues. This theme aims to map the systems onto TAME to make the framework more concrete. Then, for multi-scale competency architecture, I looked at whether and to what extent one is in place, which is the key design principle suggested in TAME. I looked at different aspects of this principle, such as designing it as an agent, whether the human adapts to the device, and whether stress in the user affects the device, which would suggest they are in the same homeostatic loop. Additional themes emerged during the interviews, especially in discussions with theoretical experts. These themes are described separately from those identified in interviews with practitioners and designers.

4. RESULTS

In total 12 interviews were conducted of which 9 are reported in the results. The three interviews that are left out did not give new insights into TAME, because the research interests of the interviewee were too different from the topics in TAME. In order to answer the research question: ‘Does Levin’s TAME framework provide sound guidelines for future developments and understanding in AI, robotics and biomedical engineering design?’ The chapter is divided into three different sections, hybrid systems, other systems, and non-binary cognition. The first and second chapter describe the findings of the interviews according to the three themes: system’s functionality and modelling, mapping, and multi-scale competency architecture. The first chapter discusses smart implants, which together with their users form hybrid systems. Then, in the second chapter, other types of systems, namely biological cells and a neuromorphic robot, are described. After the first two chapters a summary is given to remind the reader of the key points from the interviews and to highlight the main findings. Finally, the last chapter discusses themes that came up during the interviews with theoretical experts or researchers who worked on specific design or modelling techniques. These themes are different types of intelligences, plasticity in the hardware and formal models. The first two chapters help to answer the sub questions and the last chapter gives additional insights into non-binary cognition.

4.1 Hybrid systems in Groningen

4.1.1 Kottapalli: Tactile sensors

System functionality and modelling

The goal of Kottapalli’s group is to create tactile sensors that mimic the human skin, giving robotic limbs a sense of touch similar to that of human palms. Their sensors measure force and recognize textures, which helps the robotic limb to handle objects carefully. The design is inspired by the various layers of human skin, which are well documented in the literature. The stiffness, softness, and texture of the skin, including fingerprints, help to detect different textures.

The team uses biocompatible materials, specifically ‘polymer films with printed force sensors on them’. Polymer films with printed force sensors are thin layers of flexible materials designed to detect and measure mechanical forces, such as pressure or touch. These films have sensors embedded or printed onto them, allowing them to respond to applied forces. However, there are challenges, such as the inability to sense vibrations, temperature and humidity. Additionally, Kottapalli mentions the lack of integration between different senses, e.g. the combination of vision, touch and memory. Imagine grabbing a cup that lies in soapy water. You see little bubbles and remember that this indicates slipperiness. When you touch it, you can feel just how slippery it is and adjust the amount of force you need to use accordingly. He highlights that integration is an important limitation of artificial systems because ‘this sort of intelligence is not possible in artificial systems.’ Kottapalli’s team develops the electronic materials of the sensor, in order to make sense of the data, they have to work together with people from different disciplines, such as machine learning or neuromorphic computing.

Mapping

The primary goal of the tactile sensors is to provide sensory feedback to the users, which helps them to grab objects. There are two types of sensory feedback: force sensing and muscle activity sensing. In force sensing, the polymers in the sensors respond to pressure by producing a small current or changing their electrical resistance, which provides a sense of force. In force sensing, sensors can use electromyography (EMG) to detect muscle contractions. Motor neurons transmit electrical signals that cause muscles to contract. EMG sensors pick up these signals and use them to control the movement of a prosthetic limb. You can see this as a memory that remains in the amputated limb. Even though the current systems do not make autonomous decisions, there is potential for future developments through interdisciplinary collaboration.

Multi-scale competency architecture

Kottapalli views the human and the prosthetic limb as two separate systems because the tasks performed by the prosthetic are automatic for humans. However, users might perceive the prosthetic as part of their body since it helps them perform daily tasks. He also notes that his view might change depending on how intelligent the system becomes. For now, the system remains ‘a tool that allows an activity to happen’. In the lab environment, the device is seen as a task-doer, but this perception might change once it integrates with the human body. At the moment, the prosthetics work mechanically and lack perception. Kottapalli describes cognition as a kind of awareness, that is the result of ‘hundreds of senses put together’. His sensors are still in development and have not been tested on humans, so he cannot yet comment on how people perceive the sensors or whether stress affects the sensors.

4.1.2 Carloni: Lower limb prosthetics

System functionality and modelling

Carloni’s group is working on lower limb prosthetics that aim to mimic the natural movement of human legs. They use inertial measurement units (IMUs) to detect the acceleration, velocity, and position of the user’s upper leg. Based on these measurements, they determine whether the prosthetic knee should bend or extend. This approach involves creating rules and control architectures that guide the leg’s movements, ensuring it performs tasks such as walking and climbing stairs effectively. The artificial lower-limb is attached to the upper leg of the user, the sensors that read the movements are on the upper part of the prosthetic.

She views trust as the most important aspect of designing a prosthetic device. The prosthetic has to work all the time because one of the first reasons users stop using it is that they fall. Even a simple stick could work as a prosthetic, but this requires a lot of energy. You can solve this problem by adding motors, but you have to ensure that the device keeps working all the time. Carloni highlights modelling is not difficult in the sense that the prosthetic is ‘a mass that moves’. So, you can just use Newton’s laws to model the leg’s movements. However, modelling the environment can be problematic. You do not always know what the underground looks like and this can create disturbances in your model, e.g. outside there might be little holes and bumps. Moreover, the sensor needs to be able to read the movements in the upper leg, when this signal is not good enough, it affects how the leg moves. This can have a negative impact on the user’s trust in the device.

Her team is also looking at new materials, namely polymers, to design new actuators. Even though the current results with polymers are not good enough, because they can not exert a lot of force, this is a promising material for the future. Using polymers might allow the user to sense the environment or forces, which would be a great addition to the prosthesis.

Mapping

The primary goal of the prosthetic systems is to provide intuitive and seamless use for the wearer, reducing cognitive load. Users often report difficulties performing dual tasks, such as walking and talking simultaneously, because they must constantly monitor their prosthetic limb. By improving the prosthetic’s responsiveness and reliability, Carloni’s team aims to make these tasks easier and more natural for users.

There are two levels of performance for the prosthetic system: the mechanical and electrical level, and the control and performance level when the human is involved. Starting with the mechanical and electrical level. The team first evaluates whether the prosthetic allows the user to stand and move effectively. They then assess whether the prosthetic moves easily or if the user needs to use a lot of energy to compensate for the missing muscles. So they examine the symmetry in the user's gait by looking at the length of the steps and the movement of the pelvis. Additionally, they measure the energy cost of walking by monitoring the user’s oxygen intake. After building the hardware, adjustments can be made to the control parameters, such as tuning the prosthetic to be less stiff or more compliant based on user feedback. A critical aspect of this phase is educating the user. A medical physician explains how the prosthesis works, which helps the user understand its functionality and feel more confident in using it.

Multi-scale competency architecture

Carloni explains her team works with a human-centred mechatronic approach, so ‘the system, which is the robot plus the human, is seen as a unique entity.’ The robot and human form one system and they are trying to control the dynamic interaction between the two. She mentions that users adapt their motor skills to the prosthetics, this is a learning process she calls familiarisation. Not every user adapts at the same rate; some adapt fast, while others adapt slowly. This depends on how your central nervous system reacts and individual differences in learning, e.g. young people might adapt quicker.

She mentions that future research lies in having full embodiment, which means the prosthetic feels like a natural part of the body. This depends on improving both the mechanical parts and the control systems. It means designing prosthetics to mimic the features of real muscles and fine-tuning them based on feedback from the user. She hopes that these prosthetics will become smarter in the future, making it easier and more intuitive for users to interact with them. They are looking into creating recurrent neural networks so that the device has memory that it can use to predict the user's intentions. However, this has yet to be tested, and it needs to go through careful ethical discussions and follow strict protocols before this is implemented.

4.1.3 Araiza Illan: Cochlear implants

System functionality and modelling

Cochlear implants consist of two main parts: an internal implant and an external transmitter. The internal implant is an electrode array that is in touch with the auditory nerve and has to be surgically placed into the cochlea. The internal implant and external transmitter are connected via small magnets on both sides of the skin. The transmitter has a microphone that picks up signals, translates them into electrical impulses and sends these to the electrode array.

One major challenge in modelling is the ‘cocktail party effect’, in which the background noise makes it difficult for people with a cochlear implant or hearing aids to focus on specific sounds. For normal hearing people this is an easy task, they filter out the background noise, but the microphone picks up every sound. This problem can be solved by developing machine learning algorithms that detect and adjust the settings of the device based on the environment, e.g. at home or a party. Moreover, she states it is very challenging to model all the things ‘embedded in your voice’, such as an accent or emotion, in just one electrical signal. In fact, many children with cochlear implants struggle with emotion recognition.

The implants are made from biocompatible materials, such as titanium, so they integrate smoothly with the body without causing inflammation or rejection. She thinks the hardware can not be improved; it is already very tiny and soft. Some hardware improvements have been explored, such as increasing the number of electrodes. However, this did not help because the different frequencies interact. Hence, the interaction becomes too complex to model because the resulting wave cannot be mapped to individual frequencies.

Mapping

The main goal of the cochlear implant is to help users hear and process sounds. Each user’s sensitivity to electrical stimulation varies, making it difficult to create a one-size-fits-all solution. Once the device is implemented, you have it for your whole life. During the follow-ups after implementation, users can give feedback, e.g. the sounds are too loud, and a clinician or doctor can fine tune the parameters for the user. When working with children this is a challenge because they do not always comply and clinicians only have limited time. A possible solution to this problem is working with a robot that conducts the test because the children can interact with him.

The device does not get to make any decisions, but this might change in the future, when smart algorithms can automatically adjust to a certain setting. Besides, an essential part for children with cochlear implants is to inform their parents about the positive results. Children with cochlear implants can improve their speech significantly and their speech becomes indistinguishable from that of their normal-hearing pears. This remains a challenge because some parents do not want to be part of the normal-hearing community; there is a strong community feeling among deaf people. Additionally, even with hearing aids or a cochlear implant, it can be challenging for individuals to fully integrate into the normal-hearing community, especially in noisy school settings with lots of shouting and background noise.

Multi-scale competency architecture

Araiza Illan classified the cochlear implants as tools rather than agents. The device and user function together as a system, with the implant providing sensory input that the user learns to interpret. She believes that designing the device as an agent can be ‘creepy' and ‘scary’ for people. Trust is an important aspect in the healthcare domain; she does not think people will trust a device ‘that moves around as you are developing’.

She explains that brain plasticity plays an essential role in successful use of cochlear implants. Younger brains adapt to the device more easily, while adults find it harder to adjust because their brains are less flexible. As a result, children often perform better in hearing tests a few years after implantation compared to adults. However, adults tend to be better in emotional recognition because they have already learned to recognize emotions and can use corresponding visual cues more effectively. It remains difficult for the brain to adapt to these new sounds. We don't know if you'll fully adapt or how fast this process will be. She states this plasticity is very interesting but also that it is ‘like a black box, which is doing stuff and then we do not know what is going on and then it works or it does not work.’ Moreover, describing this new perception is nearly impossible. She mentions there are some recordings to model how you hear with cochlear implants and ‘it sounds horrible … very mechanical … creepy to an extent’. But the problem remains, we do not know whether these recordings are accurate.

She is not sure whether stress in the user affects the workings of the device. However, she explains that stress can have an impact in a very simple way. When children have severe hearing loss, they do not have any stimulation. The overstimulation from the implant can cause stress, ‘like shopping in IKEA’ and some people remove the transmitter. Children have to get used to the novel sensory input, which can be done by increasing the usage of the device over time.

4.2 Other systems

4.2.1 Milias Argeitis: Model cell growth and division

System’s functionality and modelling

Miles Argeitis and his lab combine experimental and computational methods to understand how cells coordinate growth and division. They use optogenetics: a method which uses light to control the specific proteins within cells. By adding light-sensitive elements to these proteins, they can switch them on or off with light, allowing them to manipulate the cell precisely. Then, he uses knowledge from systems and control theory to model the behaviour of the cell on the computer. The goal is to fully model how cells grow and divide, but he is not very optimistic about reaching this goal as the process is very complex. He describes it as a ‘terribly, terribly complicated situation … we do not really have a good grasp on it’ and mentions the different aspects that make it so complex. First, there are many components and interactions that are difficult to represent in the model. Second, the cellular processes involve interactions at multiple layers. E.g. proteins interact with proteins, but also with metabolism, which is at a different level. Third, we still do not know all the details of the working of the cell, making it difficult to abstract.

Mapping

Milias Argeitis and his team manipulate living cells. Hence, they do not give the cells goals. Instead, the cells have their own goals. He gives a real-life example of mycelium, a unicellular organism whose goal is to produce another mycelium. He does not know whether this really is a decision, but they can change their behaviour. He states: ‘Even these single cells have three hundred different ways of sensing what is happening around them chemically mostly, but also mechanically … they understand if they are under mechanical pressure [and] if they are under osmotic pressure’.

There are various strategies to change the behaviour of the cell. First, you can remove or alter genes to see how the cell adapts. However, a disadvantage of this technique is that cells can adapt or compensate when something is removed or altered; the cell ‘sort of masks what you did’. Second, you can use chemicals that can target specific proteins and block their function. However, chemicals hit many targets together, which makes analysis difficult because ‘you do not know exactly what you hit.’ To solve these problems, the team prefers to work with optogenetics, which allows them to see what is happening in the cell in real time.

Multi-scale competency architecture

Milias Argeitis describes cells as ‘super flexible and plastic’ and ‘fully-fledged adaptive systems’ because they can grow in various ways depending on their environment. He expresses some difficulty with terms like ‘agency’ or ‘decision making’, explaining that these are human-defined concepts. Instead, he prefers the term ‘cognition’, which he sees as a more fundamental function. According to Milias Argeitis, cognition involves gathering information about the environment or internal state and responding to it, which cells do constantly. He notes that it is possible to ‘fool’ the cell by introducing abnormal conditions and ‘basically confuse them in a way’. He compares this to humans, who, when placed in a situation without any visual stimuli or crazy illusions, would also become disoriented. This comparison emphasises the cells’ inherent ability to adapt to their surroundings, much like humans adapt to their environments. He believes that cognition is widespread, but does not occur below the level of the cell. He uses viruses as an example, these are just ‘passive parasites … they do not have sensing … [or] decision making, all they do is hitch a ride inside the cell.’

Milias Argeitis states he would never assign cognition to a model because the model is ‘just describing’. He does not think it is possible to build something with cognition that is ‘out of my limits and out of my knowledge.’ believes that only living things possess cognition. He reasons that this ability has developed over billions of years of evolution, resulting in incredibly complex molecular structures within living organisms. While he admits he does not fully understand how this complexity translates to cognition, he feels it is ‘qualitatively different from just a big neural network.’ He also explains that while it is possible to mimic evolution by selecting or creating models, it is not the same as real evolution. Real evolution is a highly complex process that operates on many levels. He holds that the abstraction or understanding of evolution as ‘mutating and selecting’ is a greatly simplified version ‘a caricature of what is happening’.

4.2.2 Chicca: Motion flow in robots

System functionality and modelling

Chicca and her team work in neuromorphic engineering; this field focuses on creating physical systems that mimic brain functions to understand how the brain works. There are two main branches in their current research. The first is sensory processing, which tries to understand how animals sense and generate behaviour based on those senses. This research covers senses such as vision, audition, touch and olfaction, but not taste, and uses biologically inspired sensors to study them. The goal is to encode the sensory information similarly to the brain. The second stream focuses on understanding how learning arises in the brain. The team studies how synapses change their strength through learning and how these changes integrate within neural architectures. They use recurrent neural networks with feedback loops to mirror the biological learning processes more closely.

Her current project is vision processing in insects; together with a PhD student, she developed a model of insect behaviour and built a small robot that used this model to navigate. She describes this project in detail. They used event-based sensors, which means that the sensor samples the environment based on changes in the light intensity rather than based on traditional time based sensors. This sampling method ensures that data is not lost when changes happen quickly or that there is much data without any changes. For this project, they designed an architecture with a very clear goal in mind: insect navigation. Insects use motion flow to estimate how far away objects are: objects close by appear to move fast, while objects far away appear to move slowly. However, once you start spinning, the whole field moves, and you have to separate this translation motion from the motion flow. This is too computationally expensive for flies, so they solve it behaviorally: they fly straight, make a turn, and fly straight; this is called saccadic flight. In this way, they do not have to deal with translational motion. The robot uses a model based on insect navigation but does not fly. Instead, it has wheels. This makes the problem simpler because you only have two dimensions to move in. The neural network is designed in a winner-takes-all fashion based on research on the cerebral cortex. This means that neurons excite and inhibit each other, and eventually, one neuron wins. With this architecture the robot looks at the deepest space in front of him and then goes there.

The team uses the platform SpiNNaker, which stands for Spiking Neural Network Architecture, to prototype their ideas before they commit to hardware. The robot has an event based camera that sends input to the SpiNNaker, whose output is used to control the motors. So, currently, software is running on hardware, however, Chicca states she hopes that this will change in the future. She envisions that the algorithm or program is physically implemented in the hardware.

The robot integrates sensory input and motor output, creating a closed-loop system where its actions influence its sensory experiences. This interplay between sensory input and motor output helps the robot to adapt to its environment, much like biological organisms. The physical implementation of the robot is crucial for testing hypotheses and understanding how neural architectures translate into real-world behaviours. However, Chicca mentions that doing these experiments is time-consuming, so they also work with a robot in a simulated environment. She states it is crucial to start the experiments with physical robots before doing simulations; this prevents a 'reality gap'. When people start with simulated robotics, the simulation can work perfectly. However, once they do the real-life experiments, they find out they miss something in the environment, or the robot and 'it is a disaster'.

Mapping

She describes that the robot showed all kinds of behaviours that the team did not program. They only programmed the fundamental behaviour: motion flow. However, just like insects, the robot tends to go to the centre in a tunnel. The parameters in the winner take all structure to determine how far the robot goes into the tunnel. She mentions that other experiments also show that robots explore more when the tunnel is wider and stay closer to the centre when the tunnel is narrow. Another behaviour that popped up was that the robot picks gaps based on the size of the gap. When there is a wall with two gaps in the experiment, the robot picks the biggest gap and when the gaps are comparable in size, he becomes indifferent. She states that this is the beauty of engineering: 'complex behaviours that are observed in insects, and they emerge in our agent, without us doing anything specific.'

She states that to change the robot's behaviour, she has to adjust the parameters in the algorithm. This also happened in the experiment because the robot first moved at a constant speed. The team changed this so the robot adjusts its speeds based on what is happening in the environment, e.g., when an object is close. She admits that the robot does not learn its body size; there is a fixed architecture, which is a possible limitation. On the other hand, insects can fly from the moment they are born, so a fixed architecture makes sense. Next to this, the purpose of the current research was to understand sensory processes; adding the ability to learn would be a different but exciting research.

Multi-scale competency architecture

Chicca's method shows the emergence of complex behaviour from simple rules. She admits her research does not directly address cognition, but the fact that neuromorphic systems show behaviours similar to those of insects suggests an underlying cognitive-like process. Chicca considers the robot as the primary agent, but she states that the neural networks can also be thought of as smaller agents, with individual neurons competing to influence the robot's decision. This multi-scale view of agency reflects the complex interactions between the neurons and the overall behaviour of the robot.

4.3 Results summary

System’s functionality
Kottapalli: The sensors mimic human skin by measuring force and recognizing textures.
Carloni: The prosthetics mimic natural leg movement using sensors.
Araiza Illan: Implants convert sound into electrical signals.
Milias Argeitis: Uses light to control cell proteins.
Chicca: Models insect navigation with event-based sensors and recurrent neural networks.

Challenges in modelling
Kottapalli: The sensors do not detect vibrations, temperature or humidity. Moreover, the senses do not integrate with other senses like vision and memory.
Carloni: Modelling all the different types of terrains remains a challenge.
Araiza Illan: The implants sound very mechanical and do not capture nuances in speech such as accents or emotions. Besides, researchers still do not know much about the interaction between the device and the human.
Milias Argeitis: Biological cells are so complex that it is difficult to model all the different types of interactions in the cell.
Chicca: The robot only has one task and navigates in a simplified environment, on the ground rather than in the air. It is very time-consuming to build a robot, so you have to know that your idea works before you give up on flexibility.

Mapping: Goals
Kottapalli: Provide tactile feedback
Carloni: Intuitive use, not too much energy or cognitive consumption
Araiza Illan: Help users hear

Mapping: Methods to Change System Behavior:
Carloni: Adjust control parameters based on user feedback to improve adaptation.
Araiza Illan: Personalised adjustments by clinicians based on feedback from users.
Milias Argeitis on Cell Growth: Uses optogenetics to manipulate cell behaviour.
Chicca: Adjusts algorithm parameters to change the robot’s behaviour based on environmental inputs.

Multi-scale competency architecture: Build the system as an agent
Kottapalli: The sensors are tools, the human body is much smarter, but this view might differ for the users.
Carloni: One unified system, the system is not an agent.
Araiza Illan: Function as tools working with the user, but not making autonomous decisions.
Milias Argeitis: Adaptive systems, but not decision makers.
Chicca: The robot acts as the main agent, with neural networks functioning like smaller agents influencing behaviour.

Multi-scale competency architecture: Integration with the human
Kottapalli: Provide tactile feedback but lack comprehensive sensory integration.
Carloni: Designed for seamless integration, but no sensory feedback. Users need to trust the device.
Araiza Illan: Fit well within the body, but sound is perceived differently. Users might get overwhelmed by all the sounds.

4.4 Non-binary cognition

How to interact with artificial agents and compare them to natural agents is a widely discussed topic, and each theoretical expert had different ideas on how to approach this. However, none of them believed that artificial and natural agents could currently be put on the same scale. The main themes that emerged were that the current hardware is not flexible enough to support the type of learning seen in biological agents. Instead, artificial systems are often trained for one specific task and do not exhibit this level of flexibility. Also two alternative approaches are discussed: one for dealing with different aspects of intelligence and one for scaling up and down between different levels in biology.

4.4.1 Hardware

The current hardware is not flexible enough to create truly intelligent or cognitive agents. The interviewees still see a clear divide between artificial and natural systems. Biological agents can make new connections, which helps to solve problems and be creative in many situations. For artificial systems to be truly intelligent, they need to go beyond the fixed hardware, which does not allow for forming new connections. Kappen explores quantum mechanics as a possible solution to this problem. Quantum mechanics studies nature at a level smaller than atoms. A key feature of this theory is that you cannot always precisely predict what will happen; instead, you work with probabilities. This approach allows for forming non-local connections, unlike classical physics, which operates more like building a clock where you know what will happen precisely. Elaborate on Jaeger and add quotes.

Covia and her team also aim to create plasticity in the hardware. They work on materials that can change their microscopic properties. These small changes translate into a macroscopic change in the device's electrical properties. The amount of electricity that flows between two neurons changes. This mimics what happens in the brain when a connection between two neurons strengthens or weakens. Another team member mentions that this also means you train the hardware to learn with local information only. This differs from traditional neural networks, where you have global information to update the weights. They emphasise that the co-design of hardware and software is crucial for efficient processing; at the moment, there is a gap between what hardware and software can do. The current software algorithms are very advanced and have many parameters. Realising extensive neural networks with electrical circuits is not feasible since the circuit gets complicated, becomes unstable, and costs a lot of money. For this reason, the team is working on making the two compatible. They try to create software that uses particular hardware features. Consequently, the hardware and software are co-designed; you do not have general-purpose software that you can run on any computer or general-purpose hardware on which you can run any program. Another benefit of this strategy is that integrating the software and hardware results in more efficient processing. Covi gives an example of cooking. Imagine you are cooking a meal in the kitchen, but your recipe book is in the garden. So you constantly have to move back and forth. This is very inefficient. Then, by putting memory and computation in one place, 'you bring the book back to the kitchen'.

4.4.2 Different aspects of intelligence

Instead of unifying different complex systems, Haselager made a distinction between different aspects of intelligence. These aspects are will, smartness or performance, and subjectivity or understanding. He explained this idea with a chess computer. Even though chess computers perform very well during chess games, they do not have a will. You can just put it in 'moron mode', and it gives away all its pieces. In contrast, a baby wants all kinds of things right from the moment it is born. He thinks that attributing will and subjectivity or understanding to such a computer is dangerous. Instead, we have to find a new way to deal with systems that have one, two or all aspects of intelligence. Other interviewees also mentioned that will seems a distinctive aspect of biological systems. Jaeger mentioned he is missing all types of embodiments, which gives rise to a perspective, whereas neural networks miss this; he stated, 'These neural networks do not have plans, they do not have objectives, they have nothing, they are just data in and out.' Covi and her team agreed that their devices always 'get a goal …. That is the essential part, you train it to do something.' Covi also stated she would not assign cognition to her device, or at least not in the sense how humans perceive cognition because 'there is no awareness at all, it is just a mathematical function that gives us an output.'

4.4.3 Formal model to scale up and down

Jaeger aims to unify the different levels of biology. He holds that behind a good theory is a formal model, which gives you a proper understanding of a phenomenon. Unlike computational theory, where you can easily scale up and down between different levels, this is not possible yet in biology. We can not scale up from the cellular level to the behavioural level. This means that specialists on one level, e.g., psychologists, can not talk meaningfully to biochemists. The connections between these levels are not transparent. Instead, we have many different types of models. However, this approach allows us to explain many aspects of physical phenomena. On the contrary, computational theory only uses control switching (so binary code, a 0 or 1). Therefore, Jaeger argues for a reinterpretation of computation: structure a process rather than process a structure. What is the difference between the two? In the traditional interpretation, you have a model shaped after a mathematician's reasoning, with clear, logical steps. In contrast, Jaeger aims to model pre-rational intelligence and focuses on processing vast amounts of sensory data. He aims at working with intelligent materials and exploiting their physical properties. His goal is to create 'a new mathematical language that can serve as, you could say, the translation glue between all of these levels.' He was the only interviewee who had read the paper on TAME and also mentioned this as an explicit criticism of Levin's framework, namely that there was no formal model behind it. 'The physicists, biologists, chemists, sociologists, all those people describe the world. Mathematicians give them the languages in which they can express themselves.' He thinks Levin misses this robust background story. Even though his experiments are solid, he overgeneralizes and is too optimistic about other fields, such as robotics, computer science, and machine learning, which is far less powerful than what he says.'

5. DISCUSSION AND CONCLUSION

To answer the research question, ‘Does TAME provide sound guidelines for understanding and developing AI, biomedical engineering, and robotics?’ I conducted interviews with researchers in these fields. The research question can be broken down into three sub-questions, each addressing a different aspect to provide a comprehensive answer. In each section, the findings are connected to the subquestions:

1. Where do current systems fall on the persuadability scale, and what type of goals do they pursue?
2. Do hybrid systems form a bridge between artificial and living systems?
3. If and to what extent is there a multi-scale competency architecture in place?

5.1 Mapping systems

The human and the smart implant work as one system, but controlling them as a whole happens in different stages. Usually, the persuasion is directed at either the human or the device, not the hybrid system. Changing the behaviour of the implant or prosthetic depends on your research stage. It starts with hardware modifications only and builds up to changing parameters. Once the device is built, it has a fixed structure; the only flexibility lies in the software. From this moment, professionals change the parameters based on experiments or user feedback. Since the software is already trained, changing the parameters is a setpoint modification. Thus, you make small changes to the device settings. Once the device is added or implanted, an essential factor is trust. The user needs to trust the device and the doctor. Trust in the device starts with robustness, meaning the device always works. For example, if a user falls because his prosthetic leg blocks, he might refuse to use it. Trust in the clinician is also essential. Users need to feel like they are in good hands. The doctor explains how the device works or why the device could be helpful for the user. This involves communicating reasons, which is on the far right on the persuadability scale. Since researchers were at different stages in their research, this also influenced their answer. For example, the researchers who work on materials or are in an early stage of their research also purely do hardware inventions.

The goal of the system was described by the researchers in a similar manner, namely, helping the user to walk, hear or feel. The device always gets a goal from its maker, it does not have its ‘own’ goals. Levin might argue that distinguishing between ‘true’ and ‘fake’ goals is not meaningful. However, before the device is added to or implanted by the user, it does have a pre-defined goal.

While predicting and controlling systems are fundamental, trust plays a specific and critical role in how users interact with the device. This is especially important in the medical field, where reliability directly impacts user safety and well-being. Therefore, promoting trust should be considered an essential aspect alongside prediction and control in designing and implementing complex biomedical systems.

5.2 Bridge between artificial and living systems

Currently, researchers still see a clear divide between artificial and natural systems. Theoretical experts clearly stated that classical physics is insufficient to create true intelligence, which requires plasticity and self-directed goals. Moreover, hybrid systems do not perfectly integrate living and artificial parts. Even though the artificial and living parts work closely together, they remain separable.

One reason for this divide is the challenge of perfectly modelling body parts. We cannot capture every detail or integrate all the senses from the original part, like the many sensations in a human palm. Plus, we cannot predict every situation a body part will face, and artificial parts do not learn and adapt over time as natural ones do. However, users adapt to these devices to some degree, but the specifics of this process are unknown. Araiza Illan also described this as a ‘black box’ so it is difficult to improve. Additionally, Carloni noted that users of prosthetic legs experience a heavier cognitive load when they walk and talk simultaneously as if they are performing two distinct tasks. In addition, the human body is very complex, and we have yet to understand it fully. Milias Argeitis is also critical of mimicking evolution or perfectly modelling cells. Even within a single cell, many interactions exist at different levels; the same goes for evolution. This is much more than selecting and randomization.

Finally, current prosthetics lack sensory feedback and act only as mechanical tools, so they do not feel like genuine body parts. Kotapalli is exploring this area and aims to make mechanical devices feel more like natural body parts by adding sensory feedback. This potential is also seen in cochlear implants, which help users process sounds. Children with cochlear implants can adapt to the device and use the sensory feedback to reach the same speaking level as normal-hearing individuals.

The inability to capture every detail, integrate all senses, and predict every situation a body part will encounter highlights the limitations of a functionalist approach. Even though the device becomes part of the user's cognition, this extended cognition still has a different feel. The need for sensory feedback and the importance of users' adaptation point towards an embodied approach that considers the body's interactions and feedback mechanisms as crucial in designing artificial systems. Besides, smart implants are always made of biocompatible materials to prevent rejection by the body. This is an example of more embodied design principles since you look at the specific details of the interaction between the device and the body. Still, a completely embodied approach might be too complex since the body is not entirely understood.

5.3 Multi-scale competency architecture

Currently, there is no multi-scale competency architecture in place for hybrid systems. First of all, researchers working with smart implants do not consider their devices to be agents; they view them as tools. However, other researchers do see their systems as agents. For instance, Milias Argeitis believes that agency begins at the cellular level because an agent needs to be active. Chicca agrees with this perspective, viewing different neurons as individual agents competing to 'win' in their tasks. Designing devices that work more like independent agents, meaning they get to make decisions, might help with intuitiveness. However, it can also be too scary for users and make them uncomfortable using the technology.

Next to this, users primarily adapt to the devices rather than the devices adapting to the users. This means that while users adjust to the prosthetics or implants, the cognition does not bleed over into the device. Whether stress in the user also affects the device could be a good indicator that the user and the device are in the same homeostatic loop, which Levin calls the 'glue of agency.' Thus, whether stress in the user affects the device could indicate how tightly the artificial and living parts work together. However, none of the researchers had performed experiments on this. Only Araiza Illan mentioned that stress can impact device usage. However, straightforwardly, users might stop using the device when they feel overwhelmed. Trust could be a psychological aspect of scaling up, enhancing the smooth interaction between the user and the device. When users trust their devices, it improves how well they work together.

One restraining factor for creating a multi-scale architecture is that the current hardware does not offer the flexibility or plasticity needed. Researchers are working towards creating this plasticity, aiming for seamless hardware and software integration or working with agential materials.

5.4 Limitations and future research

This research focused on research in Groningen and Nijmegen, so the results are only a partial reflection of the application of the TAME framework in the Netherlands. Moreover, I only spoke with one research team member. Other team members might have other opinions and insights about how the framework applies to their work. As a result, the findings might reflect individual viewpoints rather than a collective stance from the research groups. Next to this, researchers were often in different stages of their research, which influenced how much they could answer specific questions. Therefore, future research should use a more homogeneous or larger sample group to ensure consistency and comparable insights. This could entail gathering viewpoints from researchers within the same field or study group, such as focusing solely on cochlear implants and offering a more thorough understanding of specific technologies. Including firsthand accounts from individuals who use biomedical devices or smart implants would provide valuable insights into how the systems feel for the user and their development in working together. This would create more consistent and comparable insights.

Another limitation of this study is that I did not directly discuss the TAME framework with the researchers. The interview time was limited, and the framework was too extensive to discuss in detail with the researchers. Because of this, analysing how TAME fits into their systems is biassed because it is based on my viewpoint. The researchers did not directly address the TAME framework. However, they did offer insights into whether particular features applied to their systems. As a result, the more futuristic aspects, e.g., whether it is possible to use the multiscale competency architecture in future designs of TAME, are not discussed. Future research should explicitly discuss the design principles of the TAME framework with engineers to provide direct feedback on its applicability and effectiveness.

My study was very exploratory and sometimes unfocused, partly because so many different themes and experiments were discussed in the framework. Therefore, future research can build on the study by focusing on a more specific aspect or feature, such as only discussing the MSC or the persuadability scale. They could also address other points of interest not discussed in my interviews. These could be investigating what could serve as the ‘glue’ in artificial systems, helping to improve their integration and functionality. Additionally, experiments to understand how stress affects the user and the device could provide helpful information about the device’s ability to scale up with the user’s goals.

Furthermore, since this was my first time conducting qualitative research, my questions changed as I went along, making it more difficult to compare answers. I also did not always follow up properly, which led to varying levels of detail in the responses. This affected the overall usefulness of the answers for the analysis, as the depth and specificity of the information collected were not always consistent.

5.5 Conclusion

Hybrid systems do not yet form a bridge between artificial and living parts. The potential of a multi-scale competency architecture is promising for future design. However, current hardware lacks the flexibility to support this approach. Additionally, designing devices as independent agents could be intimidating for users.

A significant problem persists in accurately modelling all aspects of a body part. The need for sensory feedback and biocompatible materials underscores the importance of embodied design principles. These elements are crucial for developing systems that integrate seamlessly with the human body, ensuring they feel like real body parts. Even though the device becomes part of the user’s cognition, there are clear differences in how it feels, indicating that the specific physical implementation is important. The persuadability scale could be used to compare different ways of predicting and modelling, but whether all systems should be placed on the same scale is not necessarily a good idea. In conclusion, Levin’s TAME framework provides valuable insights into biological systems but has limitations when applied to current AI, robotics, and biomedical engineering design.

6. REFERENCES

Bickle, J. (2020). Multiple realizability. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy (Summer 2020 Edition).* Retrieved from https://plato.stanford.edu/archives/sum2020/entries/multiple-realizability/

Bongard, J., & Levin, M. (2021). Living things are not (20th century) machines: Updating mechanism metaphors in light of the modern science of machine behaviour. *Frontiers in Ecology and Evolution, 9.* https://doi.org/10.3389/fevo.2021.650726

Bongard, J., Zykov, V., & Lipson, H. (2006). Resilient machines through continuous self-modelling. *Science, 314* (5802), 1118–1121. https://doi.org/10.1126/science.1133687

Brett, B. M., & Wheeler, K. (2021). *How to do qualitative interviewing*. SAGE.

Emmons-Bell, M., Durant, F., Hammelman, J., Bessonov, N., Volpert, V., Morokuma, J., Pinet, K., Adams, D. S., Pietak, A., Lobo, D., & Levin, M. (2015). Gap junctional blockade stochastically induces different species-specific head anatomies in genetically wild-type Girardia dorotocephala flatworms. *International Journal of Molecular Sciences, 16* (11), 27865–27896. https://doi.org/10.3390/ijms161126065

Foglia, L., & Wilson, R. A. (2013). Embodied cognition. *Wiley Interdisciplinary Reviews: Cognitive Science, 4* (3), 319–325. https://doi.org/10.1002/wcs.1226

Godfrey-Smith, P. (2023). Nervous systems, functionalism, and artificial minds [Video]. NYU Mind, Ethics, and Policy Program. https://www.youtube.com/watch?v=4r9Y01H2sEU

Kriegman, S., Blackiston, D., Levin, M., & Bongard, J. (2020). A scalable pipeline for designing reconfigurable organisms. *Proceedings of the National Academy of Sciences of the United States of America, 117* (4), 1853–1859. https://doi.org/10.1073/pnas.1910837117

Kwiatkowski, R., & Lipson, H. (2019). Task-agnostic self-modelling machines. *Science Robotics, 4* (26). https://doi.org/10.1126/scirobotics.aau9354

Levin, M. (n.d.). Frequently asked questions. *Frequently Asked Questions - Latest Version*. Retrieved from https://drmichaellevin.org/resources/ Accessed 25 June 2024.

Levin, M. (2022). Technological approach to mind everywhere: An experimentally-grounded framework for understanding diverse bodies and minds. *Frontiers in Systems Neuroscience, 16,* 768201. https://doi.org/10.3389/fnsys.2022.768201

Lyon, P., Keijzer, F., Arendt, D., & Levin, M. (2021). Reframing cognition: Getting down to biological basics. *Philosophical Transactions of the Royal Society B: Biological Sciences, 376* (1820), 20190750. https://doi.org/10.1098/rstb.2019.0750

Rouleau, N., & Levin, M. (2023). The multiple realizability of sentience in living systems and beyond. *eNeuro, 10* (11). https://doi.org/10.1523/eneuro.0375-23.2023

Rowe, M. L., Özçalışkan, Ş., & Goldin-Meadow, S. (2008). Learning words by hand: Gesture's role in predicting vocabulary development. *First Language, 28* (2), 182–199. https://doi.org/10.1177/0142723707088310

Seifert, G., Sealander, A., Marzen, S., & Levin, M. (2024). From reinforcement learning to agency: Frameworks for understanding basal cognition. *BioSystems, 235,* 105107.

Shapiro, L., & Spaulding, S. (2024). Embodied cognition. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Summer 2024 Edition). Retrieved from https://plato.stanford.edu/archives/sum2024/entries/embodied-cognition/ Accessed on 20 June 2024.

Tracy, S. J. (2019). *Qualitative research methods: Collecting evidence, crafting analysis, communicating impact*. John Wiley & Sons.

APPENDIX

APPENDIX 1: CATEGORY 1 INTERVIEW QUESTIONS

Introductory questions

1. What are you currently working on?
2. How did you become involved in the project?
3. How do you describe ‘intelligence’?
4. How do you describe ‘artificial intelligence’, is this any different for you than ‘natural intelligence’?
5. How does this reflect in your current project?
6. How does your device mimic body part X of cognitive function X?

The framework

My thesis is about the work of Michael Levin, who developed the framework ‘technological approach to mind everywhere’ or TAME in short. With this framework, he tries to analyse and compare the new kinds of intelligence coming up in artificial intelligence and biology and their intersection. Some examples of this are robots made of smaller robots, worms with two heads or a smart knee prosthesis. He is interested in what goals these systems can achieve and how we can persuade them to do something different. He also sees growing and maintaining a body as simple or basal cognition. Even though no cell knows you should grow five fingers, this process often goes right. These changes during the body’s lifetime might also tell us more about how minds come into existence in the physical world.

7. Do you agree that goal-directed behaviour is the primary invariant of all intelligence, which means we do not make distinctions based on origin or material composition?
8. How is this present in your work?
9. What experiments do you do to determine the intelligence of your system?
10. Are there any hindrances in measuring its intelligence?
11. Where would you place your system on this persuadability scale? The persuadability scale entails the kind of interventions (rewiring, setpoint editing, training, logical arguments, etc) that are optimal for prediction and control of your system.
12. Does your system adapt or modify its ‘body’ or ‘cognition’ during its lifetime?
13. Do you consider the different parts in your system ‘agents’ on their own, if we define agent as something that can pursue goals?
14. What kind of goals is your system concerned with? How far in the future are these goals?
15. Do you think that morphogenesis and anatomical homeostasis growing and maintaining your shape) is a type of ‘cognition’?
16. Do you think that the term ‘collective intelligence’ could be a useful term to describe the intelligence in your system? Why or why not?

Ending questions

17. What do you feel is the most important thing we talked about in the interview?
18. Do you have any tips?

APPENDIX 2: CATEGORY 2 INTERVIEW QUESTIONS

Introductory questions
1. Who are you?
2. What are you currently working on?

Design
3. Can you describe how your system works?
4. How do you model this function / body part?
5. What are challenges in modelling or what do you leave out?
6. Do you think we are able to compute the human body, now or in the near future?
7. What material is the system made of?
8. Do you think working with different materials will change the system? In what ways?

Capabilities
9. What goals is your system concerned with?
10. How do you change the behaviour of the system?
11. Does the system make decisions?

Interaction with the body
12. Can you describe how your system interacts with the body? How do the biological and artificial parts interact?
13. Do you consider the human and the medical device one system?
14. Do you consider the device an agent or a tool?
15. How do you describe cognition? Would you ascribe it to your device?
16. Does the cognition of the human change when using the device? If so, in what ways?
17. Does stress in the user affect the device?

Ending questions
18. What question or topic did you like the most?
19. Is there something you think I should have discussed that I did not?

References:

Bickle, John, "Multiple Realizability", The Stanford Encyclopedia of Philosophy (Summer 2020 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/sum2020/entries/multiple-realizability/>.

Blackiston, D. J., & Levin, M. (2013). Ectopic eyes outside the head in Xenopus tadpoles provide sensory data for light-mediated learning. The Journal of experimental biology, 216(Pt 6), 1031–1040. https://doi.org/10.1242/jeb.074963

Brett, B. M., & Wheeler, K. (2021). How to do qualitative interviewing. SAGE.

Bongard, J., & Levin, M. (2021). Living Things are not (20th century) Machines: Updating Mechanism Metaphors in light of the modern science of Machine Behavior. Frontiers in Ecology and Evolution, 9. https://doi.org/10.3389/fevo.2021.650726

Bongard, J., Zykov, V., & Lipson, H. (2006). Resilient machines through continuous self-modeling. Science (New York, N.Y.), 314(5802), 1118–1121. https://doi.org/10.1126/science.1133687

Emmons-Bell, M., Durant, F., Hammelman, J., Bessonov, N., Volpert, V., Morokuma, J., Pinet, K., Adams, D. S., Pietak, A., Lobo, D., & Levin, M. (2015). Gap Junctional Blockade Stochastically Induces Different Species-Specific Head Anatomies in Genetically Wild-Type Girardia dorotocephala Flatworms. International journal of molecular sciences, 16(11), 27865–27896. https://doi.org/10.3390/ijms161126065

Foglia, L., & Wilson, R. A. (2013). Embodied cognition. Wiley Interdisciplinary Reviews. Cognitive Science, 4(3), 319–325. https://doi.org/10.1002/wcs.1226

Kriegman, S., Blackiston, D., Levin, M., & Bongard, J. (2020). A scalable pipeline for designing reconfigurable organisms. Proceedings of the National Academy of Sciences of the United States of America, 117(4), 1853–1859. https://doi.org/10.1073/pnas.1910837117

Kwiatkowski, R., & Lipson, H. (2019). Task-agnostic self-modeling machines. Science Robotics, 4(26). https://doi.org/10.1126/scirobotics.aau9354

Levin, M. (n.d.). *Frequently asked questions*. Frequently Asked Questions - Latest Version. https: drmichaellevin.org/resources/

Levin M. Technological Approach to Mind Everywhere: An Experimentally-Grounded Framework for Understanding Diverse Bodies and Minds. Front Syst Neurosci. 2022 Mar 24; 16:768201. doi: 10.3389/fnsys.2022.768201. PMID: 35401131; PMCID: PMC8988303.

Lyon, P., Keijzer, F., Arendt, D., & Levin, M. (2021). Reframing cognition: getting down to biological basics. Philosophical Transactions - Royal Society. Biological Sciences, 376(1820), 20190750. https://doi.org/10.1098/rstb.2019.0750

P. Godfrey-Smith, “Nervous Systems, Functionalism, and Artificial Minds” [Video]. (2023), NYU Mind, Ethics, and Policy Program. https://www.youtube.com/watch?v=4r9Y01H2sEU

Rouleau, N., and Levin, M. (2023), The Multiple Realizability of Sentience in Living Systems and Beyond, eNeuro, 10(11), doi:10.1523/eneuro.0375-23.2023

Rowe, M. L., Özçalışkan, Ş., & Goldin-Meadow, S. (2008). Learning words by hand: Gesture’s role in predicting vocabulary development. First Language, 28(2), 182–199. https://doi.org/10.1177/0142723707088310

Seifert, G., Sealander, A., Marzen, S., and Levin, M. (2024), From reinforcement learning to agency: Frameworks for understanding basal cognition, *BioSystems*, 235(1): 105107

Shapiro, Lawrence and Shannon Spaulding, "Embodied Cognition", The Stanford Encyclopedia of Philosophy (Summer 2024 Edition), Edward N. Zalta & Uri Nodelman (eds.), forthcoming URL = <https://plato.stanford.edu/archives/sum2024/entries/embodied-cognition/>.

Tracy, S. J. (2019). Qualitative research methods: Collecting Evidence, Crafting Analysis, Communicating Impact. John Wiley & Sons.