DISSERTATIONES
INFORMATICAE
UNIVERSITATIS
TARTUENSIS
22

MADIS VASSER

Testing a Computational Theory of Brain Functioning with Virtual Reality

DISSERTATIONES INFORMATICAE UNIVERSITATIS TARTUENSIS
22

DISSERTATIONES INFORMATICAE UNIVERSITATIS TARTUENSIS

22

MADIS VASSER

Testing a Computational Theory of Brain Functioning with Virtual Reality

UNIVERSITY OF TARTU
Press

Institute of Computer Science, Faculty of Science and Technology, University of Tartu, Estonia.

Dissertation has been accepted for the commencement of the degree of Doctor of Philosophy (PhD) in computer science on 10th of November, 2020 by the Council of the Institute of Computer Science, University of Tartu.

Supervisor

PhD	Jaan Aru University of Tartu Tartu, Estonia

Opponent

Dr	Jakub Limanowski Technische Universität Dresden Dresden, Germany
Dr	Michael Gaebler Max Planck Institute for Human Cognitive and Brain Sciences Leipzig, Germany

The public defense will take place on 21st of December, 2020 at 14:15 in
Narva mnt 18–1019, Tartu, Estonia.

The publication of this dissertation was financed by the Institute of Computer Science, University of Tartu.

ISSN 2613-5906
ISBN 978-9949-03-508-3 (print)
ISBN 978-9949-03-509-0 (PDF)

University of Tartu Press
`http://www.tyk.ee/`

To an uncertain future

ABSTRACT

Over the past decade the free energy principle has been proposed as an unifying theory of brain functioning, stating that all adaptive biological systems try to minimize uncertainty. This principle has been implemented as a process theory, active inference, that lends itself to some counter-intuitive predictions that we aim to test experimentally using modern immersive virtual reality (VR) systems. While gaining insights about active inference, we also developed new understanding and learned about VR research.

The dissertation begins with designing the overall experimental framework and discussing the limits of VR software and hardware solutions. Next, we cover a conducted study on attenuation of self-generated hand movements, validating both the theory and our chosen practical methodology through reaction time measurements. We then further refine our experiment to allow the study of higher-order effects such as contrast perception and self-rated confidence related to visual attenuation.

Our results show great promise for using VR in neuroscience and confirm previous findings in the active inference framework literature. The thesis culminates with a review and assessment of the current wider field of psychological and neuroscientific immersive VR research and offers guidelines for future work.

CONTENTS

Abstract **6**

List of original publications **9**

Introduction **10**

1. Computational neuroscience and virtual reality **13**
1.1. The free energy principle 13
1.1.1. Active inference 14
1.1.2. Predictive Bayesian inference 17
1.1.3. Criticism 20
1.1.4. The free energy principle: conclusion 21
1.2. Immersive Virtual Reality 21
1.2.1. Immersion 22
1.2.2. Presence 25
1.3. Our research approach 26

2. Developing Virtual Reality
(publication I) **27**
2.1. Goals of the program 27
2.2. Design decisions and implementation 27
2.3. Lessons learned 29

3. Self-movement and attention (publication II) **30**
3.1. Methods 31
3.2. Results and conclusions 32

4. Self-movement and visual sensitivity (publication III) **34**
4.1. Methods 34
4.2. Results and conclusions 36

5. Guidelines for VR research
(publication IV) **38**
5.1. Reliability and validity issues 38
5.2. Guidelines 39

Summary **45**

Bibliography **47**

Acknowledgements **56**

Sisukokkuvõte (Summary in Estonian) **57**

Publications **59**
VREX: an open-source toolbox for creating 3D virtual reality experiments 61
Attention is withdrawn from the area of the visual field where the own hand is currently moving 71
Waving goodbye to contrast: self-generated hand movements attenuate visual sensitivity 81
Guidelines for Immersive Virtual Reality in Psychological Research . . 93

Curriculum Vitae **101**

Elulookirjeldus (Curriculum Vitae in Estonian) **102**

LIST OF ORIGINAL PUBLICATIONS

Publications included in the thesis

I Vasser, M., Kängsepp, M., Magomedkerimov, M., Kilvits, K., Stafinjak, V., Kivisik, T., Vicente, R., & Aru, J. (2017). VREX: an open-source toolbox for creating 3D virtual reality experiments. BMC Psychology. DOI: 10.1186/s40359-017-0173-4

II Laak, K-J., Vasser, M., Uibopuu, O J., & Aru, J. (2017). Attention is withdrawn from the area of the visual field where the own hand is currently moving. Neuroscience of Consciousness 2017; 3 (1): niw025. DOI: 10.1093/nc/niw025

III Vasser, M., Vuillaume, L., Cleeremans, A., & Aru, J. (2019). Waving goodbye to contrast: self-generated hand movements attenuate visual sensitivity. Neuroscience of Consciousness. DOI: 10.1093/nc/niy013

IV Vasser, M. & Aru, J. (2020). Guidelines for Immersive Virtual Reality in Psychological Research. Current Opinion in Psychology. DOI: 10.1016/j.copsyc.2020.04.010

The author of the present dissertation contributed to these publications as follows:

I Came up with the idea and developed the overall program design, wrote most of the article.

II Participated in refining the experimental design, collecting the data and co-wrote the manuscript.

III Developed the experimental design, collected the data and created necessary analysis scripts, co-wrote the manuscript as the first author.

IV Conceptualizing the general idea and writing the original draft.

INTRODUCTION

The functioning of the brain has puzzled humans probably since the beginning of thought. Aristotle reportedly believed that the brain was a cooling mechanism for blood and that the bigger it was, the more rational or "cool-headed" the owner (Bear, Connors & Paradiso, 2016). Using the scientific method, modern neuroscientific research has pinpointed the brain as the unequivocal seat of perception, control, consciousness and cognition (and yes, rationality) among other properties. Yet, it is still surprisingly unclear how exactly the brain achieves these properties on a general computational level with such remarkable efficiency (Bear, Connors & Paradiso, 2016). While advancements in brain imaging and the study of single neurons have given us a relatively fine-grained understanding of the basic building blocks of the brain on a cellular level, it has so far been relatively uncertain how the brain works as an integrated system. How does generalized meaningful perception, cognition and behaviour arise from the combination of 86 billion (Azevedo et al, 2009) neurons? What algorithm is running on the brain's biological hardware?

To be sure, there is no lack of competing theories about what algorithms the brain might be running. And one can doubt whether there even is a single all-encompassing approach that can explain the different facets of brain functioning. Nevertheless, over the recent decade one such theory has risen and has been claimed to be a "unifying principle of brain functioning" (Friston, 2010). This is the free energy principle, which has been formulated only recently (Friston, 2005; 2010) but that has a long history. The general idea about the brain as a hypothesis-testing system was proposed by von Helmholtz (1867) and developed further by Neisser (1967) and Gregory (1980), to name a few. These early theories focused on perception and were lacking a concrete computational underpinning, and these shortcomings were evident to the researchers at the time. As Gregory (1980) notes at the end of his article: "It is very curious that we can think conceptually with such effects 'outwards' but not 'inwards'. It may be that developments in artificial intelligence will provide concepts by which we shall see ourselves." The free energy principle broadly says that adaptive biological systems try to minimize surprising situations, i.e. free energy or uncertainty (Corcoran, Pezzulo, & Hohwy, 2020).

This is not to say that there are no other theories. More recent approaches to brain functioning with a clear mathematical component include the Integrated Information Theory that deals mainly with consciousness (Oizumi, Albantakis & Tononi, 2014), cortical columns network model to explain learning (Hawkins, Ahmad & Cui, 2017), the idea of redundant neural population codes (Pitkow & Angelaki, 2017), and Dendritic Integration Theory looking at the cellular mechanisms of conscious processing (Aru, Suzuki, & Larkum, 2020). This is by no means an exhaustive list, but over the past decade the Bayesian brain and predictive coding paradigms have become predominant in cognitive neuroscience (Clark, 2015; Hohwy, 2013; Friston, 2010; Friston et al, 2017; Friston, 2018), explained in more detail in chapter 1.

A formal, hierarchical understanding is vital in order to begin solving the many puzzles of the brain as a complex biological system. A similar issue has been raised in the witty essay “Can a biologist fix a radio?” (Lazebnik, 2002) over the issue of deriving biological models mostly from experiments and not theory. As Lazebnik notes (2002): “Even if a diagram [based on purely experimental research] makes overall sense, it is usually useless for a quantitative analysis, which limits its predictive or investigative value to a very narrow range.” A more recent paper asks: “Could a neuroscientist understand a microprocessor?” (Jonas & Kording, 2017). They note: “If the brain is actually simple, then a human can guess a model, and through hypothesis generation and falsification we may eventually obtain that model. If the brain is not actually simple, then this approach may not ever converge” and “When studying a complex system like the brain, methods and approaches should first be sanity checked on complex man-made systems that share many of the violations of modelling assumptions of the real system." Alas, in order to confirm or debunk a certain theory about the functioning of the brain, appropriate tools are needed.

Over the course of the history of experimental psychology a large suite of different study methods have been developed, ranging from simple behavioural tasks of counting beans in a basket (Jevons, 1871) to functional magnetic resonance brain imaging and in-silico simulations (Einevoll et al, 2019). Visual stimuli have traditionally been either tightly controlled still images, 2D movies or real-life (sometimes ad-hoc or quasi) environments, forcing the experimenter to choose between ecological validity and experimental control. As most researchers opt for the latter, these highly controlled, simplistic environments with minimal confounding variables might be a part of the reason for why the field has so far been lacking a systems level theory of the brain. Ideally we would want the best of both worlds - a complex and realistic environment that is at the same time fully under the experimenters guidance. Pitkow and Angelaki (2017) note: “To reveal the most important aspects of these neural computations, we must study large-scale activity patterns during moderately complex, naturalistic behaviours.”

With the advancement of the field of virtual reality such a paradigm is beginning to emerge, bringing with it a new set of opportunities and also challenges. While virtual reality systems have been touted as “the ultimate display”, promising to eventually deliver convincing simulated environments already for over 50 years (Sutherland, 1965), only in the last decade have we seen a huge increase in computing power and decrease in costs starting to make this vision a possibility for a broad range of researchers. Using commercially available and relatively cheap immersive head mounted displays, researchers can now create convincing computer-generated spatial 3D scenes that allow perfect reproduction of the experimental setting within and between participants. This allows for novel approaches in testing computational theories of brain functioning.

The key aims of the current dissertation are developing immersive virtual reality software for neuroscience research and experimentally studying the active inference framework that relates to the bigger unified brain theory of the free-energy principle proposed by Karl Friston (Friston, 2010). The dissertation comprises different stages, beginning with developing the overall experimental approach, then running experiments related more specifically to the active inference framework (Friston et al, 2017) and finally evaluating the feasibility of current psychological VR research and providing guidelines for future work.

Chapter 1 gives a more detailed overview of the elements of the free-energy principle and also delves into the short history of virtual reality. We then describe the potentials, but also the problems that one faces when beginning to set up a psychological VR experiment. Here we also explain the rationale behind the general approach to our chosen experimental study design.

Chapter 2 then introduces VREX (publication I), an open-source attempt to develop our own experimental methods and also push the field along by helping other researchers conduct their study paradigms. With encouraging results and also valuable lessons learned, we progressed to experimental human research.

In chapter 3 we describe the design and execution of the first study to probe the active inference framework regarding the attenuation of self-generated body movements (publication II). We expand the concepts of attention and what behaviours we can predict from current theory. The methods section introduces new approaches regarding optical hand tracking and putting it all together in order to study the human brain. The chapter concludes that while the results supported our hypothesis, studying perception through reaction time responses is not ideal and the paradigm should be developed further.

Learning from the previous study, chapter 4 focuses on examining the effects of attenuation on subjective contrast reports (publication III). We used a well-known visual paradigm with certain modifications to again probe the active inference framework. We also introduced a new higher-order dimension to the research, looking also at the confidence of which study subjects rate their perceptions. While the overall methods did not differ much from the previous study, the new stimuli and data allowed us to make much broader conclusions.

Chapter 5 deals with the questions of overall reliability and validity of VR in psychological research (publication IV). We provide a literature review from recent years and also tap into our own working experience to evaluate different VR paradigms and offer guidelines to ensure reliable and valid outcomes in immersive experiments.

The dissertation ends with a general summary and broader discussion about the possible futures of neuroscientific brain research, and VR applications.

1. COMPUTATIONAL NEUROSCIENCE AND VIRTUAL REALITY

All living organisms try to remain in an internal homeostatic balance to be able to fend off entropy as best as one can. This goal could also be framed as the system's inherent expectation to minimise average surprise (Allen & Tsakiris, 2018), successfully by-passing situations that the agent has not evolutionarily adapted for while being active in the world. In the present context "surprise" does not have a classical dictionary definition of an astonishing event, but rather quantifies the improbability of some real outcome for the living system (Corcoran, Pezzulo, & Hohwy, 2020). In literal terms, a fish does not want to suddenly find itself in a surprising situation, e.g. out of the water.

So, in order to enjoy a worry-free existence, the agent must calculate every possible "out of the water" situation beforehand and then avoid it. But there is a problem: knowing every possible situation that might surprise an organism in the future is computationally intractable, as its direct evaluation requires the agent to have perfect knowledge of the external world (Friston, 2009). A more realistic and effective approach is needed. This chapter introduces the main concepts behind one dominant approach, the free energy principle and the way it is implemented as active inference and predictive coding. A simplified overview of the paradigm can be seen on figure 1, based on recent conceptualizations by different authors (Friston, 2018; Parr & Friston, 2019; Corcoran, Pezzulo & Hohwy, 2020; Safron, 2020). Figure 1 also acts as a roadmap for the first part of this chapter: In the following the topics from top to bottom will be briefly explained.

1.1. The free energy principle

The concept of free energy minimisation from the field of physics is useful as an approximation for the amount of surprise elicited by sensory inputs of biological systems (Friston, 2010), framed as a function of the agent's sensory and internal states. A system in this context is defined by a Markov blanket that consists of active and sensory states. A state in turn comprises of some set of physical patterns of neural activity (Ramstead, Kirchhoff, & Friston, 2019). Minimizing the free energy forms an upper bound on sensory surprise, enabling the agent to indirectly and efficiently evaluate the surprise associated with its possible sensory states (Friston & Stephan, 2007; Bogacz, 2017) and select actions that reduce, avoid or suppress surprise (Friston et al, 2015). This implies that organisms must generally develop a probabilistic internal generative model of the external environmental causes of its sensory flows (Corcoran, Pezzulo, & Hohwy, 2020), which is both computationally advantageous and biologically plausible even in simple systems (Bogacz, 2017).

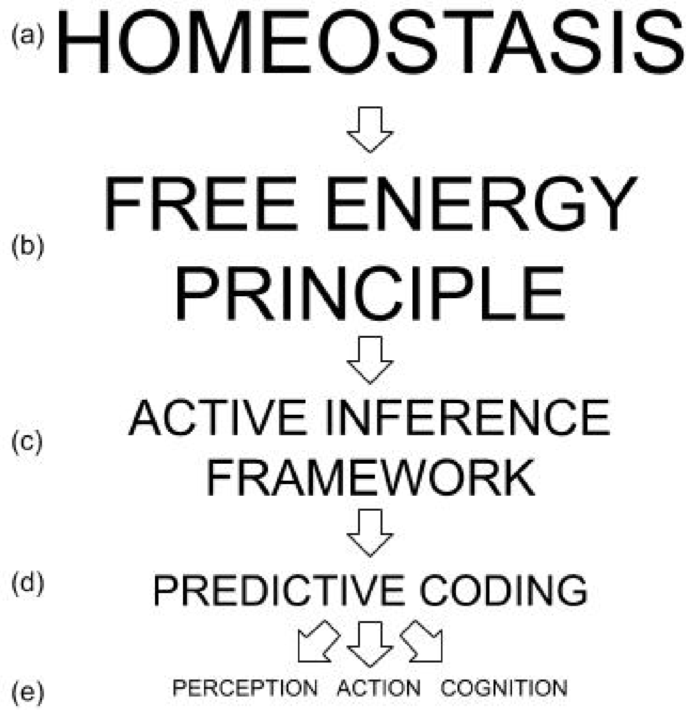

Figure 1. A hierarchy of terms stemming from the ultimate goal of homeostasis (a) to the overarching principle of free energy minimization (b), the underlying process of active inference (c), implemented computationally as predictive coding (d) to achieve perception, action and cognition (e).

The flip side of surprise minimization is expectation maximisation, confirming the quality of the agent's internal model. So the free energy principle posits that "adaptive biological systems conserve their own integrity through free energy minimising interactions which, over the long-term time average, minimise entropy (i.e. resolve uncertainty) and maximise self-evidence." (Corcoran, Pezzulo, & Hohwy, 2020). Such reduction of uncertainty greatly increases the probability of picking out relevant signals from the noise (Linson et al, 2018). The actual process that achieves all this is called active inference (Friston, 2018).

1.1.1. Active inference

Active inference comprises two basic processes: action and perception (Corcoran, Pezzulo, & Hohwy, 2020). Imagine that you are a prisoner in a dark cell. Suddenly you hear some knocking on the outside wall. You are uncertain as to what it is. Your prior knowledge serves up several predictions - it is that annoying woodpecker again, or perhaps fellow inmates throwing rocks at the wall, or just maybe it is Santa Claus trying to bust you out. While a quick probability assessment rules out the third option, other explanations still stand. Passively waiting yields no new information to reduce the uncertainty, so what are the options for the prisoner? Active inference suggests that there are two ways to reduce uncertainty. You could either act by stepping on the chair and climb up to the small window in order to peek outside, or you could alter the model in your brain in favour of one or the other option (e.g. set a high prior that the sound is coming from a bird). You decide to step on the chair, peek outside and get hit on the head with a small pebble, followed by laughter down below. Congratulations, you have used active inference to revise your internal model of the world. This is analogous to what the brain does constantly, as our reality is "ultimately built in the dark, in a foreign language of electrochemical signals" (Eagleman, 2015).

Let's take a closer look at the decision and act of stepping up on the chair, as this can give us a better understanding of the active inference process at work in the sensorimotor domain (Palmer, Zapparoli & Kilner, 2016). It also prepares us for the reasoning behind studies II and III of this thesis. Here active inference prescribes a quite peculiar situation: as the decision of stepping up emerges in the agent's mind, an mismatch is promptly created between the current state (one leg up, other down) and the predicted state (both legs up) (Figure 2).

Figure 2. According to the active inference framework, when preparing a movement, e.g stepping up (left), the brain generates a prediction of what the sensory input of this movement should be once completed (right) and this creates an acute prediction error between the actual current and the predicted sensory states, as the body is still in the initial pose (left). Adapted from (Palmer, Zapparoli & Kilner, 2016).

There are again two broad ways to resolve the error by increasing the mutual information between the internal prediction model and the sensory signals. Firstly, the agent can stay still and revise the prediction accordingly (e.g. they actually did not wish to climb up to the window anyway), so there would be no mismatch between the model and sensory reality. Secondly, the agent can move the legs up, making the sensory input match the predicted sensory input. However, for this to happen, the uncertainty in the current sensory state must actually be increased so that the agent would shift to the predicted sensory state with relatively lower uncertainty (Palmer, Zapparoli & Kilner, 2016). The inverse must also be true - if uncertainty about the first pose can not be increased (e.g. attention cannot be withdrawn from the real sensory input for some reason), action that further increases the prediction error should not take place (Hohwy, 2013). This somewhat counter-intuitive concept is also presented graphically on figure 3.

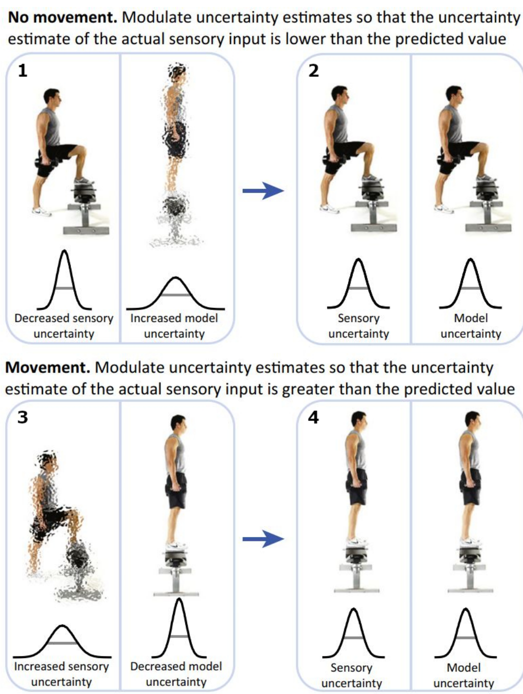

Figure 3. Movement Initiation within the Active Inference Framework. The character shows the action that is currently being performed (left side of all panels) alongside the predicted future action (right side of all panels). The width of the distributions and the clarity of the character illustrate the uncertainty in these values. To minimize prediction errors (panels 1,3), an individual can: stay still and update their prior beliefs so that the predicted sensory input matches the actual sensory input (panels 1,2); or move, so that the actual sensory input matches the predicted sensory input (panels 3,4). Adapted from (Palmer, Zapparoli & Kilner, 2016).

Action is then an efficient way of minimizing prediction error by changing the world in order to fit the brain's generative model to the internal prediction (Clark, 2015). When action is either unavailable or undesirable in a given situation, perception can be used to change or update one's internal states in order to resolve uncertainty about the hidden external causes of its sensory fluctuations, an operation akin to learning. In our prison example, in the absence of a chair to reach the window, one could mistakenly decide that what they heard was indeed a bird pecking away at the wall and with no further evidence available from the world, the internal uncertainty would be (erroneously) resolved. A common neurocomputational implementation of such a process is called predictive coding (Friston, 2018). While the main contributions introduced in this thesis (publications II and III) stay at the broad level of the active inference framework, for a clearer picture of the whole theory it is worthwhile to also give a short primer on the further subtopic of predictive coding. We hope that this gives the unprepared reader a better grasp on the different keywords used by various researchers, as there can be some confusion. It is argued and we agree that there currently exists a systematic misrepresentation in the literature about active inference as being a part of a larger theory to "unify life and mind" and related, but distinct Bayesian formulations, centred more specifically on the brain (Ramstead, Kirchhoff, & Friston, 2019). The latter are usually named as either predictive processing (Clark, 2015; 2016), predictive coding (Rao & Ballard, 1999) or the prediction error minimisation framework (Kiefer & Hohwy, 2018; 2019). Next we will look at three intertwined concepts: generative models, predictive coding and Bayesian inference.

1.1.2. Predictive Bayesian inference

An agent attempting to understand as many relationships among its experiences as possible without external supervision or reinforcement, may learn an internal generative model about the environment. The benefits of such an approach in the brain are many - to better interpret sensory inputs, to predict and prepare actions and to support flexible knowledge transfer to novel situations (Berkes et al, 2011; Hassabis et al, 2017; Kriegeskorte & Douglas, 2018).

A great explanation of a neuronally implemented generative model has been given by Linson et al (2018). They describe a first-time visit to a new university campus. Before the visit, one might generate a basic internal model from prior experience of other campuses about what to expect in the new location. On an actual visit to the campus, our generative internal model allows us to extrapolate with a high probability of there being a cafe, even if in actuality we are mistaken. This error from our exploratory process updates or further details our generative model for the particular campus, allowing us to exploit the model for our purposes, for example to find lunch. If finally every sensory impression meets our predictions while navigating the campus, we have successfully inverted our generative model. What we predict is true. This is the process of learning to recognize the external causes in relation to their context-dependent sensory consequences.

A similar process has also been shown in animals, with internal models of the external world progressively optimizing with age and leading to spontaneous brain activity that seems to predict the natural environment of the animal (Berkes et al, 2011). Berkes et al measured the cortical signals of awake ferrets in darkness, showing that the adult animal's brain activity while resting was significantly better matched to neural activity evoked by natural images than the activity recorded from young animals. This demonstrates the progressive adaptation of internal generative models to the statistics of natural stimuli at the neural level. The mechanism of building generative models that provide abstract prior knowledge is best described by Bayesian inference (Kriegeskorte & Douglas, 2018).

Predictive coding can be seen as one plausible implementation of perception. In the simplest terms, the predictive coding framework proposes that experiencing the world is a two-way street. We have incoming external ("bottom-up") sensory input, but also internal ("top-down") predictions that rely on previous experiences and expectations, that together form a percept. Computationally this algorithm takes inspiration from the "Wake-Sleep" algorithm from Hinton and colleagues (1995) that consists of bottom-up "recognition" and top-down "generative" connections. Constant updates are made to the model to be as precise as possible using Bayesian computations. In order to introduce the basics of Bayesian inference, let's consider a simple example by observing a picture of a duck (figure 4).

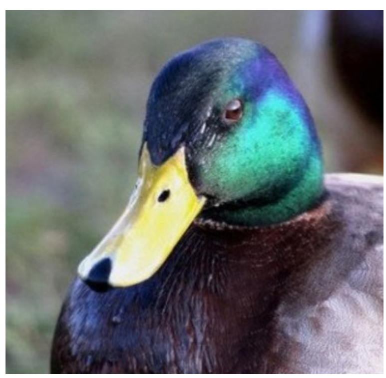

Figure 4. An ordinary duck. Image from public domain.

The experience of perceiving this particular bird is probably effortless to most of us at first - it is unmistakenly just a picture of a duck, perhaps even a boring one. The image (bottom-up sensory information) fits well with our overall internal idea (top-down prediction) of what a duck should look like. The only peculiar aspect of the picture is possibly the fact that it had found its way inside a computer science doctoral dissertation. The photo is also of acceptable quality, minimizing any problems with possible ambiguity due to noisy incoming data. However, our conscious experience is quite malleable and updated top-down predictions can have a big influence on our final perception (posterior), even when the bottom-up information seems solid at first. For example, consider the following little-known extra piece of information: all ducks are actually wearing tiny dog masks (figure 5). Then observe the original photo again and enjoy the mismatch between bottom-up and top-down information streams.

Figure 5. All ducks are actually wearing dog masks. Image from public domain.

If the exercise was successful, the reader may have felt their predicted duck model being temporarily updated with new evidence about the omnipresence of dog masks, causing a significant error between two streams of information. The long term effects of this new information however depends on the precision or credibility (likelihood) assigned to it - one can dismiss it as simply an unfunny joke or have their internal duck model forever changed, alluding to the popular internet meme “what has been seen, cannot be unseen”.

Predictive coding is founded on the idea that top-down and bottom-up information flow across the hierarchical structure of the cortex implements hierarchical probabilistic inference about the causes of sensory data. Within this framework predictions represent the hidden causes in the world that cause sensory states in the cortex. Prediction errors signal the discrepancy between the incoming sensory data and the predictions. Prediction errors are weighted by the precision of the predictions and sensory signals. These weights determine how strongly a given prediction error influences the updating of prior beliefs. A simplified example of this process is seen on figure 6, while a more in-depth discussion on related functional brain anatomy can be found elsewhere (Brown & Friston, 2012; Seth, 2013).

Figure 6. Simple example of Bayesian inference. The curves represent probability distributions over the value of a sensory signal (x-axis). Here low precision-weighted (noisy) likelihood of sensory signals (red) diminish their influence on the posterior belief (green) as compared to the more precise prior expectation (blue). Prediction error can be quantified as the difference between the distributions of likelihood and prior. Adapted from Seth (2013).

Rao and Ballard note in their seminal work (1999): "The general idea of predictive coding may be applicable across different brain regions and modalities, providing a useful framework for understanding the general structure and function of the neocortex." Indeed, the last two decades have seen a large increase in research regarding the theoretical study and practical application of the predictive brain framework. The approach has been used to explain how processing happens in different modalities like visual, auditory, somatosensory, interoceptive and proprioceptive (Adams, Shipp, & Friston, 2013) and psychopathologies like autism (Palmer, Lawson, & Hohwy, 2017) and schizophrenia (Corlett, Honey, & Fletcher, 2016). In the case of autism, one of the hallmark symptoms is overwhelming of novel social situations - while a person's appearance does not usually change dramatically, small differences like a different haircut or a slightly lower voice are common things that people with autism spectrum disorder pick up on. The proposed mechanisms for this has been over-reliance on bottom-up sensory precision or the high imprecision of top-down predictions (Friston, Lawson, & Frith, 2013; Lawson, Rees, & Friston, 2014; Pellicano & Burr, 2012; Van de Cruys et al, 2014; Tulver et al, 2019). In schizophrenia patients, one common symptom is the appearance of hallucinations - perceiving something that is not there in objective reality, although the brain might receive some ambiguous input. This has been attributed to abnormally strong priors (Powers et al, 2017; Schmack et al, 2013; Teufel et al, 2015; Tulver et al, 2019).

1.1.3. Criticism

The free energy principle is not without objections. One common criticism questions if surprise minimization is really the driving goal of all biological systems, as most creatures do not prefer to spend their lives in perfectly predictable quiet corners. This thought experiment has been dubbed the "Dark-Room Problem" (Friston, Thornton & Clark, 2012; Klein, 2016; Sun & Firestone, 2020a). Supposedly in a perfectly dark room free of surprising stimuli an organism can avoid all prediction errors until it dies from starvation, an example that boarders on the absurd. Such vignettes are seen by some as proof of active inference not being a good theory of action (Sims, 2016). Other critical views stem from the seeming misconceptions of the free energy principle - that it only describe the mental operations of "the best organisms" (Klein, 2016) or that the principle does not apply to organisms less complex than humans (Sims, 2016). According to opponents, the Dark Room Problem is unsolvable by current formulations of the free energy principle. Firstly, relying on self-predictions (organisms predicting that their overall tendency to hang around in dark rooms voluntarily is low, so the situation is actually very surprising) is deemed not convincing and secondly, adding endless caveats to predictions (organisms do like to be in dark rooms, unless they are hungry, or they hear some strange noise outside etc.) runs the risk of making this a post-hoc explanation (Sun & Firestone, 2020a).

The critics proposed solution to this conundrum is to back away from claiming the grand ambitions of the free energy principle as a theory-of-everything (Klein, 2016; Sun & Firestone, 2020a). Proponents however claim that prior predictions are indeed specific to different species and even particular individuals, thus in case of most humans, the dark room is simply not an attractor (Friston, Thornton & Clark, 2012). Others have proposed that organisms try to minimise surprise over the long term, so "an agent needs to be a curious, sensation-seeking agent in the present." (Seth et al, 2020). This argument has been countered by the lack of specific time frames, alluding ironically from a day to a millennia (Sun & Firestone, 2020b). It must also be noted that while many contrived example situations for negatively evaluating the free energy principle can be readily thought up (Sims, 2016), few of our actual real life goals are absolute, varying greatly in their susceptibility to encountered evidence (Van de Cruys et al, 2020). Lastly, one must remember to differentiate between basic elements of cognition and more complex phenomena that arise from these elements - otherwise unrealistic expectations are set for the explanations of the former to also perfectly encapsulate the latter, e.g. "personal growth, aesthetic experience or moral worth" (Sun & Firestone, 2020b).

1.1.4. The free energy principle: conclusion

We have briefly described how the free energy principle could be realized by active inference of action and perception that implement predictive coding with the help of generative models and Bayesian computations. As the field is developing rapidly, the current outline will of course not be the final truth. For example Parr and Friston (2019) have recently introduced a further concept of generalised free energy, or perception can also be considered as inseparable from action (Ramstead, Kirchhoff, & Friston, 2019).

The free energy principle itself is a broader framework, not a testable hypothesis in and of itself (Seth et al, 2020). The implementations of the theory at its current form however already provides many interesting claims that can be tested and validated experimentally, thus we now turn our attention to introducing our main tool of choice for the thesis, immersive virtual reality.

1.2. Immersive Virtual Reality

To avoid discussing the undoubtedly interesting, yet distant histories of paleolithic cave paintings, wild psychedelic drugs and Victorian era wooden stereoscopes, it is useful to first define the term “virtual reality”. In our research it has been the following: a technologically generated experience where the user can and will act as if in reality (Vasser, 2018). This definition highlights the three crucial components of a true modern virtual reality experience: created through screen technology (preferably head-mounted displays), allowing the user to move and interact naturally with their environment (at least some level of free-roaming ability and tracked hand controllers) and react to the virtual content instinctively (e.g. hesitating when stepping over a virtual cliff, jumping back when confronted with a frightening situation).

While the first headset to fulfil all the aforementioned criteria at least on some very basic level was built already in the seventies (Sutherland, 1968), progress in the field was initially slow due to limits on computing power, display technology and real-time computer graphics. Another hurdle was the high price of the systems, making early VR accessible to only high-end medical, military, automotive and scientific visualization industries. Progress picked up in the 90s. On the academic front, the IEEE Virtual Reality conference has been held since 1993. The decade saw commercial releases (and flops) of Sega VR and Nintendo's Virtual Boy among other systems. These technologies were crude and generally did not live up to the advertised hype.

A turning point arrived in 2011, when researchers from the University of Southern California unveiled their rough design for an immersive smartphone-based head mounted display (Olson et al, 2011). The approach to use abundantly available smartphone displays to screen VR content was a leap in terms of lowering costs and improving visual quality. A very similar design was used for the influential Oculus Rift Development Kit 1 that was released for the developer community in 2013 for about 300 dollars. Another important milestone was in 2016, when HTC rolled out the commercial Vive system, which included positionally tracked hand controllers and allowed for "room-scale" tracking.

Alongside the development of the hardware there has also been progress in the software front in terms of optimizing, new interactive content, development tools, and a high demand for VR developers.

It is easy to see the promise that virtual reality systems bring to psychology and neuroscience. Since the beginning of VR the aim has been to produce an ultimate display - a screen to replace and emulate all other viewing devices. VR has also been hailed as the empathy machine, being able to elicit a wide range of emotions in people (Schutte & Stilinović, 2017). Research has also shown that VR experiences can also engage the most primal parts of the brain, raising heartbeats and producing cold sweat in virtually dangerous situations (Slater et al, 2009b). It is possible to embody virtual avatars as if they were real extensions of one's body (Kilteni et al, 2012) even if the digital self is not the correct shape, size or species (Schettler, Raja, & Anderson, 2019). In order to understand the potential and limits of VR better, it is worthwhile to dive into the terms immersion and presence.

1.2.1. Immersion

Although the terms "immersion" and "presence" are commonly used interchangeably by enthusiasts, they are not the same phenomena (Slater et al, 2009a). By immersion we mean all the physical hardware and virtual software that enables any sort of virtual reality experience to occur. This covers VR headsets, controllers, earphones, computers, game mechanics and features. Anything that can enhance immersing the player into the virtual world, to “take over” their senses. More specifically, an immersive VR system fills the following criteria: true stereoscopy, wide field of view, 6 axis positional tracking, low latency, high refresh rate and convincing content. Let’s look at each component in slightly more detail.

Stereoscopy is how most of us perceive the world - as both eyes receive a slightly different retinal image, depth perception emerges in the brain. Good VR systems work much in the same way - HMDs have two lenses that both show a slightly different view from the split screen(s). This is the reason why first-time users usually shout something in the lines of “wow, it is so 3D!” when trying on a headset. This effect is not a given however. For example, many 360 videos are very much two-dimensional and even though the content is shown in 3D space, projected on the inside of a virtual sphere, the experience is that of watching a very large curved 2D screen. Another reason why current screen technology is problematic stems from the fact that the brain uses several cues for depth perception other than binocular disparity (eyes receiving different images). As HMDs present objects of varying depth objects on a fixed depth screen, many of these cues break down (Harris et al, 2019).

Wide field of view extends to over 100 degrees and gives the user an illusion of being visually completely immersed in the virtual world with no reference frames of the screen visible. This is usually achieved through high quality optics using convex lenses in front of the headset displays. The side-effects of such image magnification is lower pixel density and several artefacts, such as image warping and chromatic aberration (colour dispersion at the edges of the lens). These issues can be mitigated by super-sampling the original image and initially distorting in the opposite direction with software, so the lens actually corrects the final image. Newer headsets use both hardware and software tricks to arrive at “human-eye resolution” output (Varjo Technologies).

6 axis or degrees of freedom tracking allows the user not only to look around (pitch, roll, yaw) but also to move around in 3D space (X, Y, Z). The technology to achieve this can be broadly split in two categories: outside-in and inside-out tracking systems. Outside-in solutions usually encompass some external cameras (e.g. Oculus Rift) or beacons (e.g. HTC Vive) to determine the exact position of the users headset and controllers. Due to sensor positioning, here a further distinction is made between 180 and 360 degree or “room scale” tracking, specifying if the user can turn and stand in arbitrary ways (360) or is required to mostly face only one certain direction (180). Inside-out tracking relies on cameras positioned on the headset itself to determine the movement related to the external environment, so in a sense the whole world becomes a tracking marker. While the loss of tracking during a VR experience can be the most jarring experience, avoiding it is not guaranteed in any system. The user can either stand in the play area in a certain way to accidentally block the sensors on the headset, move too fast for the camera systems to track or experience sunny weather that causes trouble for all optical tracking systems.

Low latency signifies the fast "motion-to-photon" time, that is the amount of milliseconds it takes the VR system between registering user motion (e.g. head turning) and displaying the corresponding new frame on the headset screen. The positional information has to travel from the headset to the program, which must calculate the new image and send it all the way back to the display. In common practice it is agreed that latency lower than 20 ms is necessary for a convincing VR experience, and anything higher will produce cybersickness (i.e. nausea, motion sickness, lightheadedness). The latter is comparable to experiencing the dreaded network lag in a very visceral way - taking a step forward in virtual reality, and having the visual effects arrive a second later. The bottlenecks to latency lie mostly in software, so big strides have been made to develop drivers that are more directly in contact with hardware, bypassing rendering steps with "time warping" techniques, predicting user movements and overall optimizing the experiences.

High refresh rate is the technical capability of the given headset screen to display many frames quickly. The current industry standard refresh rate is at least 90Hz, although even 60 frames per second can be acceptable and some devices have settled on 72Hz. While more frames sounds better in strictly hardware terms, here the software becomes a limiting factor, as rendering out images with high resolution and extreme frame rates can be demanding, especially when the graphical processing unit has to draw images for both eyes, doubling the overall load. The solutions lie in both possibly upgrading the hardware and applying some software optimizations. Common approaches are instanced stereo (drawing only some of the objects separately for both eyes), lower super sampling, static lighting and minimal transparency in the virtual scene (as these are computationally heavy operations to perform). As VR developers say: "Fake everything you can" - as dynamic shadows and pretty materials "cost" a lot in the rendering pipeline, optimization is key to reach the target frame rates of modern headsets. With a 90Hz display the time to deliver each frame is about 11 milliseconds and this window can not be missed if immersion is concerned.

Finally, convincing content is paramount. If the VR system is flawless in all other respects, but the content is dull or poorly designed (albeit optimized), immersion breaks down. But if the content is engaging, many shortcomings in other components can be ignored. This is evident from the myriad of internet videos showing someone wearing the historic Oculus Development Kit 1 with a sub-optimal field of view, only rotational head tracking, poor screen resolution and 60Hz display, yet still screaming their lungs out when riding the virtual roller coaster. While early content for modern VR systems were mostly clunky experimental builds of some short experiences or direct ports from classical desktop games, over the years the developer know-how on how to extend the "suspension of disbelief" of the virtual world as real has greatly increased. A great VR experience should include an ergonomic locomotion system, clear tutorials, in-game instructions and many interactions similar to real-life (Kourtesis et al, 2019). Unplayable content most likely implements smooth locomotion, complex control schemes and deep menu systems, small hard-to-read text and stressful interactions (Vasser & Lomp, 2018).

1.2.2. Presence

We try to enhance immersion to achieve a longer uninterrupted flow of presence. By presence we mean the psychological perception of one's surroundings as mediated by both automatic and controlled mental processes, an experience of a different reality (Pillai, Schmidt & Richir, 2013). The mapping of different states between the real and the virtual has been historically described as a continuum (Milgram et al, 1994). A more recent approach focusing on evoked presence was offered by Pillai et al (2013). Their framework helps to answer the question of what level of presence is useful for conducting VR research. On their proposed three pole model of a Reality-Presence Map we find Primary Reality in the middle with self- and media evoked realities at the extremes (figure 7).

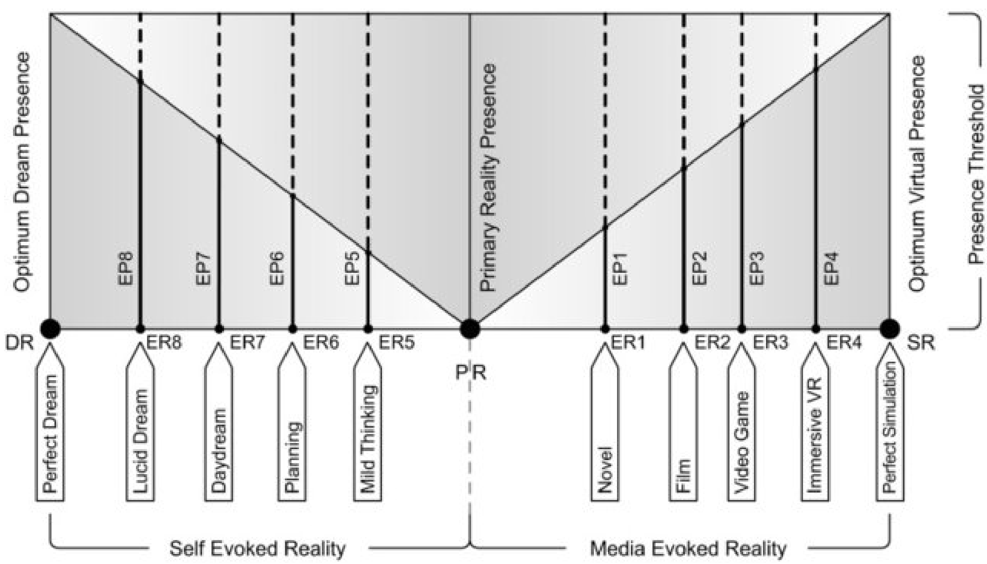

Figure 7. An example range of different possible realities evoked either by media (ER1-ER4) or self (ER5-ER8). EP1-EP8 denote the levels of evoked presence. PR stands for primary reality. The end-points of the scale stand for dream reality (DR) and simulated reality (SR). Note the position of Immersive VR (ER4). From (Pillai, Schmidt, & Richir, 2013).

A common hype surrounding VR is the proposed ability of commercial systems to induce a "perfect simulation" (figure 7, SR), so the user would mistake the digitally created world as real. However, for neuroscientific research this level of absolute immersion is not actually favourable and a level of immersive VR (figure 7, EP4) should be strived for. We do not want for the study subject to completely forget their real world surroundings, leading to some possibly unwanted behaviours, for example trying to sit on a virtual chair that does not exist in reality or walk through a real wall that does not exist in the VR environment. The idea is perhaps easier to explain when looking at the opposite end of the evoked reality scale. Dream reality denotes a regular dream where the sleeping person uncritically accepts what is happening as real and reacts accordingly, with no knowledge of the Primary Reality (their body actually lying in a bed). An experience slightly lower in evoked presence would be a lucid dream (figure 7, EP8) - the phenomena of becoming fully conscious of the primary reality, while still remaining in the dream state (LaBerge, 2009). This allows the person to take full advantage of the very convincing dream reality, while having higher-order knowledge about the primary reality. Thus they are able to remember and conduct intricate experiments while in this "hybrid" state. In a perfect dream scenario the sleeper loses voluntary control over their actions. The experimentally useful VR appears to follow the same logic - immersive reality allows the player to know that the surrounding is unreal, albeit still convincing. In a perfect simulated reality, called superrealism (Slater et al, 2020) or popularly "the Matrix", the user loses all touch with primary reality. While it would still be technically possible to communicate our study instructions through virtual screens, phones or pop-up messages, the experience of simulated reality could have unexpected psychological side effects when returning to Primary Reality.

1.3. Our research approach

There were three considerations that influenced the choice of our concrete research approach. First, the background of the researchers was mainly in psychology, neuroscience and software development. We had ample experience with experimental human studies, but were less familiar with mathematical simulations. Second, the available hardware and budget at the time greatly narrowed down the physical tools to widely available commercial VR headsets. While we did explore the options to combine free-roam VR with electroencephalographic brain measuring devices, this was deemed either too clunky, costly or noisy to justify the effort. Third, given the need to have a sufficiently large sample size for statistical rigor, we decided to not venture into the realm of directly recruiting and studying people with mental disorders.

Thus our main interests in the current work lie in the subjective and behavioural responses of normal study subjects, and less in cell-level neural activity, akin to the distinction by Ward (2019). In the aforementioned free energy principle paradigm our level of inquiry places in and corresponds to functional-computational modelling of behaviour, i.e. active inference (Safron, 2020).

2. DEVELOPING VIRTUAL REALITY (PUBLICATION I)

To begin our work in the newly founded Computer Graphics and VR lab in the Institute of Computer Science in the University of Tartu and harness the potential and combat some of the problems of VR development for ourselves and others, we set a goal to develop an user-friendly open-source toolbox for VR research in the fields of experimental psychology and neuroscience. Similar programs available at the time (2016-2017) were either too general-purpose, complex, of low visual quality or costly. We will cover the aims of the program, explain some design decisions and conclude with lessons learned from this project. For more details, please see publication I.

2.1. Goals of the program

Our main aim was to create an open-source add-on to the popular game engine Unity (Unity Technologies) that would be powerful in terms of graphics capabilities and features, yet usable without much specialized programming knowledge. The former was important for increased presence, the latter stemmed from the fact that we had both computer science and psychology students working on developing the program. The development followed the release of Oculus Rift CV1 headset, as we figured there is an increased interest among psychology and neuroscience researchers to use this novel technology, but a lack of experience with real-time game engines (in particular, knowledge of 3D modelling and texturing, game logic and scripting). Using Unity Game Engine allows for powerful built-in features, high visual quality and fast code iterations. Possible topics that could benefit from VR include spatial navigation, perception and motor control (Scarfe & Glennerster, 2015; Shinoda, Hayhoe, & Shrivastava, 2001).

2.2. Design decisions and implementation

Even today it is difficult to construct large outdoor scenes in VR that are detailed, convincing and also perform well in computational terms. VREX was designed for virtual indoor experiments to circumvent these problems, thus the basic elements of the toolbox are environments consisting of single or multiple rooms without windows. Multiple environments can be sequenced to create experiments. A typical pipeline for an experiment can be graphically seen on figure 8.

While the program is written in the language of C#, VREX provides a dedicated graphical user interface inside the toolbox to give the user intuitive access to common operations within the program. Simple menus allow creating and modifying environments and build experimental plans with different stages.

An useful feature for randomizing the layout of different rooms is the ability

Figure 8. A common pipeline for an experiment. Creating the environments (a), setting up the experiment structure (b) and a point-of-view from running the study in VR (c).

to auto-generate the environment, as this feature both saves time and provides an unique environment layout. After an environment is created either procedurally or by hand, it can then be populated with 3D objects from the available library either automatically or manually. There are a number of openly licenced objects available in the toolbox. For fine-tuned control VREX comes with an editor where all the objects in the room can be manually adjusted in the 3D world or new objects added. VREX originally supports two experimental paradigms - change blindness and false memory. These paradigms were chosen due to our prior experiences showing that applying VR in these sub-fields can generate many new experimental approaches.

Change blindness is the phenomenon of not noticing big visual changes in the scene, if the change happens out-of-sight (Simons & Rensink, 2005). Objects marked in the change blindness experiment in VREX can be modified whenever the object falls outside the field of view of the VR headset. The researcher can choose to change the object's visibility, colour, or location. The chosen appearance change will alternate between two set options. For example, a virtual chair in the corner of a room can be seen, and after momentarily looking away and turning back, the chair will be hidden. Repeating the movement, the chair will be again visible. After identifying the changing elements in an experiment, the participant can click the response button, which logs the response time.

In false memory experiments VREX supports logging all the objects seen by the participant during a trial while moving through the rooms and later modifying their position in the same environment or presenting them differently for cued recall. Recall only contains items seen by the participant and optionally distractor items that did not appear in the experiment, chosen by the experimenter. Recall can also have a set time limit.

Additional features include adding custom models, spatial audio integration, four different VR locomotion systems, experiment data logging and timing, template experiment for easy custom developments, compatibility and extensibility.

2.3. Lessons learned

Our aim was to focus on specific types of experiments in order to simplify the workflow and available options both to ourselves and fellow researchers. Thus the experience might feel limited and the user can not access the full potential of Unity through the toolbox menus. While the program gives many options to configure experiments, it does not cover all possible variations and even in our lab we did not choose to build more specific follow-up experiments using VREX as our thoughts about the active inference framework developed in a different direction. Thus we have discontinued the project for now. Often citing our work, other similar projects have recently emerged, e.g. UXF (Brookes et al, 2019), NavWell (Commins et al, 2019), VRmaze (Machado et al, 2019) and others (Rodríguez et al, 2017; Wiener et al, 2019).

3. SELF-MOVEMENT AND ATTENTION (PUBLICATION II)

Having gained valuable and necessary experience with VREX in developing VR applications with real-time game engines for neuroscientific studies, next we started looking for ways to specifically probe the active inference framework in novel ways. As explained in chapter 1, active inference posits among other things that a movement can only occur when attention is withdrawn from the current sensory input. For an everyday example, consider the act of running, where one's own hands (two relatively big objects) move through the lower part of one's visual field. If there was an external object moving at the same position (e.g. a big dog) it would immediately capture one's attention quite involuntarily. Nonetheless, our own hand movements in the same visual area are not usually noticed — they appear to be suppressed from our experience.

How can this be? It is well known that the brain predicts the sensory consequences of its own movement and that these predictions result in attenuation of the respective sensory signals (Von Helmholtz, 1867; Sperry, 1950; Blakemore et al, 1998; Bays et al, 2006; Friston, 2010; Clark, 2015). Also, several studies have shown that the perception of external tactile stimuli is suppressed during the movement of the subject's own hand (Juravle et al, 2010; Juravle & Spence, 2011). As noted above, own body movements can also lead to visual changes in one's field of view that have to be attenuated as well, in order to minimize surprise. Put differently, it would be very difficult to act when always paying full attention to own movements, thus taxing the perceptual system and possibly missing potential real threats in the external world.

From this notion we devised the basic idea for the experiment: having some stimuli change at the exact visual location of the participant's moving limb and assessing whether the perception of these stimuli is measurably impaired. The active inference framework proposes that the reduction of sensory precision that is needed for optimal processing works through withdrawal of attention from the sensory prediction errors (Brown et al, 2013, Hohwy, 2013; Clark, 2015). This means that in our study the subject's own hands should be attenuated from vision by withdrawing attention from the area of the visual field where the hands are predicted to move to reach their desired destination. Thus, we hypothesized that the hand movement that we asked out study participants to perform is maintained by continuous attenuation of sensory input from the moving hand's position. Using VR and related technologies we could precisely track the participants hand position and render it invisible to them in order to rule out the effects of visual occlusion or disturbance and study the effects of predictions on attention.

3.1. Methods

Using the Oculus Rift CV1 (Oculus VR, LLC) headset, Leap Motion (Leap Motion, Inc;) hand controller and specialized software we constructed a visual search task that allowed us to control the whole visual field of the participants, precisely track their hand movements and display stimuli behind virtually invisible hands. The typical experimental design can be on figure 9.

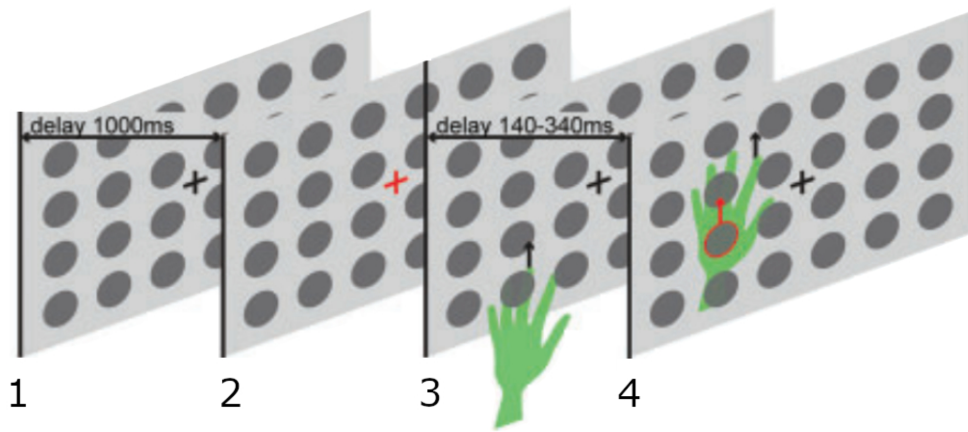

Figure 9. Overview of the experimental design. Each trial started with the participant fixating their gaze on the black cross in the middle of the field of view (1), the cross turning red and signalling the hand movement to start (2), participant moving the hand up on a pre-trained trajectory (3), the stimulus changing either direction or colour behind the hand or next to it (4) and finally the participant clicking the mouse button with the other hand as soon as the stimulus change was detected. Note that the hand was invisible to the participant in VR and the stimuli were much smaller and placed randomly.

The stimuli consisted of multiple small horizontally oscillating spheres. In the first two experiments the targets of the visual search were defined through a sudden change to horizontal movement direction, and the third experiment utilized rapid colour changes. The latter condition was important to rule out the possibility of explaining our results through the efference copy theory (e.g. Blakemore et al, 1998; Clark, 2015), claiming that the brain uses a “copy” of the motor commands (e.g. direction and speed) to subtract the predicted sensory consequences from the actual sensory input. A change in stimulus colour should however be wholly unrelated to internal motor commands. We measured the reaction time of reporting the noticing of abruptly changing target stimulus that was triggered shortly after the start of the hand movement. Crucially, the target stimuli could appear either behind the invisible hand or next to it, the latter being another control condition.

Physically the participants were wearing the Oculus Rift HMD and sat behind a regular table in the darkened lab room with their left hand executing the pre-trained movement on each trial and right hand holding the computer mouse to give responses.

3.2. Results and conclusions

Data preprocessing consisted of removing trials with reactions times below 100ms and above two median absolute deviations, trials where the stimulus appeared too far from the eye gaze fixation point and participants with extremely long reaction times or with less than 10 trials in one condition after preprocessing. The final analysis was done with 53 participants in total. Our results were the first to show that the brain indeed significantly attenuates visual perception in the part of the visual field where the own hand movement is predicted to occur in line with the active inference account. Our measured reaction times to the suddenly moving targets behind the invisible hand were significantly slower than in the control conditions (mean difference 10ms; $t(14) = 2.62$, $P = 0.010$, $d = 0.23$;). The results held even when the target changed color instead of direction (mean difference 17ms; $t(23) = 2.15$, $P = 0.021$, $d = 0.14$). The mean reaction times between conditions can be seen on figure 10. Crucially, there was no difference between the control conditions.

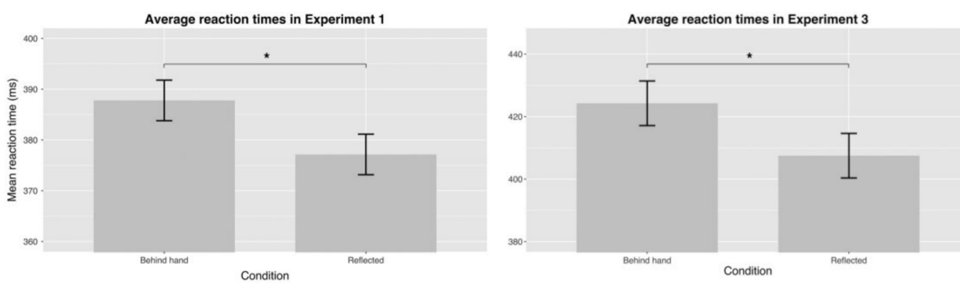

Figure 10. Mean reaction times. Experiment 1 (left) used movement change as stimulus, while Experiment 3 used colour change. In conditions "Behind hand" the stimulus appeared at the position of the moving hand, while "Reflected" showed the stimulus on the opposite side of the visual field.

The results of the first experiment could be explained by the efference copy theory (e.g. Clark, 2015), as both the hand and the stimulus were moving simultaneously and in the same direction, thus allowing the brain to simply subtract the predicted sensory consequences from the input without attenuating attention. However, in our third experiment the stimulus had a sudden colour change with still significant results on mean reaction time, ruling out the efference copy theory, as colour is not related to the motor commands of the hand movement. Taken together, these findings provided direct support for the active inference account of sensory attenuation. This framework posits sensory attenuation of even colour processing due to sensory precision in general being reduced in order for movement to occur (Friston, 2010; Hohwy, 2013; Clark, 2015). More generally, the results demonstrated that the usage of novel VR tools opens up new exciting avenues of neuroscientific research related to attention, sensory attenuation, and agency. However, reaction times are not the most informative measures of perception in general, relying on many different processes to get from initial perception to finally clicking the response button. So next we wanted to go beyond the paradigm used in publication II and test visual perception more directly. For example, by measuring the subjective sensitivity to low-contrast stimuli (e.g. Burr et al, 1982).

4. SELF-MOVEMENT AND VISUAL SENSITIVITY (PUBLICATION III)

As shown in chapter 3, participants are slower to detect sudden stimulus changes (both movement and colour) in the condition where the change occurred behind their precisely tracked invisible hand in VR (publication II). But the effect of lowered attention in previous work was measured only through longer reaction times — an indirect measure of the phenomenology or quality of perception. In publication III, we sought to extend these results to probe the effect of self-generated movement prediction on perception more directly by studying the effect of self-movement on perceived contrast of experimental stimuli. In addition we also wanted to investigate the effect of self-generated movements on metacognition, a higher-order process that underlies the ability to monitor and control one's own mental states (Koriat, 2007). In essence we studied how confident one felt in their response to the contrast estimations.

In experimental psychology, subjective contrast studies have often used specific stimuli called "Gabor patches" that consist of black and white sinusoidal gratings with a Gaussian envelope, different wavelengths and orientations. A key feature to also modify is the luminance contrast, making the Gabor patch either clearly visible or so faint as to stay below or near the subjective detection threshold. When the stimuli is subjectively difficult to detect, metacognition, that is the ability to monitor and control one's own mental states (Koriat, 2007), becomes crucial. Previous studies focusing on the interplay between attention and metacognition had yielded mixed results, with either positive, negative or no correlations (Wilimzig et al, 2008; Kanai et al, 2010; Rahnev et al, 2011; Sherman et al, 2015). Thus we decided to directly measure the influence of sensory attenuation on metacognition.

Taking in consideration our previous results from publication II that attention is withdrawn from the part of the visual field where one's own hand is moving and knowing that the deployment of attention affects perceived contrast (Carrasco, Ling, & Read, 2004), we reasoned that the subjective contrast of objects in the region of the visual field where the hand is moving should be significantly reduced. If so, this would again give some evidence to the active inference framework. Asking participants to rate their confidence in their response on each trial, and assessing the relationship between objective performance and subjective confidence (Galvin et al, 2003) allowed us to additionally explore if self generated movements also alter metacognitive accuracy in the visual region under scrutiny.

4.1. Methods

We largely based our study design on our previous work (publication II) and the experimental approaches used by Carrasco et al (2004). Our participants once again found themselves wearing a VR headset, sitting behind a desk, repeatedly raising their hand in a controlled manner and this time reporting the orientation of one out of the two very briefly (133ms) shown Gabor patches with the subjectively stronger contrast. We also alternated the hands used for movement and responses between left and right in the middle of the long experiment so as to balance for possible preferences for one hand over the other and to mitigate fatigue. Two experiments were conducted where participants moved their invisible virtual hand to overlap one of the two quickly shown target stimulus and they were then asked to perform a forced choice task to report the grating orientation of the stimulus they had perceived as more contrast. In experiment 2 we additionally probed metacognitive accuracy as participants were asked to also report the confidence in their decisions after every trial on a 4-point scale. An overview of the paradigm can be seen on figure 11.

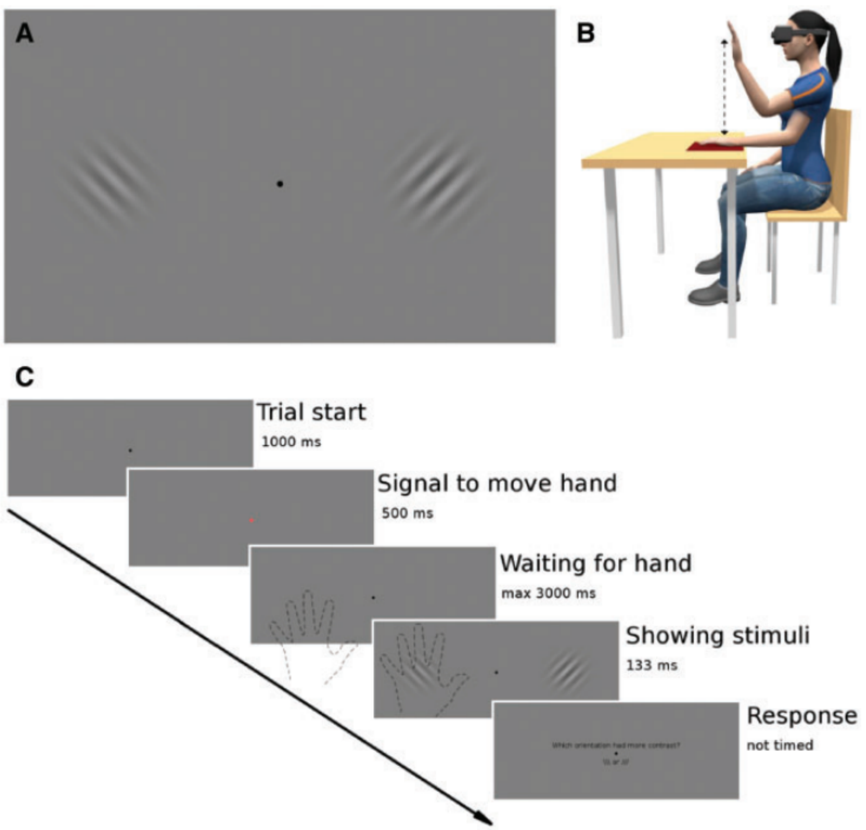

Figure 11. (A) Approximate example of the Gabor patches used as stimuli in the study. (B) Physical setup of the experiment, with both rest and raised hand positions and movement trajectory visualized. (C) General design for a single trial with time delays in milliseconds. The final panel reads: "Which orientation had more contrast?". Hand outlines and Gabor patch sizes on the figure are illustrative.

4.2. Results and conclusions

In data preprocessing we discarded trials where the hand movement was not within the allowed spatial constraints (56% from Experiment 1, 34% from Experiment 2). Also excluded were participants that showed a random answering pattern. The final sample consisted of 6 participants for Experiment 1 and 44 participants for experiment 2. Our results showed that self-generated movements influence visual sensitivity while leaving metacognitive accuracy intact. The contrast of the stimulus behind the invisible hand was indeed significantly reduced both in Experiment 1 ($F(1, 5) = 15.07$, $P = 0.01$) and Experiment 2 ([$F(1, 29) = 11.13$, $P = 0.002$). However, there was no significant difference in metacognitive accuracy between the experimental and control groups in Experiment 2 ($t(21.54) = 0.49$, $P = 0.63$, $BF = 0.35$), as measured by type II area under the receiver-operating curve (AROC) values. This method is commonly used to measure the ability to link subjective confidence to perceptual performance (Galvin et al, 2003; Fleming et al, 2010). The results suggest that participants in the experimental group were able to adjust their subjective confidence to their objective performance. In other words, if a low contrast stimulus was attenuated they could tell that the decision was easier— and vice versa. This is in line with previous work on metacognition (Kanai et al, 2010; Sherman et al, 2015). The results can be seen graphically on figure 12. Being able to tell that there is sensory attenuation occurring might be a crucial cue in recognizing an action as being internally triggered and not external.

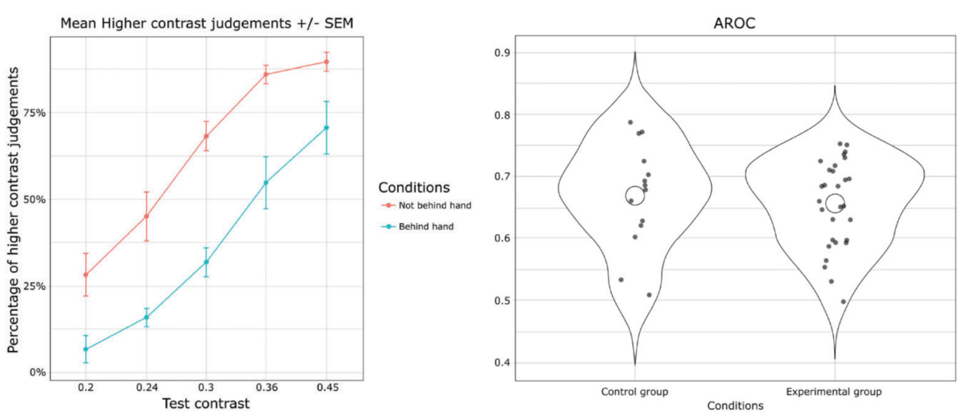

Figure 12. Main results of Experiments 1 and 2. Left panel - statistically significant difference in contrast judgements between participants within different conditions. Right panel - no significant difference in metacognitive abilities between groups as measured by type-II AROC values. The dots represent each participants metacognitive ability between values 0-1 and the violin plot also visualizes this same distribution.

The main results confirm those of publication II, but go a step further by indicating that self-generated movements do not only impact reaction times through attenuating attention, they also influence visual sensitivity directly. These findings thus offer further support for the active inference account, which posits sensory attenuation of visual processing due to sensory precision being reduced during movement (Friston 2010; Hohwy 2013; Clark 2015). Interestingly, as our results on metacognitive abilities showed, sensory attenuation appears to only influence first-order processes, e.g. direct perception.

Using this paradigm, one could further investigate the characteristics of visual attenuation caused by self-generated hand movements and answer fundamental questions about the computations that the brain is running. A big limitation of both publications II and III is the lack of eye gaze tracking. This means that we can not guarantee that the participants did not look away from the fixation point and directly towards the stimuli during the trials, confounding the results. To mitigate this, we used a proxy measure of general headset view direction to remove any suspicious trials and also designed the experiments so that looking away from the fixation point brings no benefit. The current approach can be further improved by employing the latest VR technology that provides accurate eye gaze tracking (e.g FOVE Inc, 2018; Tobii, 2018; HTC, 2019). But this all only applies when the VR research in general is conducted in a rigorous way and follows set guidelines.

5. GUIDELINES FOR VR RESEARCH (PUBLICATION IV)

Over the years we have gained experience with developing VR programs, experimenting with different hardware setups and study protocols. Everything that could technically go wrong, eventually did - from positional tracking anomalies to critical study design flaws. As the number of annually published articles in peer reviewed academic journals with "virtual reality" as a keyword has skyrocketed between 2016 and 2019 from around 6000 entries to over 12 000, according to EBSCO Information Services, this poses several questions. Are all these publications quality research that follow some general guidelines? Are there even any set guidelines? What are the known pitfalls of VR research to avoid? As VR holds a huge potential for basic science (Pan & Hamilton, 2018a; Miller et al, 2019) and therapeutic approaches (Rizzo & Koenig, 2017), we collected our own experiences and also conducted a literature review of the past few years up to 2020 in order to ensure reliable and valid VR guidelines for both ourselves and the wider community going forward (publication IV).

5.1. Reliability and validity issues

Over the years the term "immersive VR" has become such an umbrella concept as to render itself nearly useless. Under the same label we can find both cheap cardboard-based mobile phone enclosures and full-body backpack-PC powered systems. So two studies, both using "immersive VR" might arrive at different results because the setups provided are actually hardly comparable. During the writing of this thesis we made an attempt to solve this confusing situation by trying to introduce a novel classification of "quasi-immersive VR" to weed out the almost-but-not-quite-immersive VR systems from the immersive approaches, but in the end we were still overwhelmed by too many edge-case situations and decided to discontinue pursuing such a classification. A good current recommendation would be to really study the method sections of different VR papers to determine if the setups are even physically comparable to begin with, not all immersive VR systems provide the same level of presence (Sinesio et al, 2019).

Content wise, it is still debated whether it is worthwhile to increase visual realism of experimental VR experiences (Pan & Hamilton, 2018a; Rizzo & Koenig, 2017; Slater et al, 2020). This raises questions of ecological validity (Kulik, 2018) - are we really using VR to study what we claim we are studying? This also relates to the fundamental problem of always knowing that one is wearing a HMD and in a virtual world, referred to as the device-gap (Slater et al, 2020). This is at odds with the claim we laid out in the beginning of the thesis, framing the ability of participants to never forget the primary reality as a virtue. As this paradox currently seems unsolvable, it is useful to remember that all VR experiments take place in "dual realities" (Pan & Hamilton, 2018b).

In our short review (publication IV) we further pointed out potential problems regarding suspension of disbelief (Kulik, 2018), representative neural and behavioural responses (Harris et al, 2019), cybersickness (Kourtesis et al, 2019) and overall replicability (Lanier et al, 2019). We also highlighted the potential of novel reality-breaking “Virtual Unreality” experiments and different existing evaluation paradigms (Birckhead et al, 2019; Krohn et al, 2020; Tcha-Tokey et al, 2018).

5.2. Guidelines

Based on the issues discussed throughout this chapter and experienced while conducting this thesis work, we composed a list of concrete practical recommendations for designing and reporting VR experiments. Here we present a condensed version (figure 13), with more detailed descriptions available in publication IV.

Figure 13. A graphical summary of the guidelines for immersive VR research.

While publication IV is abundant with appropriate literature references of interest, here we would like to offer a more personal explanation to the eleven proposed guidelines. We take inspiration from our own past work and many VR projects developed while supervising dozens of students over the years. Hence, the references in the following paragraphs are often to the theses done by BSc and MSc students at the University of Tartu Computer Graphics and Virtual Reality lab.

Playing VR games (1) is essential for researcher just starting out in the field, as it can really show the full potential of the system when pushed to it's very limits in terms of graphical quality, interactions and optimization - features that novice developers are unable to achieve at first, but can be highly motivational aspirations. Additionally, it is useful to study background materials on the development side of the games, as most commercial programmers are keen to publicly share their tips and tricks on how certain implementations were built. One can also reverse-engineer some popular games to arrive at these principles (Ussanov, 2019). Another suggestion is to also play some really bad games - this will be a friendly lesson about how simply adding VR to a 2D experience will not automatically yield a usable end result (Kroon, 2019).

The problems surrounding dual realities (2) may seem daunting at first, but can actually be an exciting topic of study in itself. During the past four years the author has encountered several occasions where dual realities briefly gave way to a perfect simulation, where the person wearing the HMD forgot their primary reality. Once a friend was walking on a virtual ledge with a corresponding physical plank also present. But the real ledge was slightly shorter than the virtual, thus resulting in an unexpected loss of footing and one of the loudest screams the author has ever heard and a subsequent claim of a near-death experience by the friend. Another time, the author's beloved girlfriend was flying in a virtual simulator and got instructed to fly straight down into the ground. The loss of dual realities occurred about a second before expected virtual impact and a return to primary reality occurred slowly over the next hour. The humble claim here is that defeating dual realities is technically already possible, but (fortunately) still only transiently. With this capability come huge ethical questions, so care must be taken not to induce this experience accidentally. However, clever study designs can utilize the phenomenon of dual realities to explore tasks where the solution requires operating in both realities at once (Adamson, 2020).

Choosing the setup (3) is easy in theory, but practically speaking one might be stuck with some old hardware and has to make due for the time being. Since "immersive VR" can mean pretty much anything, researchers setting up a new lab should just keenly look at the method sections of some great contemporary VR papers and note what devices are being used. There is always a temptation to acquire the latest and greatest headset on the market, but many of the more exotic devices can be problematic, inflating one parameter (e.g. field of view) at the expense of another (e.g. frame rate). Even with a tight budget one can always set up the systems in different ways to open new research avenues or combine available technologies, for example by duck-taping a cheap peripheral to the headset (Kulu, 2016) or a 360 camera on a radio-controlled toy car (Traumann, 2017). If money is not an issue, there will still be constant trial-and-error learning about how new systems behave in different configurations related to the host PC (if there even is one) or with the physical world. Can the apparatus work in sunlight? In the presence of reflecting surfaces? What are the causes and solutions to some seemingly unexplainable technical failures? Most VR systems need a highly tuned environment to work properly and even then surprises can occur after thousands of hours of regular up-time, as the author can attest from being involved in three years of daily operations in one of the world's leading VR arcades Futuruum (SpringboardVR, 2018).

Designing with human-centric principles (4) is probably new to psychologists who are used to study subjects sitting quietly in darkened rooms, staring at the dim computer screen for hours. VR is more demanding in many ways - the whole field of view is obscured, the headset is somewhat heavy, the screen is right in front of the eyeballs and sometimes the participants have to actively navigate around and use controllers they have never seen before. Thus, any uncomfortable movement will become a burden quickly. Since most VR experiments are quite short, avoiding a user interface altogether and simply dropping the participant in the virtual world would be ideal. When thinking about user experience, all precautions should be taken to avoid cybersickness or accidental physical wall bumps. There are many design mistakes that even esteemed commercial developers still make (Vasser & Lomp, 2018). If nausea-inducing smooth locomotion can not be avoided, at least make it comfortably slow (Kivisik, 2016). The study subject trusts the experimenter and this trust should not be exploited. Not only will a dizzy subject ruin the results, but they will also make sure that none of their friends sign up for the experiments.

Beautifying (5) VR experiments has been a long-term dream of the author. It is then ironic that our own published experimental studies take place in a uniform grey void (publications II and III). It goes to show that ultimately the choice depends on the research question at hand. And even in the grey void, visuals play a crucial role (Koppel, 2017). However, while developing other VR projects related to more general research topics, raising the visual quality has both made the work widely more enjoyable to do and also teached a lot about the ins and outs of computer graphics, 3D modelling and performance optimizing (Magomedkerimov, 2016; Stafinjak, 2016). Labs employing students as developers should really think about challenging them to match the visual standard of the next experiment with for example that of Half-Life: Alyx (Valve Corporation). As VR experiments are usually confined to small virtual spaces, there is no need to construct large open-world scenarios, but effort can be put in designing a few truly immersive rooms. If the students succeed, this poses another problem - with VR developers in very high demand, these individuals may be drawn to the industry before even finishing their studies. Even the author of this dissertation was tempted to take part in several game jams and published a VR title on the digital marketplace Steam during doctoral studies. Thus it is an obvious challenge to keep talented developers in academia. It is not rare to find VR experiments published even today, where the visual quality seems severely outdated. We have been searching for explanations for this trend and the answers have included the inherent low fidelity needs of certain experiments and the general lack of funding to (re-)employ the good developers (Vasser, 2016). Fortunately real-time game engines such as Unity, Unreal or CryEngine have been progressing rapidly in recent years, making visually stunning content development easier by the day.

Gamifying (6) can be approached in two ways. Either by actually using a commercial title to do the experiments on, or by adding entertainment value to an existing experiment. It is true that for many people VR is in itself already very exciting. You can place a table in a virtual room and first-time VR users will probably marvel, step through and crawl under it. For others, the table must be interactive to hold attention. Ideally a VR experiment would be one where the study subject (i.e. player) has no idea that it is an experiment, rather it looks like an interesting game (Adamson, 2020; Kask, 2020). We started exploring such avenues with a virtual 4x4 floor of coloured tiles, where the player had to get from one corner of the floor to the other. Importantly, the player was shown a score that sometimes decreased while moving from place to place. The task was to get from start to finish without a penalty, thus discovering the hidden rules of our virtual universe. Even in this very simple game, one participant spent almost two hours on separate days to beat it.

Picking your participants (7) seems easier said than done as VR technology becomes more widespread. If one is looking for people with no prior experience with VR (Adamson, 2020), this population is declining rapidly. Another unreachable segment is the people who have previously tried some extremely nauseating VR ride and will shy away from any headsets for many years to come. And then there are people with extensive VR background - either developers or frequent VR arcade customers. This group of people are usually also not the best for studies that try to look at some general psychological phenomena. When advertising a VR study it is often quite easy to fill in the available slots. However, the question remains about the validity of such tight self-selection on the generalisability of the results. Perhaps a more ideal approach would be to design the experiments so that it could be conducted at home or at a local VR arcade, thus gaining access to a broader set of people.

Testing often and extensively (8) means that one should not leave testing to the very end of the development process and do it in a limited fashion. This will most likely result in an awkward and buggy experiment. Testing should already start when the basic “core loop” of the study is up and running, even with all the 3D assets still populated with place-holder boxes. Yes, this approach will “waste” some study subjects who can then probably not participate in the final data collection, but it can "save" the whole experiment, if time can not be squandered on developing the wrong things. By extensive testing we mean giving people ample time to play around and "break" the experience, thus giving valuable insights as to what may happen during the actual study. There might be some extraordinary talented individuals in this, such as fellow developers that will stop at nothing until your output log window is glowing red with errors. Be sure to also check the recorded data early on - there might be situations where problems detected late in the experiment will take up the better part of data cleaning (Tammsaar, 2017), if one is lucky.

Freezing updates (9) is mostly a reflex reminiscent of the early days of Oculus when every new runtime introduced with it utter destruction in terms of backwards compatibility. Modern versions of this problem can be more subtle, for example your game engine of choice renaming some critical internal functions, a mandatory middleware remapping the controller buttons or the whole operating system deciding to update to a broken build. Disconnecting from the web seems like the best option, but not always practical if the workstation has multiple users. One approach would be to ensure that the experiment is always available also as a stand-alone executable, if possible. This will eliminate at least some of the potential problems. Or, if one has a lot of capacity, it might be possible to try and stay up to date and really fix every problem as it emerges, always checking that the experiment has not fundamentally changed in the process. This is evident with updating the hardware, as swapping out the headset or graphics card in the middle of a long data collection period can introduce an unwanted situation where half of the sample will experience a lower resolution or less smooth version of the study. We have resisted such temptations half-way through one of our studies even when newer hardware would have been significantly better for the specific experiment (Luik, 2020).

Being real (10) is something that should never be forgotten. Yes, VR offers ways to significantly raise the ecological validity compared to classical highly artificial lab experiments, but it is still a long way from simulating reality, if it ever gets there at all. Researchers should not go along with the marketing hype of VR that claims "total immersion" or "escaping reality". Focus should be put on trying to tease out any possible confounding factor related to the results (or lack thereof): was it due to VR in general, a particular experience in VR or something completely different. This is especially true, when the study sample size is rather small (Gilden, 2017; Kitse, 2018).

Fostering replication (11) will probably be a crucial aspect for the field in the future. In the absence of this it will be difficult to build up significant cumulative knowledge. A case in point: due to the huge technological rift that occurred in around 2012 with the release of a commercial device that simultaneously offered a wide field of view, high frame rate and low latency, every academic paper released before that year is met with a certain skeptical scorn by the author of this thesis. Can I trust these results? Was it even proper VR? It is quite possible that in a decade from now, after another technological rift, some young researcher will feel the same about dissertations written in 2020. In order to lessen this worry, an interesting field of inquiry would be to replicate classic VR studies, but with modern technology - if the results still hold, we are standing on solid foundations. If not, then we are in for a VR replication crisis.

SUMMARY

We set out to experimentally test the active inference framework of brain functioning with novel virtual reality approaches. Along the way we had to learn to work with rapidly changing technologies, create specialized software, invent new study protocols and conduct trials with human participants. For the many years of research we have some interesting results to show: virtual reality can indeed be a useful tool to study psychology and neuroscience, the active inference framework holds up when explaining the attenuating effects of self-generated movements on visual attention and sensitivity, and great care must be taken to ensure reliable and valid psychological VR research.

Chapter 1 gave a broad overview of the intersection of neuroscience and virtual reality, focusing on the underlying logic of the active inference framework situated in the larger free energy principle and the technicalities of VR. Issues noted about VR development back in 2016 still stand in 2020 - while the technology has greatly matured and many old problems have been solved, others have stayed or new ones emerged.

Chapter 2 explained the rationale, design and development process behind the open-source toolbox VREX (publication I), as it still very much applies today - many researchers need a user-friendly approach to employ VR in psychological settings. Developing the software gave us much-needed experience in designing our next study protocols.

In chapter 3 we introduced a novel research paradigm to study active inference using VR and related technologies, namely hand-tracking devices. Our research question was simple - what is the mechanism behind the everyday experience of suppressed visual attenuation for self-generated movements? Our results (publication II) showed the statistically significant effect of longer reaction time when the presented stimuli appeared behind the participants invisible virtual hand. These results fit well with the active inference framework of brain functioning.

Chapter 4 explored the active inference account for attenuating visual perception in a more direct way, by assessing contrast judgements and secondary effects on metacognition (publication III). With a similar study design as outlined in chapter 3, but different stimuli, we again showed a strong effect of self-generated movements on attenuating visual perception. However, metacognitive abilities, thinking about thinking, were not affected by our paradigm.

With chapter 5, we shared some warnings and guidelines on continuing the important work of psychological VR experiments (publication IV). We asked tough questions from ourselves and the field about issues concerning reliability, validity and evaluating different paradigms. The proposed guidelines were aimed at defeating most of the common tropes that today's researchers might find themselves in.

In conclusion, our work in the intersection of computer science, psychology and virtual reality has given some new evidence for piecing together the larger

puzzle of which algorithms underlie biological intelligence. Needless to say that there are still many unsolved questions and neuroscience as a field has a long way to go towards a unified theory of brain functioning. Yet these results can already be applied also in many other domains, such as medicine, robotics and artificial intelligence.

Epilogue: from virtual to reality

Predictions are paramount to the brain, helping to avoid conditions that can ultimately threaten our survival. This is true for moment-to-moment existence, but can also be applied to longer time-scales. During the writing of the thesis and pondering our own future, this notion has gained quite a lot of weight for us.

It has been said that "experts who have extensive knowledge of a specific technology express much more guarded optimism [compared to lay people], if not pessimism, and generally provide a more balanced view" regarding the impacts of a given technology, the so called Huesemann's Law of Techno-Optimism (Huesemann & Huesemann, 2011). Given the critical health of the living world (IPCC, 2018; IPBES, 2019) and the foreseeable improbability of decoupling material growth from its environmental impacts (Parrique et al, 2019), we see urgent need to work with the general public on awareness raising about the possible negative impacts of modern technology, both environmental (e.g. Aavik, 2018) and psychological (e.g. Eensaar, 2018). Simply employing more technology to try and solve the many problems itself created by technology can only be a short-term techno-fix, as a more fundamental shift in world-view is now needed (Wiedmann et al, 2020).

We sincerely hope our work on active inference, more specifically the notion of not being able to perceive the consequences our own very real actions, also has value for the study of pro-environmental sustainable behaviour in the age of ecological breakdown and ecocide (Kaaronen, 2018; Kaaronen & Strelkovskii, 2020).

BIBLIOGRAPHY

Aavik, L. (2018). Futuruumi looja Madis Vasser: tehnoloogia pole alati ainult roosiline tulevik, mis meid päästab. *Postimees*, 24.10.2018 https://sobranna.elu24.ee/6436287/futuruumi-looja-madis-vasser-tehnoloogia-pole-alati- ainult-roosiline-tulevik-mis-meid-paastab (accessed on 26.07.2020).

Adams, R. A., Shipp, S., & Friston, K. J. (2013). Predictions not commands: active inference in the motor system. *Brain Structure and Function*, 218(3), 611–643.

Adamson, A. (2020). Exploring Insight Using Virtual Reality. Bachelor's Thesis, University of Tartu.

Allen, M. & Tsakiris, M. (2018) The body as first prior: Interoceptive predictive processing and the primacy of self-models. In: Tsakiris, M., De Preester, H. (eds). *The interoceptive mind: from homeostasis to awareness*. Oxford University Press, Oxford, pp 27–45

Aru, J., Suzuki, M. & Larkum, M. (2020). Cellular mechanisms of conscious processing. *Trends in Cognitive Sciences.*

Azevedo, F. A. C., Carvalho, L. R. B., Grinberg, L. T., Farfel, J. M., Ferretti, R. E. L., Leite, R. E. P., . . . Herculano-Houzel, S. (2009). Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain. *The Journal of Comparative Neurology*, 513(5), 532–541. doi:10.1002/cne.21974

Bays, P. M., Flanagan, J. R., & Wolpert, D. M. (2006). Attenuation of Self-Generated Tactile Sensations Is Predictive, not Postdictive. *PLoS Biology*, 4(2), e28. doi:10.1371/journal.pbio.0040028

Bear, M. F., Connors, B. W., & Paradiso, M. A. (2016). *Neuroscience: Exploring the Brain. 4th edition*. Wolters Kluwer.

Berkes, P., Orban, G., Lengyel, M., & Fiser, J. (2011). Spontaneous Cortical Activity Reveals Hallmarks of an Optimal Internal Model of the Environment. *Science*, 331(6013), 83–87. doi:10.1126/science.1195870

Birckhead, B., Khalil, C., Liu, X., Conovitz, S., Rizzo, A., Danovitch, I., ... & Spiegel, B. (2019). Recommendations for methodology of virtual reality clinical trials in health care by an international working group: iterative study. *JMIR mental health*, 6(1), e11973.

Blakemore, S. J., Wolpert, D. M., & Frith, C. D. (1998). Central cancellation of self-produced tickle sensation. *Nature Neuroscience*, 1(7), 635–640. doi:10.1038/2870

Brookes, J., Warburton, M., Alghadier, M., Mon-Williams, M., & Mushtaq, F. (2019). Studying human behavior with virtual reality: The Unity Experiment Framework. *Behavior Research Methods*. doi:10.3758/s13428-019-01242-0

Brown, H., Adams, R. A., Parees, I., Edwards, M., & Friston, K. (2013). Active inference, sensory attenuation and illusions. *Cognitive Processing*, 14(4), 411–427. doi:10.1007/s10339-013-0571-3

Brown, H., & Friston, K. (2012). Dynamic causal modelling of precision and synaptic gain in visual perception—an EEG study. *Neuroimage*, 63(1), 223-231.

Bogacz, R. (2017). A tutorial on the free-energy framework for modelling perception and learning. *Journal of Mathematical Psychology*, 76, 198–211. doi:10.1016/j.jmp.2015.11.003

Burr, D. C., Holt, J., Johnstone, J. R., & Ross, J. (1982). Selective depression of motion sensitivity during saccades. *The Journal of Physiology*, 333(1), 1–15. doi:10.1113/jphysiol.1982.sp014434

Carrasco, M., Ling, S., & Read, S. (2004). Attention alters appearance. *Nature neuroscience*, 7(3), 308-313.

Clark, A. (2015). *Embodied prediction*. Open MIND. Frankfurt am Main: MIND Group.

Clark, A. (2016). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press.

Commins, S., Duffin, J., Chaves, K., Leahy, D., Corcoran, K., Caffrey, M., ... Thornberry, C. (2019). NavWell: A simplified virtual-reality platform for spatial navigation and memory experiments. *Behavior Research Methods*. doi:10.3758/s13428-019-01310-5

Corcoran, A. W., Pezzulo, G., & Hohwy, J. (2020). From allostatic agents to counterfactual cognisers: active inference, biological regulation, and the origins of cognition. *Biology & Philosophy*, 35(3). doi:10.1007/s10539-020-09746-2

Corlett, P. R., Honey, G. D., & Fletcher, P. C. (2016). Prediction error, ketamine and psychosis: An updated model. *Journal of psychopharmacology*, 30(11), 1145-1155.

Eagleman, D. (2015). *The Brain: The story of you*. Canongate Books.

Eensaar, L. (2018). Neuroteadlane Jaan Aru: nutiseade ei tee teid oma töös paremaks. *Novaator*, 31.07.2018. https://novaator.err.ee/850321/neuroteadlane-jaan- aru-nutiseade-ei-tee-teid-oma-toos-paremaks (accessed on 26.07.2020)

Einevoll, G. T., Destexhe, A., Diesmann, M., Grün, S., Jirsa, V., de Kamps, M. & Schürmann, F. (2019). The scientific case for brain simulations. *Neuron*, 102(4), 735-744.

Fleming, S. M., Weil, R. S., Nagy, Z., Dolan, R. J., & Rees, G. (2010). Relating Introspective Accuracy to Individual Differences in Brain Structure. *Science*, 329(5998), 1541–1543. doi:10.1126/science.1191883

Friston, K. (2005). A theory of cortical responses. *Philosophical transactions of the Royal Society B: Biological Sciences*, 360(1456), 815–836.

Friston, K. (2009) The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences*, 13(7):293–301

Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127.

Friston, K., Thornton, C., & Clark, A. (2012). Free-Energy Minimization and the Dark-Room Problem. *Frontiers in Psychology*, 3. doi:10.3389/fpsyg.2012.00130

Friston, K. (2018). Active Inference and Cognitive Consistency. *Psychological Inquiry*, 29(2), 67–73. https://doi.org/10.1080/1047840X.2018.1480693

Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active Inference: A Process Theory. *Neural Computation*, 29(1), 1–49. https://doi.org/10.1162/NECO_a_00912

Friston, K. J., Lawson, R., & Frith, C. D. (2013). On hyperpriors and hypopriors: comment on Pellicano and Burr. *Trends in Cognitive Sciences*, 17(1), 1.

Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T. & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive Neuroscience*, 6(4):187–224

Friston, K. & Stephan, K. (2007). Free-energy and the brain. *Synthese*, 159(3):417–458

FOVE Inc. (2018) Eye Tracking VR dev kit. https://www.getfove.com (accessed 14.07.2020).

Galvin, S. J., Podd, J. V., Drga, V., & Whitmore, J. (2003). Type 2 tasks in the theory of signal detectability: Discrimination between correct and incorrect decisions. *Psychonomic Bulletin & Review*, 10(4), 843–876. doi:10.3758/bf03196546

Gilden, T. (2017). Paralysed by Surprise: Testing the Active Inference Theory with VR. Master's Thesis, University of Tartu.

Gregory, R. L. (1980). Perceptions as hypotheses. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, 290(1038), 181–197.

Harris, D. J., Buckingham, G., Wilson, M. R., & Vine, S. J. (2019). Virtually the same? How impaired sensory information in virtual reality may disrupt vision for action. *Experimental brain research*, 237(11), 2761-2766.

Hassabis, D., Kumaran, D., Summerfield, C., & Botvinick, M. (2017). Neuroscience-Inspired Artificial Intelligence. *Neuron*, 95(2), 245–258. doi:10.1016/j.neuron.2017.06.011

Hawkins, J., Ahmad, S. & Cui, Y. (2017). A Theory of How Columns in the Neocortex Enable Learning the Structure of the World. *Frontiers in Neural Circuits*, 11(2017). DOI:10.3389/fncir.2017.00081

Helmholtz, H. v. (1867). *Handbuch der Physiologishen Optik*. Leipzig: Leopold Voss.

Hinton, G. E., Dayan, P., Frey, B. J., & Neal, R. M. (1995). The" wake-sleep" algorithm for unsupervised neural networks. *Science*, 268(5214), 1158-1161.

Hohwy, J. (2013). *The predictive mind*. Oxford University Press

HTC. (2019). Setting a New Standard for Enterprise Virtual Reality. https://blog.vive.com/us/2019/06/06/setting-new-standard-enterprise-virtual-reality-vive- pro-eye-now-available-north-america/ (accessed 14.07.2020)

Huesemann, M., & Huesemann, J. (2011). *Techno-fix: why technology won't save us or the environment*. New Society Publishers.

IPBES (2019). Bongaarts, J. Summary for policymakers of the global assessment report on biodiversity and ecosystem services of the Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services. *Population and Development Review*, 45(3), 680-681.

IPCC (2018). *Global Warming of 1.5 OC: An IPCC Special Report on the Impacts of Global Warming of 1.5° C Above Pre-industrial Levels and Related Global Greenhouse Gas Emission Pathways, in the Context of Strengthening the Global Response to the Threat of Climate Change, Sustainable Development, and Efforts to Eradicate Poverty (p. 32).* V. Masson-Delmotte (Ed.). Geneva, Switzerland: World Meteorological Organization.

Jevons, W. S. (1871). The Power of Numerical Discrimination. *Nature*, 3(67), 281–282. doi:10.1038/003281a0

Jonas, E., & Kording, K. P. (2017). Could a neuroscientist understand a microprocessor? *PLoS computational biology*, 13(1).

Juravle, G., Deubel, H., Tan, H. Z., & Spence, C. (2010). Changes in tactile sensitivity over the time-course of a goal-directed movement. *Behavioural Brain Research*, 208(2), 391–401. doi:10.1016/j.bbr.2009.12.009

Juravle, G., & Spence, C. (2011). Juggling reveals a decisional component to tactile suppression. *Experimental Brain Research*, 213(1), 87–97. doi:10.1007/s00221-011-2780-2

Kaaronen, R. O. (2018). A theory of predictive dissonance: Predictive processing presents a new take on cognitive dissonance. *Frontiers in psychology*, 9, 2218.

Kaaronen, R. O., & Strelkovskii, N. (2020). Cultural Evolution of Sustainable Behaviors: Pro-environmental Tipping Points in an Agent-Based Model. *One Earth*, 2(1), 85–97. doi:10.1016/j.oneear.2020.01.003

Kanai, R., Walsh, V., & Tseng, C. H. (2010). Subjective discriminability of invisibility: a framework for distinguishing perceptual and attentional failures of awareness. *Consciousness and cognition*, 19(4), 1045-1057.

Kask, K. J. (2020). Innovation in the Laboratory: Developing a New Experimental Paradigm in Virtual Reality, which Requires an Innovative Solution. Bachelor's Thesis, University of Tartu.

Kiefer, A., & Hohwy, J. (2018). Content and misrepresentation in hierarchical generative models. *Synthese*, 195, 2387–2415.

Kiefer, A., & Hohwy, J. (2019). Representation in the prediction error minimization framework. In J. Symons, P. Calvo, & S. Robins (Eds.), *Routledge handbook to the philosophy of psychology, 2nd ed.* Oxford, UK: Routledge.

Kilteni, K., Normand, J.-M., Sanchez-Vives, M. V., & Slater, M. (2012). Extending Body Space in Immersive Virtual Reality: A Very Long Arm Illusion. *PLoS ONE*, 7(7), e40867. doi:10.1371/journal.pone.0040867

Kitse, K. (2018). Probing the internal world model of the brain. Bachelor's Thesis, University of Tartu.

Kivisik, T. (2016). What Makes a Good Location in the Method of Loci:

Meaningfulness of Locations and Environmental Segregation Studied with Virtual Reality. Master's Thesis, University of Tartu.
Klein, C. (2016). What do predictive coders want? *Synthese*, 195(6), 2541–2557. doi:10.1007/s11229-016-1250-6
Kriegeskorte, N., & Douglas, P. K. (2018). Cognitive computational neuroscience. *Nature Neuroscience*. doi:10.1038/s41593-018-0210-5
Krohn, S., Tromp, J., Quinque, E. M., Belger, J., Klotzsche, F., Rekers, S., ... & Villringer, A. (2020). Multidimensional Evaluation of Virtual Reality Paradigms in Clinical Neuropsychology: Application of the VR-Check Framework. *Journal of medical Internet research*, 22(4), e16724.
Kroon, R. (2019). Application for Psychophysics Experiments in Virtual Reality. Bachelor's Thesis, University of Tartu.
Koppel, T (2017). Reduced Contrast in the Area of the Visual Field Where the Own Hand Is Moving: A Virtual Reality Study. Bachelor's Thesis, University of Tartu.
Koriat, A. (2007). Metacognition and Consciousness. In: Zelazo. P. D., Moscovitch, M. and Thompson, E. (eds), *The Cambridge handbook of consciousness*. Cambridge, UK: Cambridge University Press,2007, 289–325.
Kourtesis, P., Collina, S., Doumas, L. A. A., & MacPherson, S. E. (2019). Technological competence is a precondition for effective implementation of virtual reality head mounted displays in human neuroscience: a technological review and meta-analysis. *Frontiers in Human Neuroscience*, 13, 342.
Kulik, A. (2018). Virtually the ultimate research lab. *British Journal of Psychology*, 109(3), 434-436.
Kulu, S. (2016). The Human Octopus: controlling supernumerary hands with the help of virtual reality. Bachelor's Thesis, University of Tartu.
LaBerge, S. (2009). *Lucid Dreaming: A Concise Guide to Awakening in Your Dreams and in Your Life*. Sounds True, Boulder, Colorado.
Lanier, M., Waddell, T. F., Elson, M., Tamul, D. J., Ivory, J. D., & Przybylski, A. (2019). Virtual reality check: Statistical power, reported results, and the validity of research on the psychology of virtual reality and immersive environments. *Computers in Human Behavior*, 100, 70-78.
Lawson, R. P., Rees, G., & Friston, K. J. (2014). An aberrant precision account of autism. *Frontiers in Human Neuroscience*, 8, 302.
Lazebnik, Y. (2002). Can a biologist fix a radio?—Or, what I learned while studying apoptosis. *Cancer Cell*, 2(3), p 179-182
Linson, A., Clark, A., Ramamoorthy, S., & Friston, K. (2018). The active inference approach to ecological perception: General information dynamics for natural and artificial embodied cognition. *Frontiers Robotics AI*, 5(MAR). doi.org/10.3389/frobt.2018.00021
Luik, S. A. (2020). Study of the reliability of eye-witness testimony in virtual environment. Bachelor's Thesis, University of Tartu Institute of Law.

Machado, M. L., Lefèvre, N., Philoxene, B., Le Gall, A., Madeleine, S., Fleury, P., … Besnard, S. (2019). New software dedicated to virtual mazes for human cognitive investigations. *Journal of Neuroscience Methods*, 327, 108388. doi:10.1016/j.jneumeth.2019.108388

Magomedkerimov, M. (2016). Extending the VREX toolbox for studying false memory in virtual reality. Bachelor's Thesis, University of Tartu.

Miller, L. C., Shaikh, S. J., Jeong, D. C., Wang, L., Gillig, T. K., Godoy, C. G., … Read, S. J. (2019). Causal Inference in Generalizable Environments: Systematic Representative Design. *Psychological Inquiry*, 30(4), 173–202. doi:10.1080/1047840x.2019.1693866

Milgram, P., Takemura, H., Utsumi, A. & Kishino, F. (1994). Augmented Reality: A class of displays on the reality-virtuality continuum. *Proceedings of Telemanipulator and Telepresence Technologies*. pp. 2351–34.

Neisser, U. (1967). *Cognitive psychology*. Appleton-Century-Crofts.

Oizumi, M., Albantakis, L., & Tononi, G (2014). From the Phenomenology to the Mechanisms of Consciousness: Integrated Information Theory 3.0. *PLoS Computational Biology*. 10 (5): e1003588. doi:10.1371/journal.pcbi.1003588.

Olson, J. L., Krum, D. M., Suma, E. A. & Bolas, M. (2011). A design for a smartphone-based head mounted display. *2011 IEEE Virtual Reality Conference*, Singapore, pp. 233-234. doi: 10.1109/VR.2011.5759484

Palmer, C. J., Lawson, R. P., & Hohwy, J. (2017). Bayesian approaches to autism: Towards volatility, action, and behavior. *Psychological bulletin*, 143(5), 521.

Palmer, C., Zapparoli, L., & Kilner, J. M. (2016). A New Framework to Explain Sensorimotor Beta Oscillations. *Trends in Cognitive Sciences*, 20(5), 321–323. doi:10.1016/j.tics.2016.03.007

Pan, X., & Hamilton, A. F. D. C. (2018a). Why and how to use virtual reality to study human social interaction: The challenges of exploring a new research landscape. *British Journal of Psychology*, 109(3), 395-417.

Pan, X., & Hamilton, A. F. D. C. (2018b). Understanding dual realities and more in VR. *British Journal of Psychology*, 109(3), 437.

Parr, T., & Friston, K. J. (2019). Generalised free energy and active inference. *Biological Cybernetics*, 113(5–6), 495–513. https://doi.org/10.1007/s00422-019-00805-w

Parrique, T., Barth, J., Briens, F., Kerschner, C., Kraus-Polk, A., Kuokkanen, A., & Spangenberg, J. H. (2019). *Decoupling debunked. Evidence and arguments against green growth as a sole strategy for sustainability*. A study edited by the European Environmental Bureau.

Pellicano, E., & Burr, D. (2012). When the world becomes 'too real': A Bayesian explanation of autistic perception. *Trends in Cognitive Sciences*, 16(10), 504–510.

Pillai, J. S., Schmidt, C., & Richir, S. (2013). Achieving presence through evoked reality. *Frontiers in psychology*, 4, 86.

Pitkow, X. & Angelaki, D. E. (2017). Inference in the Brain: Statistics Flowing in Redundant Population Codes. *Neuron*, Volume 94, Issue 5. doi:10.1016/j.neuron.2017.05.028.

Powers, A. R., Mathys, C., & Corlett, P. R. (2017). Pavlovian conditioning–induced hallucinations result from overweighting of perceptual priors. *Science*, 357(6351), 596–600.

Rahnev, D., Maniscalco, B., Graves, T., Huang, E., de Lange, F. P., & Lau, H. (2011). Attention induces conservative subjective biases in visual perception. *Nature Neuroscience*, 14(12), 1513–1515. doi:10.1038/nn.2948

Ramstead, M. J. D., Kirchhoff, M. D., & Friston, K. J. (2019). A tale of two densities: active inference is enactive inference. *Adaptive Behaviour*. https://doi.org/10.1177/1059712319862774

Rizzo, A., & Koenig, S. T. (2017). Is clinical virtual reality ready for primetime? *Neuropsychology*, 31(8), 877.

Rodríguez, M. F., Lazo, M., Cifuentes, M. V., D'Amato, J. P., Serodio, M., Ballarini, F., ... & García Bauza, C. (2017). Immersive Platform for Neuroscience Experimental Studies. *XXIII Congreso Argentino de Ciencias de la Computación*.

Safron, A. (2020). An Integrated World Modeling Theory (IWMT) of Consciousness: Combining Integrated Information and Global Neuronal Workspace Theories With the Free Energy Principle and Active Inference Framework; Toward Solving the Hard Problem and Characterizing Agentic Causation. *Frontiers in Artificial Intelligence*, 3. doi:10.3389/frai.2020.00030

Scarfe, P., & Glennerster, A. (2015). Using high-fidelity virtual reality to study perception in freely moving observers. *Journal of Vision*, 15(9), 3. doi:10.1167/15.9.3

Schettler, A., Raja, V., & Anderson, M. L. (2019). The Embodiment of Objects: Review, Analysis, and Future Directions. *Frontiers in Neuroscience*, 13. doi:10.3389/fnins.2019.01332

Schmack, K., de Castro, A. G. C., Rothkirch, M., Sekutowicz, M., Rössler, H., Haynes, J. D., ... & Sterzer, P. (2013). Delusions and the role of beliefs in perceptual inference. *Journal of Neuroscience*, 33(34), 13701–13712.

Schutte, N. S., & Stilinović, E. J. (2017). Facilitating empathy through virtual reality. *Motivation and emotion*, 41(6), 708-712.

Seth, A. K. (2013). Interoceptive inference, emotion, and the embodied self. *Trends in cognitive sciences*, 17(11), 565-573.

Seth, A. K., Millidge, B., Buckley, C. L., & Tschantz, A. (2020). Curious Inferences: Reply to Sun and Firestone on the Dark Room Problem. *Trends in Cognitive Sciences*. doi:10.1016/j.tics.2020.05.011

Sherman, M. T., Seth, A. K., Barrett, A. B., & Kanai, R. (2015). Prior expectations facilitate metacognition for perceptual decision. *Consciousness and cognition*, 35, 53-65.

Shinoda, H., Hayhoe, M. M., & Shrivastava, A. (2001). What controls attention in natural environments? *Vision Research*, 41(25-26), 3535–3545. doi:10.1016/s0042-6989(01)00199-7

Simons, D. J., & Rensink, R. A. (2005). Change blindness: Past, present, and future. *Trends in cognitive sciences*, 9(1), 16-20.
Sims, A. (2016). A problem of scope for the free energy principle as a theory of cognition. *Philosophical Psychology*, 29(7), 967–980.
Sinesio, F., Moneta, E., Porcherot, C., Abbà, S., Dreyfuss, L., Guillamet, K., ... & McEwan, J. A. (2019). Do immersive techniques help to capture consumer reality? *Food Quality and Preference*, 77, 123-134.
Slater, M., Gonzalez-Liencres, C., Haggard, P., Vinkers, C., Gregory-Clarke, R., Jelley, S., Silver, J. (2020). The Ethics of Realism in Virtual and Augmented Reality. *Frontiers in Virtual Reality*, 1. doi:10.3389/frvir.2020.00001
Slater, M., Lotto, B., Arnold, M. M., & Sanchez-Vives, M. V. (2009a). How we experience immersive virtual environments: the concept of presence and its measurement. *Anuario de Psicologia*, 40, 193–210.
Slater, M., Khanna, P., Mortensen, J., & Yu, I. (2009b). Visual realism enhances realistic response in an immersive virtual environment. *IEEE computer graphics and applications*, 29(3), 76-84.
Sperry, R.W. (1950). Neural basis of the spontaneous optokinetic response produced by visual inversion. *Journal of Comparative and Physiological Psychology*, 43, 482–9.
SpringboardVR. (2018). VR Arcade Open Business Plan: Customer Experience. https://medium.com/@springboardVR/vr-arcade-open-business-plan-customer-experien ce-e747b35d10fb. Accessed 21.01.2020
Stafinjak, V. (2016). Expanding the functionalities of VREX. Bachelor's Thesis, University of Tartu.
Sun, Z., & Firestone, C. (2020a). The Dark Room Problem. *Trends in Cognitive Sciences*. doi:10.1016/j.tics.2020.02.006
Sun, Z., & Firestone, C. (2020b). Optimism and Pessimism in the Predictive Brain. *Trends in Cognitive Sciences*. doi:10.1016/j.tics.2020.06.001
Sutherland, I. E. (1965). The Ultimate Display. *Proceedings of International Federation for Information Processing* 65, vol 2, pp. 506-50.
Sutherland, I. E. (1968). A head-mounted three dimensional display. *Proceedings of AFIPS* 68, pp. 757-764
Tammsaar, A. W. (2017). The Induction and Objective Measurement of Illusory Roll-Axis Vection in Virtual Reality. Bachelor's Thesis, University of Tartu.
Tcha-Tokey, K., Christmann, O., Loup-Escande, E., Loup, G., & Richir, S. (2018). Towards a model of user experience in immersive virtual environments. *Advances in Human-Computer Interaction*, 2018.
Teufel, C., Subramaniam, N., Dobler, V., Perez, J., Finnemann, J., Mehta, P. R., ... & Fletcher, P. C. (2015). Shift toward prior knowledge confers a perceptual advantage in early psychosis and psychosis-prone healthy individuals. *Proceedings of the National Academy of Sciences*, 112(43), 13401–13406.
Tobii AB. (2018). Tobii Pro VR Integration. https://www.tobiipro.com/product-listing/vr-integration/ (accessed 14.07.2020).

Traumann, A. (2017). Semi-Automatic Method for Creating Virtual Reality Environments. Master's Thesis, University of Tartu.
Tulver, K., Aru, J., Rutiku, R., & Bachmann, T. (2019). Individual differences in the effects of priors on perception: a multi-paradigm approach. *Cognition*, 187, 167-177.
Ussanov, S. (2019). 3D user inferface design. Bachelor's Thesis, University of Tartu.
Van de Cruys, S., Evers, K., Van der Hallen, R., Van Eylen, L., Boets, B., de-Wit, L., & Wagemans, J. (2014). Precise minds in uncertain worlds: Predictive coding in autism. *Psychological review*, 121(4), 649.
Van de Cruys, S., Friston, K. J., & Clark, A. (2020). Controlled Optimism: Reply to Sun and Firestone on the Dark Room Problem. *Trends in Cognitive Sciences*. doi:10.1016/j.tics.2020.05.012
Varjo Technologies (n.d). https://varjo.com/ Accessed 14.07.2020.
Vasser, M. (2016). Blog post on VECTOR workshop in Tübingen. https://cgvr.cs.ut.ee/wp/index.php/recap-vector-workshop-2016-in-tubingen/. Accessed 21.01.2020
Vasser, M. (2018). Virtual Reality [Lecture video]. *Introduction to Speciality (LTAT.03.002)*. Retrieved from https://www.uttv.ee/naita?id=27532
Vasser, M & Lomp, K. (2018). VR arcades: what works, what breaks. Presentation at GameDev Days 2018. https://www.youtube.com/watch?v=Vf0RJECFK1M (Accessed 21.01.2020)
Ward, J. (2019). Individual differences in sensory sensitivity: A synthesizing framework and evidence from normal variation and developmental conditions. *Cognitive Neuroscience*, 10(3), 139–157.
Wiedmann, T., Lenzen, M., Keyßer, L. T., & Steinberger, J. K. (2020). Scientists' warning on affluence. *Nature communications*, 11(1), 1-10.
Wiener, J. M., Carroll, D., Moeller, S., Bibi, I., Ivanova, D., Allen, P., & Wolbers, T. (2019). A novel virtual-reality-based route-learning test suite: Assessing the effects of cognitive aging on navigation. *Behavior Research Methods*. doi:10.3758/s13428-019-01264-8
Wilimzig, C., Tsuchiya, N., Fahle, M., Einhäuser, W., & Koch, C. (2008). Spatial attention increases performance but not subjective confidence in a discrimination task. *Journal of Vision*, 8(5), 7. doi:10.1167/8.5.7

ACKNOWLEDGEMENTS

I owe my greatest gratitude to my supervisor and mentor Jaan Aru for replying to my every problem with “this is science” and then helping to thoroughly resolve the issue at hand. Who knew that one e-mail sent in the beginning of my bachelor studies would go this far. My closest co-authors over the years also deserve special credit: Markus Kängsepp for tireless debugging in the early mornings, Krisjan-Julius Laak for the experiments in the late evenings and Laurène Vuillaume for introducing me to more international work. The final version of the thesis also benefited from the insightful feedback from reviewer Dr Meelis Kull and opponents Dr Jakub Limanowski and Dr Michael Gaebler.

I also thank the Institute of Computer Science of the University of Tartu, in particular Raul Vicente, for boldly supporting me and the CGVR lab. Building up the lab from a small cramped basement to the cutting-edge facility that it is today has been an unforgettable journey.

I also could not have finished this work without the support of my inspiring friends, family and life companion Kaisa, the Estonian VR community EEVR, the pioneers at Futuruum VR and fellow geeks at the CGVR and Computational Neuroscience labs.

My work has financially been supported by The University of Tartu Foundation Estonian terminology stipend (2017) and European Union European Regional Development Fund smart specialization stipend (2017-2020).

SISUKOKKUVÕTE

Arvutusliku ajuteooria testimine virtuaalreaalsuse abil

Käesoleva töö eesmärgiks oli eksperimentaalselt testida aktiivse järeldamise raamistikku, rakendades selleks uudseid virtuaalreaalsuse (VR) lahendusi. Aktiivse järeldamise teooria püüab seletada aju toimimise arvutuslikke põhimõtteid, käsitledes aju kui üldist ennustusmasinat. Uuringute käigus õppisime töötama kiirelt muutuva tehnoloogiaga, lõime spetsialiseeritud tarkvara, arendasime välja uudse katseparadigma ning viisime läbi inimkatseid. Aastatepikkuse töö tulemusena on meil näidata mitmeid huvitavaid tulemusi. Selgus, et virtuaalreaalsus saab tõesti olla kasulik uuringutööriist nii psühholoogias kui neuroteaduses; aktiivse järeldamise raamistik suudab seletada aju enda poolt juhitud kehaliigutuste visuaalsete tagajärgede summutamist; ning psühholoogiliste VR uuringute puhul tuleb hoolikalt jälgida, et tagatud oleks nii reliaablus kui valiidsus.

Peatükk 1 andis üldise ülevaate neuroteaduse ja VR-i kokkupuutepunktidest, keskendudes nii aktiivse järeldamise alusprintsiipide tutvustamisele kui ka selle paiknemisele laiemas vaba energia teoorias. Lisaks tutvustasime üksikasjalikumalt VRi tehnilist tausta. Mitmed probleemid, mis valitsesid VR tarkvaraarenduse juures aastal 2016 on jätkuvalt relevantsed ka aastal 2020 - kuigi tehnoloogia on vahepeal oluliselt küpsenud ning paljud endised kitsaskohad on lahenenud, siis teised püsivad ning kolmandad on tulnud juurde.

Peatükk 2 põhjendab vabavaralise tööriistakasti VREX vajalikkust, disaini ja arendustööd (publikatsioon I). Valminud programm on relevantne ka täna - paljud uurijad vajavad jätkuvalt kasutajasõbralikku ja võimekat viisi VR-i kasutamiseks psühholoogiliste uuringute läbiviimisel. Taolise tarkvara arendamine andis meile väärtuslikku kogemust järgmiste katseprotokollide disainimiseks.

Peatükis 3 tutvustasime uudset katseparadigmat, mille abil VR-i ja seotud tehnoloogiaid rakendades uurida aktiivse järeldamise protsesse. Meie uurimisküsimus oli lihtne - mis mehhanism tingib meie igapäevaelu kogemuse, kus iseenda poolt tekitatud kehaliigutused on meie nägemisest justkui alla surutud? Uuringu andmed (publikatsioon II) näitasid statistiliselt oluliselt pikemat reaktsiooniaega juhul, kui katsestiimul ilmus täpselt katseisiku liikuva nähtamatu virtuaalse käe taha. Taolised tulemused sobivad kokku aktiivse järeldamise raamistiku ennustustega.

Peatükk 4 avas aktiivse järeldamise põhimõtete kontrollimist otsesemate meetoditega, mõõtes katseisikute hinnanguid stiimulite tajutud kontrastile ning uurides ka teiseseid efekte metakognitsioonile (publikatsioon III). Kasutades eelmises peatükis kirjeldatud katseprotokolli analoogi, kuid uusi stiimuleid, näitasime taaskord enese poolt kontrollitud liigutuste summutavat mõju visuaalsele tajule. Samas selgus, et metakognitisoon ehk inimeste enesekindlus oma enda hinnangute täpsuse osas ei näidanud erinevate katsetingimuste vahel erisusi.

Peatükis 5 jagasime oma õppetunde ja juhiseid, kuidas jätkata psühholoogiliste

VR eksperimentide läbiviimist ka edaspidi. Püstitasime olulisi küsimusi nii iseendale kui laiemale valdkonnale seoses reliaabluse, valiidsuse ning erinevate katseparadigmadega. Meie poolt välja pakutud juhtnöörid (publikatsioon IV) on eesmärgiga ületada enamuse probleemkohti, millega tänased VR uurijad silmitsi võivad seista.

Kokkuvõtvalt on meie töö, mis seisab informaatika, psühholoogia ning VR-i ristumiskohal, andnud uusi teadmiskildusid, mille abil jätkata bioloogilise intelligentsuse arvutuslike põhimõtete suure pildi kokku panemist. Tööd selles valdkonnas on kindlasti veel palju, et jõuda kõikehõlmava teooriani. Saadud tulemusi saab aga juba praktiliselt kasutada mitmetes eri valdkondades, nagu näiteks meditsiin, robootika ja tehisintellekt.

PUBLICATIONS

CURRICULUM VITAE

Personal data

Name:	Madis Vasser
Date of birth:	18.07.1988
Citizenship:	Estonian
Contact:	madis.vasser@ut.ee

Education

2013–2015	University of Tartu, Institute of Psychology, MA.
2010–2013	University of Tartu, Institute of Psychology, BA.
1995–2007	Kunda Gymnasium.

Employment

2018–	Estonian Green Movement, advocacy expert
2017–2020	Futuruum VR OÜ, founder, CTO
2016–2020	VirtualFX OÜ, developer
2016–2017	VR Lab OÜ, founder (exited)
2014–2016	Psühhobuss OÜ, founder, CEO
2014–2016	University of Tartu, Institute of Psychology, TA
2013–2014	University of Tartu Student Council, head of development

Scientific work

Main fields of interest:

- Attention
- Consciousness
- Virtual Reality

ELULOOKIRJELDUS

Isikuandmed

Nimi:	Madis Vasser
Sünniaeg:	18.07.1988
Kodakondsus:	Estonian
Kontakt:	madis.vasser@ut.ee

Haridus

2013–2015	Tartu Ülikool, Psühholoogia instituut, Magistrikraad.
2010–2013	Tartu Ülikool, Psühholoogia instituut, Bakalaureusekraad.
1995–2007	Kunda Ühisgümnaasium, Üldkeskharidus.

Teenistuskäik

2018–	MTÜ Eesti Roheline Liikumine, huvikaitse ekspert
2017–2020	Futuruum VR OÜ, asutaja ja tehnoloogiajuht
2016–2020	VirtualFX OÜ, arendaja
2016–2017	VR Lab OÜ, asutaja (väljunud)
2014–2016	Psühhobuss OÜ, Asutaja ja tegevjuht
2014–2016	Tartu Ülikool, Psühholoogia instituut, õppeassistent
2013–2014	Tartu Ülikool, Õliõpilasesindus, arendusjuht

Teadustegevus

Peamised uurimisvaldkonnad :

- Tähelepanu
- Teadvus
- Virtuaalreaalsus

DISSERTATIONES INFORMATICAE PREVIOUSLY PUBLISHED IN DISSERTATIONES MATHEMATICAE UNIVERSITATIS TARTUENSIS

19. **Helger Lipmaa.** Secure and efficient time-stamping systems. Tartu, 1999, 56 p.
22. **Kaili Müürisep.** Eesti keele arvutigrammatika: süntaks. Tartu, 2000, 107 lk.
23. **Varmo Vene.** Categorical programming with inductive and coinductive types. Tartu, 2000, 116 p.
24. **Olga Sokratova.** Ω-rings, their flat and projective acts with some applications. Tartu, 2000, 120 p.
27. **Tiina Puolakainen.** Eesti keele arvutigrammatika: morfoloogiline ühestamine. Tartu, 2001, 138 lk.
29. **Jan Villemson.** Size-efficient interval time stamps. Tartu, 2002, 82 p.
45. **Kristo Heero.** Path planning and learning strategies for mobile robots in dynamic partially unknown environments. Tartu 2006, 123 p.
49. **Härmel Nestra.** Iteratively defined transfinite trace semantics and program slicing with respect to them. Tartu 2006, 116 p.
53. **Marina Issakova.** Solving of linear equations, linear inequalities and systems of linear equations in interactive learning environment. Tartu 2007, 170 p.
55. **Kaarel Kaljurand.** Attempto controlled English as a Semantic Web language. Tartu 2007, 162 p.
56. **Mart Anton.** Mechanical modeling of IPMC actuators at large deformations. Tartu 2008, 123 p.
59. **Reimo Palm.** Numerical Comparison of Regularization Algorithms for Solving Ill-Posed Problems. Tartu 2010, 105 p.
61. **Jüri Reimand.** Functional analysis of gene lists, networks and regulatory systems. Tartu 2010, 153 p.
62. **Ahti Peder.** Superpositional Graphs and Finding the Description of Structure by Counting Method. Tartu 2010, 87 p.
64. **Vesal Vojdani.** Static Data Race Analysis of Heap-Manipulating C Programs. Tartu 2010, 137 p.
66. **Mark Fišel.** Optimizing Statistical Machine Translation via Input Modification. Tartu 2011, 104 p.
67. **Margus Niitsoo**. Black-box Oracle Separation Techniques with Applications in Time-stamping. Tartu 2011, 174 p.
71. **Siim Karus.** Maintainability of XML Transformations. Tartu 2011, 142 p.
72. **Margus Treumuth.** A Framework for Asynchronous Dialogue Systems: Concepts, Issues and Design Aspects. Tartu 2011, 95 p.
73. **Dmitri Lepp.** Solving simplification problems in the domain of exponents, monomials and polynomials in interactive learning environment T-algebra. Tartu 2011, 202 p.

74. **Meelis Kull.** Statistical enrichment analysis in algorithms for studying gene regulation. Tartu 2011, 151 p.
77. **Bingsheng Zhang.** Efficient cryptographic protocols for secure and private remote databases. Tartu 2011, 206 p.
78. **Reina Uba.** Merging business process models. Tartu 2011, 166 p.
79. **Uuno Puus.** Structural performance as a success factor in software development projects – Estonian experience. Tartu 2012, 106 p.
81. **Georg Singer.** Web search engines and complex information needs. Tartu 2012, 218 p.
83. **Dan Bogdanov.** Sharemind: programmable secure computations with practical applications. Tartu 2013, 191 p.
84. **Jevgeni Kabanov.** Towards a more productive Java EE ecosystem. Tartu 2013, 151 p.
87. **Margus Freudenthal.** Simpl: A toolkit for Domain-Specific Language development in enterprise information systems. Tartu, 2013, 151 p.
90. **Raivo Kolde.** Methods for re-using public gene expression data. Tartu, 2014, 121 p.
91. **Vladimir Šor.** Statistical Approach for Memory Leak Detection in Java Applications. Tartu, 2014, 155 p.
92. **Naved Ahmed.** Deriving Security Requirements from Business Process Models. Tartu, 2014, 171 p.
94. **Liina Kamm.** Privacy-preserving statistical analysis using secure multiparty computation. Tartu, 2015, 201 p.
100. **Abel Armas Cervantes.** Diagnosing Behavioral Differences between Business Process Models. Tartu, 2015, 193 p.
101. **Fredrik Milani.** On Sub-Processes, Process Variation and their Interplay: An Integrated Divide-and-Conquer Method for Modeling Business Processes with Variation. Tartu, 2015, 164 p.
102. **Huber Raul Flores Macario.** Service-Oriented and Evidence-aware Mobile Cloud Computing. Tartu, 2015, 163 p.
103. **Tauno Metsalu.** Statistical analysis of multivariate data in bioinformatics. Tartu, 2016, 197 p.
104. **Riivo Talviste.** Applying Secure Multi-party Computation in Practice. Tartu, 2016, 144 p.
108. **Siim Orasmaa.** Explorations of the Problem of Broad-coverage and General Domain Event Analysis: The Estonian Experience. Tartu, 2016, 186 p.
109. **Prastudy Mungkas Fauzi.** Efficient Non-interactive Zero-knowledge Protocols in the CRS Model. Tartu, 2017, 193 p.
110. **Pelle Jakovits.** Adapting Scientific Computing Algorithms to Distributed Computing Frameworks. Tartu, 2017, 168 p.
111. **Anna Leontjeva.** Using Generative Models to Combine Static and Sequential Features for Classification. Tartu, 2017, 167 p.
112. **Mozhgan Pourmoradnasseri.** Some Problems Related to Extensions of Polytopes. Tartu, 2017, 168 p.

113. **Jaak Randmets.** Programming Languages for Secure Multi-party Computation Application Development. Tartu, 2017, 172 p.
114. **Alisa Pankova.** Efficient Multiparty Computation Secure against Covert and Active Adversaries. Tartu, 2017, 316 p.
116. **Toomas Saarsen.** On the Structure and Use of Process Models and Their Interplay. Tartu, 2017, 123 p.
121. **Kristjan Korjus.** Analyzing EEG Data and Improving Data Partitioning for Machine Learning Algorithms. Tartu, 2017, 106 p.
122. **Eno Tõnisson.** Differences between Expected Answers and the Answers Offered by Computer Algebra Systems to School Mathematics Equations. Tartu, 2017, 195 p.

DISSERTATIONES INFORMATICAE UNIVERSITATIS TARTUENSIS

1. **Abdullah Makkeh**. Applications of Optimization in Some Complex Systems. Tartu 2018, 179 p.
2. **Riivo Kikas**. Analysis of Issue and Dependency Management in Open-Source Software Projects. Tartu 2018, 115 p.
3. **Ehsan Ebrahimi**. Post-Quantum Security in the Presence of Superposition Queries. Tartu 2018, 200 p.
4. **Ilya Verenich**. Explainable Predictive Monitoring of Temporal Measures of Business Processes. Tartu 2019, 151 p.
5. **Yauhen Yakimenka**. Failure Structures of Message-Passing Algorithms in Erasure Decoding and Compressed Sensing. Tartu 2019, 134 p.
6. **Irene Teinemaa**. Predictive and Prescriptive Monitoring of Business Process Outcomes. Tartu 2019, 196 p.
7. **Mohan Liyanage.** A Framework for Mobile Web of Things. Tartu 2019, 131 p.
8. **Toomas Krips.** Improving performance of secure real-number operations. Tartu 2019, 146 p.
9. **Vijayachitra Modhukur.** Profiling of DNA methylation patterns as biomarkers of human disease. Tartu 2019, 134 p.
10. **Elena Sügis.** Integration Methods for Heterogeneous Biological Data. Tartu 2019, 250 p.
11. **Tõnis Tasa.** Bioinformatics Approaches in Personalised Pharmacotherapy. Tartu 2019, 150 p.
12. **Sulev Reisberg.** Developing Computational Solutions for Personalized Medicine. Tartu 2019, 126 p.
13. **Huishi Yin.** Using a Kano-like Model to Facilitate Open Innovation in Requirements Engineering. Tartu 2019, 129 p.
14. **Faiz Ali Shah.** Extracting Information from App Reviews to Facilitate Software Development Activities. Tartu 2020, 149 p.
15. **Adriano Augusto**. Accurate and Efficient Discovery of Process Models from Event Logs. Tartu 2020, 194 p.
16. **Karim Baghery.** Reducing Trust and Improving Security in zk-SNARKs and Commitments. Tartu 2020, 245 p.
17. **Behzad Abdolmaleki.** On Succinct Non-Interactive Zero-Knowledge Protocols Under Weaker Trust Assumptions. Tartu 2020, 209 p.
18. **Janno Siim.** Non-Interactive Shuffle Arguments. Tartu 2020, 154 p.
19. **Ilya Kuzovkin.** Understanding Information Processing in Human Brain by Interpreting Machine Learning Models. Tartu 2020, 149 p.
20. **Orlenys López Pintado**. Collaborative Business Process Execution on the Blockchain: The Caterpillar System. Tartu 2020, 170 p.
21. **Ardi Tampuu.** Neural Networks for Analyzing Biological Data. Tartu 2020, 152 p.