



HAL
open science

Towards Perception-based Character Animation

Ludovic Hoyet

► **To cite this version:**

Ludovic Hoyet. Towards Perception-based Character Animation. Modeling and Simulation. Université Rennes 1, 2022. tel-03671690

HAL Id: tel-03671690

<https://inria.hal.science/tel-03671690>

Submitted on 18 May 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Towards Perception-based Character Animation

Ludovic Hoyet

Chargé de Recherche Inria
Inria, Univ Rennes, IRISA, CNRS

A thesis submitted for the degree of
Habilitation à Diriger des Recherches

For

Université de Rennes 1

Defended on the 3rd of May 2022 in front of

Pr. Bobby Bodenheimer	Vanderbilt University, USA	Rapporteur
Dr. Catherine Pelachaud	CNRS/UPMC, France	Rapporteuse
Pr. Anthony Steed	University College London, UK	Rapporteur
Dr. Ronan Boulic	EPFL, Switzerland	Examineur
Pr. Marie-Paule Cani	Ecole Polytechnique, France	Examinatrice
Pr. Carol O'Sullivan	Trinity College Dublin, Ireland	Examinatrice
Pr. Eric Marchand	Université de Rennes 1, France	Président



Abstract

With the development of real-time graphics, virtual humans have become a requisite to create always more life-like virtual worlds for industries ranging from entertainment to training and education. For instance, in the movie and video game industries their use ranges from creating digital doubles of actors in computer-generated movies, to populating digital worlds using thousands of virtual characters. In these industries, the current trend is to create always more realistic characters, as well as more impressive scenes involving larger and larger groups of characters interacting in real-time with each other. However, because of computational, costs, and time constraints, creating and animating such virtual characters often requires to find the best trade-off between realism and computational needs. This stresses the need for highly efficient methods to animate virtual humans, that are able to create believable motions, behaviours, interactions, as well as to easily scale depending on the number of characters displayed on screen.

To reach this objective, my research activities focus on how to create visually plausible virtual humans, by first understanding how viewers perceive some of their key characteristics. By visually plausible, we mean motions that viewers will consider to be plausible even though they might not be biomechanically or physically correct. It then becomes possible to save computation time, to only simulate a perceptually-plausible approximation of the motion, to save on production costs by selecting optimal subsets of motions, etc.

This manuscript, submitted for the *Habilitation à Diriger des Recherches* degree, presents my involvement in research activities on this topic over approximately the last ten years. Following a detailed presentation of my Curriculum Vitae (including student supervision, publications, participation in European and national projects, etc.), the presentation of our contributions is organized into four chapters, each illustrated through two research focuses. First, we present two chapters focusing on the perception of individual character motions, and on the perception of character interactions. Then, we present two chapters exploring our perceptions in immersive situations, first of virtual crowds, and second of avatars (i.e., the users' representation) which we explore through the psychological construct of virtual embodiment. A final chapter then presents our perspectives for future works, and provides the concluding remarks.

Acknowledgements

First of all, I would like to thank all the jury members who participated in the evaluation of this *Habilitation à Diriger des Recherches*. I would like to thank Eric Marchand for chairing the defense, Bobby Bodenheimer, Catherine Pelachaud and Anthony Steed for taking time in reviewing the manuscript, as well as Ronan Boulic, Marie-Paule Cani and Carol O’Sullivan for participating in the oral defense. I would also like to thank them for all their constructive comments and the exciting discussions around the work presented in this Habilitation.

Of course, all of this work is the result of numerous collaborations, and would not have been possible without the support and the expertise of numerous colleagues over a number of years. For this I would like to thank all the colleagues of the different teams in Rennes, with whom I interacted more or less closely over a period of years. I am thinking of course of people from the new VirtUs team (Anne-Hélène, Julien, Katja, Marc): I have really enjoyed working closely with you, for all the discussions, brainstorming, as well as learning that happened during such a short period of time. I am excited to now start this new adventure with all of you. I also want to thank particularly Anatole and Ferran from the Hybrid team, without whom all this exciting work on the topic of avatars would not have been possible. Again, collaborating with you over the last couple of years has been extremely stimulating intellectually, and I always enjoyed the fact that it is possible to combine both fun and science together on so many projects. Thank you also to all the members of the MimeTIC team, who warmly welcomed me when I joined Inria: Franck of course, for always supporting me in my research since I started with you during my Master’s internship, but also all the team members for the good time we spent together (Charles, Georges, Laurent, Richard, Armel, Anthony, Fabrice, etc.). A big specific thank also to Nathalie, for all the work she does behind the scene to ensure that everything runs as smoothly as possible in the different teams! I would also like to thank all the other persons that have in one way or another participated in these great experiences: Claudio, Pierre, François, Quentin, Fabien, Philippe, Bruno, Valérie, and all those that I forget.

This work is the result of collaborations, but also from a big novel experience for me in terms of supervising a number of people who have been crucial in these successes. I am particularly thinking of Rebecca, Florian, Diane and Benjamin, who have been my first official PhD students to co-supervise. This period with the four of you was particularly precious, for the time I had in being deeply implied with your work, all that I learned by participating in your supervision, and the interactions we had in this pre-Covid time. But I

am sure that this is only the beginning of more exciting experiences, and I keep learning for all those that I have the chance to supervise and whom I would also like to thank: Nicolas, Adèle, Alberto, Lucas, Tairan, Vincenzo, Thomas, Maé, Pierre, Robin, Adrien, etc.

I also believe that I would not be where I am today if it was not for what I consider to be my biggest life-changing experience as of today, i.e., the four years I spent working as a Research Fellow in Trinity College Dublin. It was both an unforgivable professional and personal experience, and I would like to wholeheartedly thank Carol and Rachel for welcoming me in Trinity, for the time and training they gave me, and for all that I learned with them. I will always remember the vibe and excitement that was always present in the lab, in particular close to the Siggraph deadlines, and be grateful for how much I learned with the two of you in these four years.

Finally, my last thanks will of course go to my family: to my partner H el ene who followed me a number of times in my professional journey (and I am so glad that we always managed to balance professional and personal experiences in these different places) and who always has to cope with more and more deadlines to manage, as well as to my two little ones who have also brought a new experience in my life.

Contents

Abstract	iii
Acknowledgements	v
Research Summary	1
I Curriculum Vitae	3
Personal Information	5
Current Position	5
Education & Professional Experience	6
Publications	6
Supervision	11
Research Projects	14
Scientific, Editorial and Organisational Activities	19
Teaching	21
Awards	22
II Research Activity	23
1 Introduction	25
Research Objectives and Challenges	26
Document Overview & Main Contributions	29
2 Perception of Individual Virtual Characters	35
Context and Objectives	35
Focus 1: Distinctiveness and Attractiveness of Virtual Human Motions	37
Focus 2: Perception of Finger Motion	42
Chapter Conclusion	47

3	Perception of Interacting Virtual Characters	51
	Context and Objectives	51
	Focus 1: Perceiving Causality in Virtual Interactions	53
	Focus 2: Perceptual Effect of Shoulder Motions on Crowd Animations	57
	Chapter Conclusion	62
4	Perception of Virtual Crowds in Immersive Environments	65
	Context and Objectives	65
	Focus 1: Gaze Activity in Virtual Crowds	67
	Focus 2: Crowd Navigation in VR: Exploring Haptic Rendering of Collisions	71
	Chapter Conclusion	75
5	Avatars: Perception of Self Virtual Characters in Immersion	79
	Context and Objectives	79
	Focus 1: Avatar and Sense of Embodiment: Studying the Relative Preference Between Appearance, Control and Point of View	82
	Focus 2: Influence of Personality Traits and Body Awareness on the Sense of Embodiment	87
	Chapter Conclusion	91
6	Perspectives	97
	Motion Variety	98
	Animation by Demonstration using VR	100
	Avatar-Friendly Animation & Interaction Techniques	103
	Conclusion	104
	References	107

Research Summary

2007-2010 - PHD CANDIDATE, INSA RENNES / IRISA

I started my research activity at the IRISA laboratory (Rennes, France) in 2007 as a PhD candidate under the co-supervision of Franck Multon (Université Rennes 2, France) and Taku Komura (Edinburgh University, Scotland), as part of the Bunraku team (led by Stephane Donikian at that time). The Bunraku team had a strong experience in real-time animation of humanoid characters, as well as in other related topics of Computer Graphics (including Virtual Reality, Haptics, Rendering, etc.), and had worked in the past on the topics of interactive motion retargeting, synchronization, blending, etc. The general objective of my PhD thesis was to propose novel methods for interactively editing virtual human motions subject to dynamic constraints (e.g., heavier objects to push, turning accelerations). In particular, we proposed a method modeling simplified dynamics to interactively edit virtual human motions (character animation aspects), and combining an understanding of human balance behaviours (biomechanical aspect) with users' perception of differences in dynamic human motions (perceptual aspects).

2011-2014 - RESEARCH FELLOW, TRINITY COLLEGE DUBLIN

During my PhD, I progressively got more and more interested in the topic of the perception of virtual character motions, and got joined Carol O'Sullivan's GV2 group (Graphics, Vision and Visualization) in Trinity College Dublin, a leading expert in Applied Perception for Computer Graphics. I worked there from February 2011 to December 2014 as a Research Fellow, principally working on the Captavatar research project which goal was to create socially and perceptually attractive virtual characters. While working with Carol, we explored several questions related to the perception of causality in virtual character interactions, as well as to the perception of distinctiveness and attractiveness of individual motions displayed on virtual characters. I also collaborated with several other researchers and students from the group (e.g., Rachel McDonnell, Cathy Ennis, Katja Zibrek, Kenneth Ryall, Michele Vicovaro, Hanni Kiiski) on related topics. This is also the time when I acquired most of my expertise to conduct research in the Applied Perception field, which was definitely new to me, while also

consolidating my skills in Computer Graphics and Character Animation. I particularly express my gratitude to Carol and Rachel, from whom I learned a lot, for their time and guidance.

SINCE 2015 - RESEARCH SCIENTIST, INRIA RENNES

This incredible experience ended in December 2014, as I obtained a permanent Research Scientist position at Inria Rennes in the MimeTIC team (led by Franck Multon), which I joined in March 2015. Research activities in MimeTIC revolve around simulating virtual humans that behave in realistic manners, with applications ranging from Ergonomics and Sports, to Computer Graphics. Upon my arrival in the team, I continued to work on topics related to the perception of virtual characters. However, through collaborations I started to expand my research activities to the perception of virtual characters in crowds (with Julien Pettré and Anne-Hélène Olivier), including exploring novel ways of studying pedestrian behaviours using Virtual Reality with the goal of improving crowd simulation models. We also started to explore the question of creating motion variety by focusing on perceivable individual characteristics, through a national Young Researcher project (ANR JCJC Per2). In parallel, we also started to explore the topics of avatars and virtual embodiment, in collaboration with Ferran Argelaguet and Anatole Lécuyer, in the context of the Inria Research project *Avatar*, which involves 6 Inria teams in France. Since I joined Inria, I have participated in the supervision of 13 PhD students, 3 of whom already graduated, 3 Research Fellow, 2 Engineers, and 14 Master students. I am also coordinating 2 projects: the ANR JCJC Per2 and the Inria Research Challenge *Avatar*.

The following of this document is divided into two main parts. Part I details my Curriculum Vitae, including student supervision, publications, participation in European and national projects, other scientific activities, etc. Part II then presents an overview of the work we performed over approximately the last ten years on the topic of creating visually plausible virtual humans. Following a short introduction, our main contributions are organized into four chapters, each illustrated through two research focuses. A final chapter then presents our perspectives for future works, and provides the concluding remarks.

Part I
Curriculum Vitae

Curriculum Vitae

PERSONAL INFORMATION

Name	Hoyet
Firstname	Ludovic
Date of birth	December 15th, 1984
Place of birth	Vannes (56), France
Nationality	French
Sexe	Male
Email	ludovic.hoyet@inria.fr
Website	http://people.rennes.inria.fr/Ludovic.Hoyet/

CURRENT POSITION

I am currently a Research Scientist at Inria Rennes (Chargé de Recherche Classe Normale, 6th grade), since 2015, and working in the MimeTIC team.

My research activities target the efficient real-time simulation of virtual characters based on relevant perceptual features. As there is yet to find approaches providing a good balance between computation time, controllability and realism, I focus on approaches producing perceptually-plausible motions by exploiting what is perceived to be natural, in order to provide animation guidelines and drive the design of novel animation methods. To reach such goals, I have been exploring how viewers perceive biological human motion (i.e., natural movements captured on real actors) displayed on virtual characters, how they interact with such virtual characters (in particular using Virtual Reality), and used perceptual insights to enhance the visual quality of the movements of virtual characters. More recently, I have also been investigating questions related to the virtual representation of the users themselves (i.e., their avatar), with the aim of creating avatars that are better embodied and more interactive.

EDUCATION & PROFESSIONAL EXPERIENCE

- Since 2015 **Research Scientist**, Inria Rennes, team MimeTIC
Topic: Towards new generations of virtual characters: using perceptual insights for efficient real-time simulations
- 2011-2014 **Research Fellow**, Trinity College Dublin, Ireland
Supervisor: Carol O'Sullivan
Involved on the Captavatar project: *Creating socially and perceptually attractive avatars*
- 2010 **PhD (INSA Rennes)**, Inria/IRISA laboratory, team Bunraku, France
Title: Dynamic Adaptation of Human Motions
Supervisors: Franck Multon (Université Rennes 2), Taku Komura (Edinburgh University)
Examinors: Ronan Boulic (EPFL), Jean-Paul Laumond (LAAS-CNRS)
- 2007 **MSc/Engineering Degrees**, Computer Sciences, INSA Rennes, France
MSc Thesis Title: Dynamic Adaptation of Aerial Human Motions
Supervisor: Franck Multon (Université Rennes 2)

PUBLICATIONS

Over my carrier, I have participated in the publication of 22 peer-reviewed international journal articles, 32 international peer-reviewed international conference articles, and numerous other types of communications (posters, abstracts, etc.). In accordance with the French AERES criteria, as well as with practices of the Computer Graphics community recently referenced by the Société Informatique de France, 30 of these publications are of rank A, including 5 first-author publications (see prestigious and sought-after channels¹).

Three of the papers published in international journals as first-author were accepted and presented at SIGGRAPH and SIGGRAPH ASIA (two best conferences in Graphics), and published in the ACM Transaction on Graphics, best journal in Computer Science, Software Engineering. Eight other papers led by significantly supervised PhD students were also published in the IEEE Transactions on Visualization and Computer Graphics (5), the second best journal in the field, and Computer Graphics Forum (3), in the proceedings of Eurographics (also one of the best conference in Graphics). Several publications were also accepted in highly regarded and selective conferences and journals of the domains, including ACM Transaction on Applied Perception, IEEE Virtual Reality, IEEE International Symposium on Mixed and Augmented Reality, ACM Symposium on applied Perception, etc.

¹<https://github.com/societe-informatique-de-france/referentiel-pratiques-publication-2019/blob/master/analyse/gdrigrv.md>

International Journals

1. A. Colas, W. van Toll, K. Zibrek, **L. Hoyet**, A-H. Olivier, J. Pettre. Interaction Fields: Intuitive Sketch-based Steering Behaviors for Crowd Simulation. *Computer Graphics Forum (Eurographics, Honorable mention)*, 2022.
2. B. Cabrero-Daniel, R. Marques, **L. Hoyet**, J. Pettré, J. Blat. Dynamic Combination of Crowd Steering Policies Based on Context. *Computer Graphics Forum (Eurographics)*, 2022.
3. T. Yin, **L. Hoyet**, M. Christie, M-P. Cani, J. Pettré. The One-Man-Crowd: Single User Generation of Crowd Motions Using Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics (IEEE Virtual Reality)*, 2022.
4. D. Dewez, **L. Hoyet**, A. Lécuyer, F. Argelaguet Sanz. "Do You Need Another Hand? Investigating Dual Body Representations During Anisomorphic 3D Manipulation". *IEEE Transactions on Visualization and Computer Graphics (IEEE Virtual Reality)*, 2022.
5. N. Olivier, G. Kerbiriou, F. Arguelaguet, Q. Avril, F. Danieau, P. Guillotel, **L. Hoyet** and F. Multon. Study on Automatic 3D Facial Caricaturization: From Rules to Deep Learning. *Frontiers in Virtual Reality*, 2:785104, 2022.
6. L. Mourot, **L. Hoyet**, F. Le Clerc, F. Schnitzler, P. Hellier. A Survey on Deep Learning for Skeleton-Based Human Animation. *Computer Graphics Forum*, 2021.
7. B. Cabrero-Daniel, R. Marques, **L. Hoyet**, J. Pettré, J. Blat. A Perceptually-Validated Metric for Crowd Trajectory Quality Evaluation. *Proceedings of the ACM on Computer Graphics and Interactive Techniques (Symposium on Computer Animation Special Issue)*, 4(3), 42, 2021.
8. R. Fribourg, E. Blanpied, **L. Hoyet**, A. Lécuyer, F. Argelaguet. Does virtual threat harm VR experience?: Impact of threat occurrence and repeatability on virtual embodiment and threat response. *Computers Graphics*, 2021.
9. F. Berton, F. Grzeskowiak, A. Bonneau, A. Jovane, M. Aggravi, **L. Hoyet**, A-H. Olivier, C. Pacchierotti, J. Pettre. Crowd Navigation in VR: exploring haptic rendering of collisions. *IEEE Transactions on Visualization and Computer Graphics*, 2020.
10. R. Fribourg, N. Ogawa, **L. Hoyet**, F. Argelaguet, T. Narumi, M. Hirose, A. Lécuyer. Virtual Co-Embodiment: Evaluation of the Sense of Agency while Sharing the Control of a Virtual Body among Two Individuals. *IEEE Transactions on Visualization and Computer Graphics*, 2020.
11. K. Zibrek, B. Niay, A-H. Olivier, **L. Hoyet**, J. Pettre, R. McDonnell. The Effect of Gender and Attractiveness of Motion on Proximity in Virtual Reality. *ACM Transactions on Applied Perception (SAP 2020 Special Issue)*, 2020.
12. R. Fribourg, F. Argelaguet, A. Lécuyer, **L. Hoyet**. Avatar and Sense of Embodiment: Studying the Relative Preference Between Appearance, Control and Point of View. *IEEE Transactions on Visualization and Computer Graphics (IEEE Virtual Reality, Best Journal Papers Award)*, 2020.
13. A. Mucherino, J. Omer, **L. Hoyet**, P. Robuffo Giordano, F. Multon. An Application-based Characterization of Dynamical Distance Geometry Problems. In *Optimization Letters*, 2018.
14. **L. Hoyet**, A-H. Olivier, R. Kulpa, J. Pettré. Perceptual Effect of Shoulder Motions on Crowd Animations. In *ACM Transaction on Graphics (SIGGRAPH 2016)*, 35(4), 2016.
15. **L. Hoyet**, F. Argelaguet, C. Nicole, A. Lécuyer. "Wow! I Have Six Fingers!": Would You Accept Structural Changes of Your Hand in VR? In *Frontiers in Robotics and AI* 3:27, 2016.
16. K. Zibrek, **L. Hoyet**, K. Ruhland, R. McDonnell. Exploring the Effect of Motion Type and Emotions on the Perception of Gender in Virtual Humans. *ACM Transactions on Applied Perception*, 12(3), 2015.
17. H. Kiiski, **L. Hoyet**, A. Woods, C. O'Sullivan, F. Newell. Strutting Hero, Sneaking Villain: Utilising Body Motion Cues to Predict the Intentions of Others. *ACM Trans. on Applied Perception*, 12(4), 2015.
18. M. Vicovaro, **L. Hoyet**, L. Burigana, C. O'Sullivan. Perceptual Evaluation of Motion Editing for Realistic Throwing Animations. *ACM Transactions on Applied Perception*, 11(2), 2014.
19. **L. Hoyet**, K. Ryall, K. Zibrek, H. Park, J. Lee, J. Hodgins, C. O'Sullivan. Evaluating the Distinctiveness

- and Attractiveness of Human Motions on Realistic Virtual Bodies. *ACM Trans. on Graphics (SIGGRAPH Asia 2013)*, 32(6), 2013.
20. H. Shum, **L. Hoyet**, E. Ho, T. Komura, F. Multon. Natural preparation behavior synthesis. In *Computer Animation and Virtual Worlds*, 2013.
 21. **L. Hoyet**, R. McDonnell, C. O'Sullivan. Push it Real: Perceiving Causality in Virtual Interactions. *ACM Transaction on Graphics (SIGGRAPH 2012)*, 31(4), 2012.
 22. F. Multon, R. Kulpa, **L. Hoyet**, T. Komura. Interactive animation of virtual humans based on motion capture data. *Computer Animation and Virtual Worlds*, Vol. 20(5-6), pages 491-500, 2009.

Peer-reviewed International Conferences

1. P. Raimbaud, A. Jovane, K. Zibrek, C. Pacchierotti, M. Christie, **L. Hoyet**, J. Pettré, A-H. Olivier. The stare in-the-crowd effect in virtual reality. *IEEE Conference on Virtual Reality and 3D User Interfaces (IEEE VR)*, 2022.
2. R. Adili, B. Niay, K. Zibrek, A-H. Olivier, J. Pettré, **L. Hoyet**. Perception of Motion Variations in Large-Scale Virtual Human Crowds. *ACM Motion, Interaction and Games (MIG '21)*, 2021.
3. D. Dewez, **L. Hoyet**, A. Lécuyer, F. Argelaguet. Towards "Avatar-Friendly" 3D Manipulation Techniques: Theoretical Background and Practical Design Guidelines. *ACM CHI Conference on Human Factors in Computing Systems*, 2021.
4. R. Fribourg, E. Blanpied, **L. Hoyet**, A. Lécuyer, F. Argelaguet. Influence of Threat Occurrence and Repeatability on the Sense of Embodiment and Threat Response in VR. *International Conference on Artificial Reality and Telexistence & Eurographics Symposium on Virtual Environments*, 2020.
5. D. Dewez, **L. Hoyet**, A. Lécuyer, F. Argelaguet Sanz. Studying the Inter-Relation Between Locomotion Techniques and Embodiment in Virtual Reality. *IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 2020.
6. B. Niay, A-H. Olivier, K. Zibrek, J. Pettré, **L. Hoyet**. Walk Ratio: Perception of an Invariant Parameter of Human Walk on Virtual Characters. *ACM Symposium on Applied Perception*, 2020.
7. N. Olivier, **L. Hoyet**, F. Argelaguet, F. Danieau, Q. Avril, P. Guillotel, A. Lécuyer, F. Multon. The impact of stylization on face recognition. *ACM Symposium on Applied Perception*, 2020.
8. F. Berton, **L. Hoyet**, A-H. Olivier, J. Bruneau, O. Le Meur, J. Pettré. Eye-Gaze Activity in Crowds: Impact of Virtual Reality and Density. *IEEE Conf. on Virtual Reality and 3D User Interfaces (VR)*, 2020.
9. T. Duverné, T. Rougnant, F. Le Yondre, F. Berton, J. Bruneau, K. Zibrek, J. Pettré, **L. Hoyet**, A-H. Olivier. Effect of social settings on proxemics during social interactions in real and virtual conditions. *EuroVR 2020*.
10. D. Dewez, R. Fribourg, F. Argelaguet Sanz, **L. Hoyet**, D. Mestre, M. Slater, A. Lécuyer. Influence of Personality Traits and Body Awareness on the Sense of Embodiment in Virtual Reality. *IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 2019.
11. F. Berton, A-H. Olivier, J. Bruneau, **L. Hoyet**, J. Pettré. Studying Gaze Behaviour During Collision Avoidance With a Virtual Walker: Influence of the Virtual Reality Setup. *IEEE Virtual Reality*, 2019.
12. **L. Hoyet**, C. Spies, P. Plantard, A. Sorel, R. Kulpa, F. Multon. Influence of Motion Speed on the Perception of Latency in Avatar Control. *Workshop CRDH, 2nd IEEE International Conference on Artificial Intelligence & Virtual Reality (AIVR)*, 2019.
13. Z. Liu, A. Mucherino, **L. Hoyet**, F. Multon. Surface based Motion Retargeting by Preserving Spatial Relationship. *Motion, Interaction and Games*, 2018.
14. E. Carrigan, **L. Hoyet**, R. McDonnell, Q. Avril. A Preliminary Investigation into the Impact of Training for Example-Based Facial Blendshape Creation. *Eurographics Short Papers*, 2018.
15. R. Fribourg, F. Argelaguet, **L. Hoyet**, A. Lécuyer. Studying the Sense of Embodiment in VR Shared

- Experiences. *IEEE Virtual Reality*, 2018.
16. F. Elain, A. Mucherino, **L. Hoyet**, R. Kulpa. Feature Selection in Time-Series Motion Databases. *Proceedings of FedCSIS*, 2018.
 17. A. Bernardin, **L. Hoyet**, A. Mucherino, D. Gonçalves, F. Multon. Normalized Euclidean Distance Matrices for Human Motion Retargeting. *Proceedings of the Tenth International Conference on Motion in Games (MIG '17)*, 2017.
 18. A. Mucherino, D. Gonçalves, A. Bernardin, **L. Hoyet**, F. Multon. A Distance-Based Approach for Human Posture Simulations. *Proceedings of the 2017 Federated Conference on Computer Science and Information Systems*, 2017.
 19. F. Argelaguet, **L. Hoyet**, M. Trico, A. Lécuyer. The Role of Interaction in Embodiment: The Effects of the Virtual Hand Representation. *IEEE Virtual Reality*, 2016.
 20. C. Ennis, **L. Hoyet**, C. O'Sullivan. Eye-tracktive: Measuring attention to body parts when judging human motions. *Eurographics Short Papers*, 2015.
 21. C. Ennis, **L. Hoyet**, A. Egges, R. McDonnell. Emotion Capture: Emotionally Expressive Characters for Games. *ACM Motion in Games (MIG)*, 2013.
 22. K. Zibrek, **L. Hoyet**, K. Ruhland, R. McDonnell. Evaluating the Effect of Emotion on Gender Recognition in Virtual Humans. *ACM Symposium on Applied Perception*, 2013.
 23. H. Shum, **L. Hoyet**, E. Ho, T. Komura, F. Multon. Preparation behaviour synthesis with reinforcement learning. *International Conference on Computer Animation and Social Agents (CASA)*, 16-18 May, 2013.
 24. **L. Hoyet**, K. Ryall, R. McDonnell, C. O'Sullivan. Sleight of Hand: Perception of finger motion from reduced marker sets. *ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (I3D)*, 2012.
 25. M. Vicovaro, **L. Hoyet**, L. Burigana, C. O'Sullivan. Evaluating the Plausibility of Edited Throwing Animations. *ACM SIGGRAPH/Eurographics Symposium on Computer animation*, 2012.
 26. M. Pražák, **L. Hoyet**, C. O'Sullivan. Perceptual evaluation of footskate cleanup. *ACM SIGGRAPH/Eurographics Symposium on Computer animation*, 2011.
 27. **L. Hoyet**, F. Multon, T. Komura, A. Lécuyer. Can We Distinguish Biological Motions of Virtual Humans? Perceptual Study With Captured Motions of Weight Lifting. *ACM Virtual Reality and Technology (VRST)*, 2010.
 28. **L. Hoyet**, F. Multon, T. Komura, A. Lécuyer. Perception based real-time dynamic adaptation of human motions. *Motion in Games (MIG)*, 2010.
 29. X. Liang, **L. Hoyet**, W. Geng, F. Multon. Responsive action generation by physically-based motion retrieval and adaptation. *Motion in Games (MIG)*, 2010
 30. F. Multon, R. Kulpa, **L. Hoyet**, T. Komura. From Motion Capture to Real-Time Character Animation. *Motion in Games (MIG)*, 2008.
 31. F. Multon, **L. Hoyet**, T. Komura, R. Kulpa. Interactive control of physically-valid aerial motion: application to VR training system for gymnasts. *ACM Virtual Reality and Technology (VRST)*, 2007.
 32. F. Multon, **L. Hoyet**, T. Komura, R. Kulpa. Dynamic motion adaptation for 3D acrobatic Humanoids. *IEEE Humanoids*, 2007.

Others

1. V. Etien, R. Sterna, **L. Hoyet**, K. Zibrek. May I sit next to you? The effect of motion quality of virtual agents on the proximity in virtual reality. In *Motion, Interaction and Games (MIG '21) - Posters Session*, 2021.
2. P. Raimbaud, A. Jovane, K. Zibrek, C. Pacchierotti, M. Christie, **L. Hoyet**, J. Pettré, A.H. Olivier. Reactive Virtual Agents: A Viewpoint-Driven Approach for Bodily Nonverbal Communication. *Proceedings of the 21th ACM International Conference on Intelligent Virtual Agents (IVA '21)*, 2021.
3. S. Bourgaize, M. Cinelli, F. Berton, B. Niay, **L. Hoyet**, J. Pettré, A.H. Olivier. Walking speed and

- trunk sway: Influence of an approaching person’s gait pattern on collision avoidance. *Journal of Vision* September, 21(9), 2021.
4. I. Podkosova, K. Zibrek, J. Pettré, **L. Hoyet**, A-H. Olivier. Exploring behaviour towards avatars and agents in immersive virtual environments with mixed-agency interactions. *IEEE Virtual Reality Workshops 2021 - VHCIE Workshop*, 2021.
 5. T. Chatagnon, A-H. Olivier, **L. Hoyet**, J. Pettré, C. Pontonnier. Modeling physical interactions in human crowds: a pilot study of individual response to controlled external pushes. *Computer Methods in Biomechanics and Biomedical Engineering*, 2021.
 6. B. Niay, A-H. Olivier, K. Zibrek, J. Pettré, **L. Hoyet**. Walk Ratio: Perception of an Invariant Parameter of Human Walk on Virtual Characters. Oral presentation in the French Computer Graphics Days (jFIG).
 7. A. Colas, W. van Toll, **L. Hoyet**, C. Pacchierotti, M. Christie, K. Zibrek, A-H. Olivier, J. Pettré. Interaction Fields: Sketching Collective Behaviours. Poster in *ACM Motion, Interaction, and Games*.
 8. R. Fribourg, F. Argelaguet, A. Lécuyer, **L. Hoyet**. Avatar et Sentiment d’Incarnation: Étude de la préférence relative entre l’apparence, le contrôle et le point de vue. *Workshop sur les Affects, Compagnons artificiels et Interactions*.
 9. K. Zibrek, B. Niay, A-H. Olivier, **L. Hoyet**, J. Pettré, R. McDonnell. Walk this way: Evaluating the effect of perceived gender and attractiveness of motion on proximity in virtual reality. *Virtual Humans and Crowds for Immersive Environments (VHCIE), IEEE Virtual Reality Workshop*, 2020.
 10. B. Niay, A-H. Olivier, J. Pettré, **L. Hoyet**. The Influence of Step Length to Step Frequency Ratio on the Perception of Virtual Walking Motions. Poster in *ACM Motion in Games*, 2019.
 11. B. Niay, A-H. Olivier, J. Pettré, **L. Hoyet**. The Influence of Step Length to Step Frequency Ratio on the Perception of Virtual Walking Motions. Poster in *ACM Symposium on Applied Perception*, 2019.
 12. Adapter les mouvements aux personnages virtuels pour l’animation de foules. *Magazine Inria Emergences, Lettre d’information n° 56*, 2019.
 13. De meilleures incarnations de nous-mêmes. *Magazine Inria Emergences, Lettre d’information n° 58*, 2019.
 14. T. Duverne, T. Rougnant, F. Le Yondre, F. Berton, J. Bruneau, **L. Hoyet**, J. Pettré, A-H. Olivier. Analyse des réactions corporelles à la transgression des normes de proxémie en contexte de spectacle sportif : expérimentations en situations réelles et virtuelles. *ACAPS 2019 - 18ème congrès de l’Association des Chercheurs en Activités Physiques et Sportives*, 2019.
 15. F. Berton, **L. Hoyet**, A-H. Olivier, J. Pettré. Gaze Behaviour During Collision Avoidance Between Walkers: A Preliminary Study to Design an Experimental Platform. *Virtual Humans and Crowds for Immersive Environments (VHCIE), IEEE Virtual Reality Workshop*, 2018.
 16. A-H. Olivier, J.M. Auberlet, A. Domes, M-A. Granie, **L. Hoyet**, J. Pettré. Etude du piéton en réalité virtuelle: état de l’art, enjeux et perspectives. *International Conference Olympic Games*, 2018.
 17. **L. Hoyet**. Inria Project Lab “AVATAR: The Next Generation of our Virtual Selves in Digital Worlds”. *ERCIM News 115, Special theme: Digital Twins*, 2018.
 18. H. Kiiski, **L. Hoyet**, B. Cullen, C. O’Sullivan, F. Newell. Perception of traits from static and dynamic visual cues in faces and bodies. *Perception 42 ECVF 2013 Abstracts*, page 197, 2013.
 19. H. Kiiski, **L. Hoyet**, B. Cullen, C. O’Sullivan, F. Newell. Perception and Prediction of Social Intentions from Human Body Motion. Poster in *ACM Symposium on Applied Perception*, 2013.
 20. H. Kiiski, **L. Hoyet**, K. Zibrek, C. O’Sullivan, F. Newell. Audio-visual interactions in the perception of intention from actions. *International Multisensory Research Forum (IMRF)*, 2013.
 21. E. Roudaia, **L. Hoyet**, C. O’Sullivan, D. McGovern, F. Newell. Ageing reduces sensitivity to timing mismatches in the perception of human motion. *Perception 42 ECVF Abstracts*, page 139, 2013.
 22. E. Roudaia, **L. Hoyet**, D. McGovern, C. O’Sullivan, F. Newell. Effects of Ageing and Sound on Perceived Timing of Human Interactions. Poster in *ACM Symposium on Applied Perception*, 2013.
 23. M. Vicovaro, **L. Hoyet**, L. Burigana, C. O’Sullivan. Randomly interleaved staircases and ‘acceptance thresholds’ in computer graphics experiments. *Proceedings of the 29th annual meeting of the International*

Society for Psychophysics, 2013.

24. M. Vicovaro, **L. Hoyet**, L. Burigana, C. O’Sullivan. Evaluating observers’ sensitivity to errors in virtual throwing animations. *Perception* (41), European Conference on Visual Perception 2012 Abstract Supplement, page 149, 2012.
25. M. Vicovaro, **L. Hoyet**, L. Burigana, C. O’Sullivan. Evaluating Observers’ Sensitivity to Errors in Human and Physical Throws. Poster in the Symposium on Applied Perception, 2012.
26. K. Ryall, **L. Hoyet**, J. Hodgins, C. O’Sullivan. Exploring Sensitivity to Time-Warped Biological Motion. *Perception* (41), European Conference on Visual Perception 2012 Abstract Supplement, page 149, 2012.
27. **L. Hoyet**, K. Ryall, R. McDonnell, C. O’Sullivan. Sleight of Hand: Perception of finger motion from reduced marker sets. Poster in ACM Symposium on Computer Animation, 2012.
28. **L. Hoyet**, F. Multon. Comparison of models of dynamic balance in biological motions. *Computer Methods in Biomechanics and Biomedical Engineering*. Proceedings of 36th Congress of the Société de Biomécanique, 2011.
29. **L. Hoyet**, F. Multon, K. Mombaur, E. Yoshida. Balance in dynamic situations: role of the underlying model. *Computer Methods in Biomechanics and Biomedical Engineering*, 12:1, 147-149. Proceedings of 34th Congress of the Société de Biomécanique, 2009.
30. **L. Hoyet**, F. Multon, K. Mombaur, E. Yoshida. Balance in dynamic situations: role of COP, xCOM and ground reaction force. *ISPGR Satellite Symposium*, 2009.
31. **L. Hoyet**, F. Multon, K. Mombaur, E. Yoshida. Influence du choix du modèle sur l’estimation du centre de pression à partir du ZMP. *Journées Nationales de Robotique Humanoïde*, 2009.
32. **L. Hoyet**, F. Multon, T. Komura, R. Kulpa. Adaptation dynamique de mouvements aériens. *Journées de l’Association Française d’Informatique Graphique*, 2007.
33. F. Multon, **L. Hoyet**, T. Komura, R. Kulpa. Simulation biomécanique rapide de mouvements aériens respectant la dynamique. *Communication orale au congrès de l’Association des Chercheurs en Activité Physique et Sportive*, Leuven, Belgique, 2007.

SUPERVISION

PhD Students

Ms. Yuliya Patotskaya (2022-), co-supervised at 33% with Katja Zibrek and Julien Pettré. *Appealing Character Design for Embodied Virtual Reality*.

Ms. Maé Mavromatis (2021-), co-supervised at 33% with Ferran Argelaguet and Anatole Lécuyer. *Towards “Avatar-Friendly” Characterization of Virtual Reality Interaction Methods*.

Mr. Vincenzo Abichequer-Sangalli (2020-), co-supervised at 25% with Marc Christie, Carol O’Sullivan and Julien Pettré. *Humains virtuels expressifs et réactifs pour la réalité virtuelle*.

Mr. Thomas Chatagnon (2020-), co-supervised at 25% with Anne-Hélène Olivier, Charles Pontonnier, and Julien Pettré. *Modélisation énergétique multi-échelle pour les foules denses*.

Mr. Tairan Yin (2020-), co-supervised at 25% with Marc Christie, Marie-Paule Cani and Julien Pettré. *Création de scènes peuplées dynamiques pour la réalité virtuelle*.

Mr. Lucas Mourot (2020-), CIFRE InterDigital, co-supervised at 25% with Pierre Hellier, Franck Multon and François Le Clerc. *Learning-Based Human Character Animation Synthesis for Content Production*.

Mr. Alberto Jovane (2019-), co-supervised at 25% with Marc Christie, Claudio Pacchierotti and Julien Pettré. *Modélisation de mouvements réactifs et comportements non verbaux pour la création d'acteurs digitaux pour la réalité virtuelle.*

Ms. Adèle Colas (2019-), co-supervised at 25% with Anne-Hélène Olivier, Claudio Pacchierotti and Julien Pettré. *Modélisation de comportements collectifs réactifs et expressifs pour la réalité virtuelle.*

Mr. Nicolas Olivier (2019-), CIFRE Interdigital, co-supervised with Franck Multon, Anatole Lécuyer, Ferran Argelaguet, Fabien Danieau, Quentin Avril. *Customisation adaptative d'avatars pour expériences immersives.*

Mr. Benjamin Niay (2018-), co-supervised at 33% with Anne-Hélène Olivier and Julien Pettré. *A Framework for Synthesizing Personalised Human Motions from Motion Capture Data and Perceptual Information.*

Ms. Diane Dewez (2018-2021), co-supervised at 33% with Ferran Argelaguet and Anatole Lécuyer. *Avatar-based interactions in Virtual Reality - PhD defended on the 29/11/2021.*

Mr. Florian Berton (2017-2020), co-supervised at 33% with Anne-Hélène Olivier and Julien Pettré. *Immersive Virtual Crowds: Evaluation of Pedestrian Behaviours in Virtual Reality - PhD defended on the 14/12/2020.*

Ms. Rebecca Fribourg (2017–2020), co-supervised at 33% with Ferran Argelaguet and Anatole Lécuyer. *Contribution to the Study of Factors Influencing the Sense of Embodiment Towards Avatars in Virtual Reality - PhD defended on the 04/11/2020.*

Visiting PhD Students

Ms. Emma Carrigan, Trinity College Dublin, Ireland (6-month visit, 2017), co-supervised at 50% with Mr. Quentin Avril (Technicolor). *Investigating the impact of chosen training scans for example-based facial blendshape creation.*

Mr. Yijun Shen, Northumbria University, UK (4-month visit, 2016), co-supervised at 50% with Franck Multon. *Correction de mouvement à partir de graphes d'interactions.*

Michele Vicovaro, University of Padoue, Italie (6-month visit, 2013), co-supervised at 70% with Carol O'Sullivan. *Evaluating the Plausibility of Edited Throwing Animations.*

Research Fellows

Mr. Ific Goudé (June 2021 -), co-supervised at 33% with Anne-Hélène Olivier and Julien Pettré. *Improving crowd simulation through visual saliency of virtual environments.*

Pierre Raimbaud (Nov. 2020 -), co-supervised at 33% with Anne-Hélène Olivier and Julien Pettré. *Non-verbal communication and interactions: towards reactive and expressive virtual humans.*

Dr. Zhiguang Liu (Nov. 2016 - May 2018), co-supervised at 33% with Franck Multon and Olivier Sentieys. *Combining Smart Sensors and New Motion Representations for Human Performance Analysis.*

Engineers & Research Assistants

Mr. Adrien Reuzeau (Oct. 2020 -), Engineer. *AvatarReady: A unified platform for the next generation of avatars in digital worlds*, 50% co-supervised with Ferran Argelaguet.

Mr. Robin Adili (Sept. 2020 -), Engineer. *Motion Editing Developments for Character Animation Tools*

Ms. Katja Zibrek (March 2013 - June 2014), Research Assistant, 20% co-supervised with Rachel McDonnell.

Ms. Katja Zibrek (April 2012 - March 2013), Research Assistant, 70% co-supervised with Carol O’Sullivan.

Mr. Kenneth Ryall (2011 - 2014), Research Assistant, 50% co-supervised with Carol O’Sullivan.

Master Students

Mr. Vincent Etien (2021, 6 months), 50% co-supervised with Katja Zibrek. *Towards the creation of appealing realistic virtual humans in virtual reality.*

Mr. Arnaud Roger (2021, 6 months), 50% co-supervised with Benjamin Niay. *Creating Perceptual Variations in Virtual Character Locomotion.*

Mr. Robin Adili (2020, 6 months), 33% co-supervised with Anne-Hélène Olivier and Julien Pettré. *Synthesising Motion Variations in Human Crowds.*

Mr. Stéven Picard (2020, 6 months), 33% co-supervised with Anne-Hélène Olivier and Julien Pettré. *Evaluating the Influence of Character Realism on Avoidance Strategies in VR.*

Mr. Alexandre Stathopoulos (2019, 2 months), 50% co-supervised with Anne-Hélène Olivier. *Intégration de capteurs inertiels pour l’étude des évitements piétons en réel et virtuel.*

Mr. Nathan Calvarin (2019, 6 months), 33% co-supervised with Anne-Hélène Olivier and Julien Pettré. *Développement d’une plateforme pour l’étude d’évitements de piétons en réel et virtuel.*

Ms. Diane Dewez (2018, 6 months), 33% co-supervised with Ferran Argelaguet and Anatole Lécuyer. *Influence de la personnalité sur l’incarnation virtuelle.*

Mr. Clément Spies (2018, 5 months), 50% co-supervised with Franck Multon. *Influence de la latence sur la perception de son propre mouvement en réalité virtuelle.*

Ms. Tiffany Luong (2017, 6 months), 33% co-supervised with Ferran Argelaguet and Anatole Lécuyer. *Low-Cost Full-Body Control of Human Avatars for Immersive VR Systems.*

Mr. Antonin Bernardin (2017, 6 months), 50% co-supervised with Antonio Mucherino. *Human Movement Adaptation using a Novel Distance-Based Approach.*

Ms. Rebecca Fribourg (2016, 6 months), 33% co-supervised with Ferran Argelaguet and Anatole Lécuyer. *Effets d’interactions multi-utilisateurs sur la sensation d’incarnation virtuelle .*

Mr. Arnaud Biallais (2016, 7 months), 70% co-supervised with Richard Kulpa. *approches perceptives pour l'adaptation de mouvement par cinématique inverse.*

Mr. Corentin Nicole (2016, 4 months), 33% co-supervised with Ferran Argelaguet and Anatole Lécuyer. *“Wow! I Have Six Fingers!”: Would You Accept Structural Changes of Your Hand in VR?*

Mr. Michael Trico (2015, 6 months), 20% co-supervised with Ferran Argelaguet and Anatole Lécuyer. *The Role of Interaction in Embodiment.*

Mr. Damien Coin-Perard (2014, 4 months), 20% co-supervised with Ferran Argelaguet and Anatole Lécuyer. *Developing tools for motion capture and character animation.*

Mr. Mayeul de Werbier (2013, 6 months), 20% co-supervised with Franck Multon. *Simulating Virtual Characters Based on User Perception .*

Bachelor Students

Mr. Thomas Kergoat (2020, 3 months). *Développement de briques logiciel permettant la création d'humains virtuels réalistes et réactifs sous Unity3D.*

Ms. Laure Du Mesnildot (2018, 3 months). *Adaptation automatique de la morphologie de l'avatar à son utilisateur.*

Mr. Florian Elain (2017, 2 months), 50% co-supervised with Antonio Mucherino. *A triclustering approach for feature selection and sample classification prediction.*

RESEARCH PROJECTS

This section lists the national and international research projects in which I am or have been involved. For each project, a short description of the overall objectives is provided, as well as a list of the persons involved locally (PR: Professor; DR: Research Director; MCU: Assistant Professor; CR: Research Scientist; PhD: PhD candidate, PD: Postdoc). The name of the principal investigator of each project is underlined.

Funding & Coordination

ANR JCJC Per2 (coordinator, 2018-2022, budget 280k€). The objective of this project entitled Perception-based Human Motion Personalisation is to focus on how viewers perceive motion variations to automatically produce natural motion personalisation accounting for inter-individual variations. In short, our goal is to automate the creation of motion variations to represent given individuals according to their own characteristics, and to produce natural variations that are perceived and identified as such by users. Challenges addressed in this project consist in (i) understanding and quantifying what makes motions of individuals perceptually different, (ii) synthesising motion variations based on these identified relevant

perceptual features, according to given individual characteristics, and (iii) leveraging the synthesis of motion variations to explore their creation for interactive large-scale scenarios where both performance and realism are critical.

Local team: Ludovic Hoyet (CR), Julien Pettré (DR), Anne-Hélène Olivier (MCU), Benjamin Niay (PhD), Robin Adili (Engineer)

Website: <https://project.inria.fr/per2/>

Inria Research Challenge *Avatar* (coordinator, 2018-2022, budget ~1M€). This project aims at designing avatars (i.e., the user’s representation in virtual environments) that are better embodied, more interactive and more social, through improving all the pipeline related to avatars, from acquisition and simulation, to designing novel interaction paradigms and multi-sensory feedback. It involves 6 Inria teams (GraphDeco, Hybrid, Loki, MimeTIC, Morpheo, Potioc), Prof. Mel Slater (Uni. Barcelona), and 2 industrial partners (InterDigital and Faurecia).

Local team: Ludovic Hoyet (CR), Ferran Argelaguet (CR), Anatole Lécuyer (DR), Diane Dewez (PhD), Rebecca Fribourg (PhD), Maé Mavromatis (PhD)

Website: <http://avatar.inria.fr>

Rennes Metropole: Acquisition Equipement Scientifique *Interact* (coordinator, 2019-2021, budget 70k€). The AES Interact funding aims at acquiring research equipment with the goal of pushing the limits of understanding and synthesising interactions between real and virtual humans. More precisely, the goal is to enable the immersion of several users in the same virtual environment, while both increasing the ecological validity of the immersion (e.g., HMD with larger field of views, capturing subtler information such as facial or finger motions) and enabling researchers to measure simultaneously more information about the interactions (e.g., eyetracking).

Rennes Metropole: Allocation Installation Scientifique (coordinator, 2016, budget 40k€). This funding was awarded by Rennes Metropole after my arrival in the MimeTIC team, for my research project on creating perception-based interactions between virtual characters, with the goal of acquiring two Xsens Motion capture systems to conduct my research.

Participation in European Projects

H2020 ICT ITN CLIPE (2020-2024). CLIPE (Creating Lively Interactive Populated Environments) is an Innovative Training Network (ITN) funded by the Marie Skłodowska-Curie program of the European Commission. The primary objective of CLIPE is to train a generation of innovators and researchers in the field of virtual characters simulation and animation. Advances in technology are pushing towards making VR/AR worlds a daily experience. Whilst virtual characters are an important component of these worlds, bringing them to life and giving them interaction and communication abilities requires highly specialized programming combined with artistic skills, and considerable investments: millions spent on countless coders and designers to develop video-games is a typical example. The research objective of CLIPE is to design the next-generation of VR-ready characters. CLIPE is addressing the most important current aspects of the problem, making the characters capable of: behaving more naturally; interacting with real users sharing a virtual experience with them; being more intuitively and extensively controllable for virtual worlds designers. To meet our objectives, the CLIPE consortium gathers some of the main European actors in the field of VR/AR, computer graphics, computer animation, psychology and perception. CLIPE also extends its partnership to key industrial actors of populated virtual worlds, giving students the ability to explore new application fields and start collaborations beyond academia.

Local team: Julien Pettré (DR), Marc Christie (MCU), Ludovic Hoyet (CR), Tairan Yin (PhD), Vincenzo Abichequer-Sangalli (PhD)

Website: <https://www.clipe-itn.eu/>

H2020 FET-Open CrowdDNA (2020-2024). This project aims to enable a new generation of “crowd technologies”, i.e., a system that can prevent deaths, minimize discomfort and maximize efficiency in the management of crowds. It performs an analysis of crowd behaviour to estimate the characteristics essential to understand its current state and predict its evolution. CrowdDNA is particularly concerned with the dangers and discomforts associated with very high-density crowds such as those that occur at cultural or sporting events or public transport systems. The main idea behind CrowdDNA is that analysis of new kind of macroscopic features of a crowd – such as the apparent motion field that can be efficiently measured in real mass events – can reveal valuable information about its internal structure, provide a precise estimate of a crowd state at the microscopic level, and more importantly, predict its potential to generate dangerous crowd movements. This way of understanding low-level states from high-level observations is similar to that humans can tell a lot about the physical properties of a given object just by looking at it, without touching. CrowdDNA challenges the existing paradigms which rely on simulation technologies to

analyse and predict crowds, and also require complex estimations of many features such as density, counting or individual features to calibrate simulations. This vision raises one main scientific challenge, which can be summarized as the need for a deep understanding of the numerical relations between the local – microscopic – scale of crowd behaviours (e.g., contact and pushes at the limb scale) and the global – macroscopic – scale, i.e. the entire crowd.

Local team: Julien Pettré (DR), Ludovic Hoyet (CR), Anne-Hélène Olivier (MCU), Charles Pontonnier (MCU), Thomas Chatagnon (PhD)

Website: <http://crowddna.eu/>

H2020 FET-Open INVICTUS (2020-2024). INVICTUS (Innovative Volumetric Capture and Editing Tools for Ubiquitous Storytelling) aims at delivering innovative authoring tools for the creation of a new generation of high-fidelity avatars and the integration of these avatars in interactive and non-interactive narratives. The project proposes to (i) exploit the full potential of recent volumetric motion capture technologies that consist in capturing simultaneously the appearance (shape, texture, material) and motion of actors using simple RGB cameras to create volumetric avatars, and (ii) rely on these technologies to design narratives using novel collaborative VR authoring tools. INVICTUS proposes the design of three innovative authoring tools. First, a tool to perform high-resolution volumetric captures of both appearance and motion of characters and that will enable their exploitation in both high-end off-line productions (film quality) and real-time rendering productions. This will ease high-fidelity content creation and reduce costs through less manual labor. Second, a tool to perform edits on high-fidelity volumetric appearances and motions, such as transferring shapes between characters, performing stylization of appearance, adapting and transferring motions. This will reduce manual labor and improve fidelity. Third, a story authoring tool that will build on VR interactive technologies to plunge storytellers in virtual representations of their stories to edit decors, layouts and animated characters, improving productivity and creativity. By demonstrating and communicating on how these technologies can be immediately exploited in both traditional media (films/animation) and novel media (VR + AR) narratives, the INVICTUS project will open opportunities in the EU market for more compelling, immersive and personalized visual experiences, at the crossroads of film and game entertainment.

Local team: Marc Christie (MCU), Adnane Boukhayma (CR), Ludovic Hoyet (CR)

Website: <https://invictusproject.eu>

H2020 ICT-25 RIA PRESENT (2019-2023). PRESENT (Photoreal REaltime Sentient ENTity) is a Research and Innovation project to create virtual digital companions — embodied agents — that look entirely naturalistic, demonstrate emotional sensitivity, can establish meaningful dialogue, add sense to the experience, and act as trustworthy guardians and guides in the interfaces for AR, VR and more traditional forms of media. There is no higher quality interaction than the human experience when we use all our senses together with language and cognition to understand our surroundings and — above all — to interact with other people. We interact with today’s ‘Intelligent Personal Assistants’ primarily by voice; communication is episodic, based on a request-response model. The user does not see the assistant, which does not take advantage of visual and emotional clues or evolve over time. However, advances in the real-time creation of photorealistic computer generated characters, coupled with emotion recognition and behaviour, and natural language technologies, allow us to envisage virtual agents that are realistic in both looks and behaviour; that can interact with users through vision, sound, touch and movement as they navigate rich and complex environments; converse in a natural manner; respond to moods and emotional states; and evolve in response to user behaviour. The objective of PRESENT is to create and demonstrate a set of practical tools, a pipeline and APIs for creating realistic embodied agents and incorporating them in interfaces for a wide range of applications in entertainment, media and advertising. The international partnership includes the Oscar-winning VFX company Framestore; technology developers Brainstorm and Cubic Motion; Europe’s largest certification authority InfoCert; research groups from Universitat Pompeu Fabra, Universität Augsburg and Inria; and the pioneers of immersive virtual reality performance CREW.

Local team: Julien Pettré (DR), Ludovic Hoyet (CR), Anne-Hélène Olivier (MCU), Claudio Pacchierotti (CR), Katja Zibrek (CR), Adèle Colas (PhD), Alberto Jovane (PhD), Pierre Raimbaud (PD)

Website: <http://present-project.eu/>

Participation in National Projects

ANR OPMOPS (2017-2020). OPMoPS (Organized Pedestrian Movement in Public Spaces) is a Research and Innovation project, where partners representing Civil Security Forces (CSF) cooperated with academic researchers to develop a decision support tool that can help them in both the preparation and crisis management phase of urban parades and demonstrations (UPD). The development of this tool was conditioned by the needs of the CSF, but also by the results of research into the social behaviour of participants and opponents. The latter, as well as the evaluation of legal and ethical issues related to the proposed technical solutions, constitute an important part of the proposed research.

Highly controversial group parades or political demonstrations are seen as a major threat to urban security, since the opposed views of participants and opponents can lead to violence or even terrorist attacks. Due to the movement of UPD through a large part of cities, it is particularly difficult for CSF to guarantee security in these types of urban events without endangering one of the most important indicators of a free society. The specific technical problems faced by the Franco-German consortium are: optimisation methods for planning the routes of UPD, transport to and from the UPD, planning of CSF people and their location, control of the UPD using fixed and mobile cameras, as well as simulation methods, including their visualisation, with special emphasis on social behaviour. The methods are applicable to the preparation and organisation of UPD, as well as to the management of critical situations of UPD or to deal with unexpected incidents.

Local team: Julien Pettré (DR), Armel Crétual (MCU), Ludovic Hoyet (CR), Anne-Hélène Olivier (MCU), Florian Berton (PhD)

Website: <https://project.inria.fr/crowdscience/project/opmops/>

Inria Associate Team – BEAR (2019-2021). BEAR is a collaborative project between France (Inria Rennes) and Canada (Wilfrid Laurier University and Waterloo University), dedicated to the simulation of human behaviour during interactions between pedestrians. In this context, the project aims at providing more realistic models and simulation by considering the strong coupling between pedestrians’ visual perception and their locomotor adaptations.

Local team: Anne-Hélène Olivier (MCU), Ludovic Hoyet (CR), Julien Pettré (CR), Armel Crétual (MCU), Florian Berton (PhD), Benjamin Niay (PhD)

Website: <https://sites.google.com/view/inriabearproject>

SCIENTIFIC, EDITORIAL AND ORGANISATIONAL ACTIVITIES

Organization Committees

International Conferences & Workshops

MIG 2021 Co-Conference Chair of the 14th ACM Motion, Interactions and Games, Switzerland (with Ronan Boulic, EPFL).

Frontiers in Virtual Reality: Co-guest editor for the Research Topic *Creating Lifelike Digital Humans*, 2020 (with Fabien Danieau, Philippe Guillotel, Steve Tonneau, Yajie Zhao).

SAP 2019 Co-Program Chair of the Symposium on Applied Perception 2019, Barcelona (with Douglas Cunningham, Brandenburg University of Technology).

CASA 2019 Co-Workshop Chair for the 32nd International Conference on Computer Animation and Social Agents 2019, Paris (with Kerstin Ruhland, Trinity College Dublin).

GRAPP 2019 Animation & Simulation Conference Area Chair for the 14th International Conference on Computer Graphics Theory and Applications, 2019, Prague.

SAP 2013 Co-Conference Chair of the Symposium on Applied Perception 2013, Dublin (with Betsy Williams Sanders, Rhodes College).

National Conferences & Workshops

WACAI 2021 Co-Program Chair for the 2021 Workshop on Affects, Compagnons Artificiels et Interactions (France)

WACAI 2020 Co-Program Chair for the 2020 Workshop on Affects, Compagnons Artificiels et Interactions (France)

AFIG 2017 Co-Program Chair for the 2017 French Computer Graphics days (j.FIG), Rennes.

Program Committees

International Conferences

ACM SIGGRAPH Asia, 2021

ACM I3D - Symposium on Interactive 3D Graphics and Games - 2021, 2022

ACM MIG - Conference on Motion, Interaction and Games - since 2014

ACM SAP - Symposium on Applied Perception - since 2012

ACM VRST - Symposium on Virtual Reality Software and Technology - 2014

IEEE AIVR - International Conference on Artificial Intelligence and Virtual Reality, 2018

IEEE VR - International Conference on Virtual Reality and 3D User Interfaces - journal track (2021)

GRAPP - International Conference on Computer Graphics Theory and Applications (GRAPP) - 2013 to 2019

International Workshops

MARCH - Modeling and Animating Realistic Crowds and Humans, IEEE AIVR Workshop, 2019.

VHCIE - Virtual Humans and Crowds in Immersive Environments, IEEE Virtual Reality Workshops - since 2018

Reviewing Activities

Journals: ACM Transaction on Graphics, ACM Transactions on Applied Perception, IEEE Transactions on Visualization and Computer Graphics, IEEE Computer Graphics & Applications, Computer & Graphics, Computer Graphics Forum, IEEE Transactions on Games, The Visual Computer, Computer Animation and Virtual Worlds, etc.

Conferences: ACM SIGGRAPH, ACM SIGGRAPH Asia, ACM SIGGRAPH Asia Courses, Eurographics, IEEE Virtual Reality, Pacific Graphics, IEEE Artificial Intelligence and Virtual Reality, Computer Animation and Social Agents, etc.

TEACHING

INSA de Rennes (Since 2017)

-*Animation, Motion, Gesture Recognition*, 5ème année Info, Lectures/Labs, (since 2017: 10h/year). With Eric Anquetil and Richard Kulpa.

Université Rennes 1 (Since 2015)

-*Motion for Animation and Robotics*, Master SIF, Lectures (2015: 20h; since 2016: 6h/year). With Marc Christie, Julien Pettré, Fabrice Lamarche.

-*Advanced Computer Graphics*, Master SIF, Lectures/Lab (since 2019: 8h/year). With Marc Christie.

Université Rennes 2 (Since 2018)

-*Ergonomy and Engineering: an introduction to VR prototyping*, Master 2 Ingénierie de l'Ergonomie et des Activités Physiques, Project, (since 2018: 24h/year). With Franck Multon.

Trinity College Dublin, School of Computer Sciences and Statistics (2011-2014; 75h)

-*CS7033 - Real-Time Animation*, IET Master, Lectures/Labs (2011-2014: 15h/year).

-*CS4052 - Computer Graphics*, 4th year CS, Labs (2012: 12h/year).

University of Winnipeg (Canada), Applied Computer Science Seminar Series (2012)

-*Invited Lecture: Using insights from perception to understand and drive human motion on virtual characters*

Teaching Assistant (Monitorat), Université Rennes 1, IFSIC (2007-2010; 260h)

-BSc: *Scientific Programming* (Mathématica; 60h), *Introduction to Programming* (Java; 98h), *Internship Mentoring* (4h)

-MSc: *Network and Distributed Applications* (Java, C; 74h), *Programming: Refresher's Course* (Java; 16h), *Algorithms* (C++; 8h)

AWARDS

Best Journal Papers Award - IEEE VR 2020: *Avatar and Sense of Embodiment: Studying the Relative Preference Between Appearance, Control and Point of View* (R. Fribourg, F. Argelaguet, A. Lécuyer, **L. Hoyet**).

Best Papers Award - ICAT-EGVE 2020: *Influence of Threat Occurrence and Repeatability on the Sense of Embodiment and Threat Response in VR* (R. Fribourg, E. Blanpied, **L. Hoyet**, A. Lécuyer, F. Argelaguet).

Part II
Research Activity

*“What a piece of work is a man, how noble in reason,
how infinite in faculties, in form and moving how
express and admirable, in action how like an angel,
in apprehension how like a god! the beauty of the
world, the paragon of animals.”*

— William Shakespeare, *Hamlet* (2.2.295-302)

1

Introduction

Representing human bodies and movements has been one of the fixation of humanity for millennia. From the first simple representations in primitive civilizations (e.g., Lascaux painting: the Well, 15,000-17,000 years old), to the elaborate Greek and Roman sculptures (e.g., The Discobolus of Myron, Laocoon and his Sons, 500-200 BC), or the detailed paintings and anatomical drawings of Leonard da Vinci (1452-1519), individuals have strived to replicate the human body, in stillness and movement. Over the last decades, we have quickly entered a new era where the technology now enables us to create more and more life-like digital human replica, that are changing our entertainment industries (e.g., movies, games), as well as many technology-driven connected fields (e.g., education, training, reeducation).

With the swift development of real-time graphics, these virtual humans have become a requisite to create always more life-like virtual worlds. For instance, in the movie and video game industries their use ranges from creating digital doubles of actors in computer-generated movies, including long-passed famous actors such as Peter Cushing in “Rogue One: A Star Wars Story” (2016), to populating digital worlds using several thousands of virtual characters. In these industries, the current trend is to create always more realistic characters, as well as more impressive scenes involving larger and larger groups of characters interacting in real-time with each other. The difficulty might even be said to be increased by more than a hundredfold in such real-time or interactive applications (e.g., video games), where characters also need to interact with users in a real-time manner, which requires to find the best trade-off between realism and computational needs. However, decreasing computational resources is also relevant when interactivity is not required (e.g., in movies) to lower production costs. This stresses the need for highly efficient methods to animate virtual humans, that are able to create believable motions, behaviours, interactions, as well as to easily scale depending on the number of characters displayed on screen.

This need for realistic virtual humans has also been increased by the recent enthusiasm for Virtual Reality, in part due to the massive dissemination of consumer-grade Head-Mounted Displays (HMDs), e.g., from Oculus, Samsung, HTC. In such applications, users immersed in virtual environments would expect to see virtual characters interacting naturally with them, with other virtual characters, as well as with the environment itself. These applications also raise novel questions, related to the virtual representation of the users themselves (i.e., their avatar). Ensuring that users embody in an avatar then becomes crucial for making people live a truly immersive and effective experience.

RESEARCH OBJECTIVES AND CHALLENGES

The general objective of my research activities is to **create more realistic virtual humans**, by first understanding how viewers perceive some of their key characteristics, and then drawing from this understanding to **produce visually plausible virtual humans**. More precisely, most of my research targets the **creation of visually plausible virtual human motions**, i.e., motions that viewers will consider to be plausible even though they might not be biomechanically or physically correct, which is a cornerstone of the research I participated in over approximately the last ten years. This topic is particularly relevant for Computer Graphics, where by focusing on what is perceived by viewers, it then becomes possible to save computation time, to only simulate a perceptually-plausible approximation of the motion, to save on production costs by identifying the minimum number of motions necessary to ensure visual variety, etc. It is also particularly relevant in the context of immersive Virtual Reality, where virtual humans are becoming increasingly crucial for creating life-like immersive experiences. Such virtual humans enable users to interact with other (virtual) persons, ideally as they would in real life. They are also important in the context of providing users' with a virtual representation (i.e., their avatar), which serves as a proxy between their real body and the virtual environment. However, creating such visually plausible human motions raises the following challenges.

Animating Individual Virtual Characters

Characters are most often driven either by motion capture or by animations manually created by animators. These animations are used to make them move and perform actions in the virtual scene, as well as to express their intentions, actions, internal state, etc. However, covering the whole repertoire of human movements, including intra- and inter-individual variations, is infeasible. This means that virtual characters are usually driven only by a subset of generic motions, structured according to different methods (e.g., Finite State Machines, Motion Graphs, Motion Matching).

Our first objective is to provide **novel ways of animating virtual characters that take into account how viewers perceive character motions**. This objective is particularly challenging, as humans can perform an almost infinite number of motions with large intra- and inter- individual variations. Identifying the characteristics that will influence visual realism therefore typically does not generalize well, and requires to perform numerous experimental studies, tailored to each research question explored. However, ensuring that some required characteristics are indeed perceived by viewers then enables us to create characters that are more expressive and perceived to be more human-like, as they are now able to convey their emotional state, some elements of their personality, or are now moving consistently with their individual characteristics (e.g., old vs. young characters, heavy vs. slim characters). Simultaneously, focusing on what users perceive enables us to better balance visual realism and computational costs, e.g., by identifying the best set of motions that will maximize visual plausibility for a specific task, or by considering the most relevant visual characteristics while synthesizing motions.

Interacting Virtual Characters

Complementary to displaying individual character motions that look plausible, in numerous cases we need to build virtual scenes involving a number of characters interacting together to create lively situations. This includes crowd simulation approaches, that focus on creating virtual characters that navigate in a virtual scene without colliding with other characters, as well as approaches focusing more on body-based interactions (e.g., pushing another character, wrestling, interacting with objects in the environment). In both cases, the difficulty lies in handling the high dimensionality of the full-body character kinematic model. Crowd simulation approaches therefore typically separate the computation of the character trajectories and interactions from the full-body animations, to compute collision-free trajectories in a lower dimension space (usually 2D position and 1D orientation), and therefore often do not account for the interactions between characters at the animation stage. Similarly, when close interactions are required between characters in interactive applications, specific animations are usually requested from a pre-processed motion database (one-to-one mapping), e.g., playing a specific pushing motion on one character and a specific falling motion on the corresponding target character, or playing a specific greeting motion when an event is triggered.

Our second objective is therefore to provide **novel methods for creating more natural interactions between virtual characters**, by focusing on viewers' perceptions of these interactions. One specific sub-objective is also to **create animated virtual crowds that look more natural**, because they display more varied interactions and animations. This is particularly challenging, as interactions involving virtual characters can become

extremely complex to quantify because of the high dimensionality of character kinematic models. Similarly, creating crowd interactions that look visually varied, while balancing computational costs, is also challenging. Exploring these challenges also relies on approaches to create individual motions with varying characteristics, as explored in the previous objective. Creating such varied character interactions then enables us to create virtual characters that look more plausible, as well as reactive, unlike typical approaches which often lead to passive characters displaying uniform motions.

Immersive Virtual Crowds

Creating more natural interacting virtual crowds requires to model these interactions, either in terms of crowd simulation or in terms of other types of physical or social interactions. Traditionally, the knowledge necessary to design such models is drawn from real-life experiments, which can be difficult to organize, to standardize and to replicate. Instead, immersing participants in a virtual environment using Virtual Reality (VR) has been considered to be promising for exploring more complex situations in an experimentally controlled environment.

Our third objective is therefore to **explore novel ways of studying interactions between people by immersing them in interactive virtual crowds**. This objective is challenging both at the technological and scientific levels. Technologically, numerous methods and devices need to inter-operate (e.g., crowd simulation, character animation, user interaction, navigation in VR), which requires expertise in multiple complementary domains. Scientifically, exploring these questions requires to propose novel methodologies to evaluate how people move, behave and interact in these virtual situations, so that our observations enable us to study increasingly more complex and ecological situations. It also raises the question of the realism of the immersive situation, which is also often challenging to evaluate: How consistent observations made in these immersive situations are with reality? Are we able to quantify some of these biases?

Avatars

Finally, the rapid development of consumer-grade HMDs has reinvigorated research on avatars, i.e., the user representation in the virtual environment, as users typically cannot see their real body when they equip themselves with these devices. However, their usage raises numerous technological and scientific questions: How to animate avatars so that their movements correspond to the user's movements? How should we represent the user? How should users interact with the environment with their avatar? Etc. This is particularly important as current avatars often fail at providing users with natural ways of interacting

with the virtual environment, as well as at conveying a strong sense of embodiment (i.e., making users feel as if the virtual body is their real body).

In this context, our fourth objective is to **create avatars that are better embodied and more interactive**. This is challenging, as it involves both a) acquiring novel fundamental knowledge, to understand how to make avatars efficiently convey a strong sense of embodiment to the user, as well as b) proposing novel methods for interacting more efficiently with the virtual environment through the use of the avatar, which includes proposing interaction techniques that positively influence the sense of embodiment. We believe that understanding through the concept of embodiment what are the factors that enable users to become more immersed in the virtual experience is paramount to create situations in which they will act as in real life, which also contributes to our third objective.

DOCUMENT OVERVIEW & MAIN CONTRIBUTIONS

In this section, we present a short overview of our main contributions related to the objectives mentioned above, which will be further detailed in the following chapters. These contributions are organised in four chapters, and their inter-connections are illustrated in Figure 1.1. It is important to remind the reader that most of these works were performed through the principle of understanding how viewers perceive virtual characters, to contribute to the more general objective of creating perceptually plausible virtual characters.

Perception of Individual Virtual Characters

As previously mentioned, our goal is to propose novel ways of animating virtual characters that take into account how viewers perceive character motions. Over the years, we have therefore conducted a number of studies exploring how viewers perceive biological human motion displayed on individual virtual characters, in order to identify the factors that influence users' perceptions of these characters. These contributions are presented in Chapter 2.

Ideally, to replicate real-life situations we would hope for each individual character in a scene to be unique, both in terms of appearance, motion, personality, etc. As this is definitely a utopia, because of computational and economical reasons, we have performed a number of experiments exploring **how viewers perceive individual characteristics of virtual characters**, which could then be controlled to create a variety of individual motions. First, we probed into the factors that make human motion recognizable and appealing [Hoyet et al. 2013], by exploring how distinctive and attractive were the walking, jogging and dancing captured motions of 30 different actors. We also explored whether viewers were able to visually identify a biomechanical parameter of human walking when




Figure 1.1: General framework of the research presented in this manuscript. Creating plausible virtual characters requires dedicated methods to animate them, which we explore by studying how viewers perceive characteristics of individual character motions (Chapter 2). We then explore approaches focusing on the perception of character interactions (Chapter 3), which are also crucial to the creation of life-like virtual scenes. The works presented in these two first chapters all contribute to the creation of lively immersive environments, which we combine to create immersive virtual crowds. These immersive virtual crowds can be exploited to study complex human behaviours that would be difficult to evaluate in real situations (Chapter 4). Our goal is to use insights and knowledge gained from such immersive experiences to iteratively improve our human animation models both at the individual and collective levels. Finally, in Chapter 5 virtual characters are studied from the standpoint of the users’ avatar, i.e., their virtual representation in these immersive situations, which we explore through the psychological construct of virtual embodiment.

displayed on virtual characters [Niay et al. 2020], namely the Walk Ratio, which was shown to be specific to each individual and constant for different speeds, but which had never been used to drive animations of virtual characters.

Concurrently, we have also been interested in evaluating existing animation techniques to understand the **perceptibility of motion artefacts** commonly encountered when animating virtual characters. We explored the perceptibility of footsliding artefacts [Pražák et al. 2011], and demonstrated that people perceive even low levels of footsliding ($< 2cm$). Similarly, we found that corrected animations were always preferred by viewers, and that a simple approach of lengthening limbs was preferred to a more complex approach. We are also interested in proposing novel methods for animating virtual characters, with the goal of providing ways to **balance visual realism and computational complexity**. We investigated the perception of hand animations computed from a number of reduced marker sets [Hoyet et al. 2012b], and proposed a number of guidelines for balancing visual realism with the complexity of simultaneously capturing body and finger motions.

Perception of Interacting Virtual Characters

Interactions between characters are important to create life-like virtual scenes, as they enable us to create characters that react to their surrounding, as well as to events that might happen to them. Our goal is therefore to propose novel methods for creating more natural interactions between virtual characters, which we explored by studying how viewers perceive virtual character interactions. These contributions are presented in Chapter 3.

To better understand how viewers perceive physical interactions between characters, we studied the factors that affect the **perception of causality in human interactions**, i.e., the impression that the action of one individual causes an appropriate reaction in another person. In particular, we investigated the effects of timing, force, and angular anomalies on the perceived plausibility of two-character interactions [Hoyet et al. 2012a]. We explored similar questions in the context of one character interacting with its environment, to understand how viewers perceive modifications of throwing animations introduced by state-of-the-art editing methods [Vicovaro et al. 2012; Vicovaro et al. 2014], where the throwing motion of the character and the trajectory of the thrown object were edited simultaneously or separately.

We also explored the **perception of interactions in the context of crowd animations**, where the quality of interactions between characters is paramount to creating visually compelling virtual crowds. As these approaches typically focus on collision avoidance, for computational reasons they separate the computation of the character trajectories and interactions from the full-body animations. This often leads to residual collisions between the character meshes, as well as to characters walking as if they were alone, showing no sign to the influence of others. To solve some of these issues, we conducted two perceptual experiments investigating the value of adding shoulder motions to characters passing at close distances on the perceived visual quality of crowd animations, in terms of both perceived residual collisions and animation naturalness [Hoyet et al. 2016b]. We demonstrated that such a strategy has a significant impact on the visual quality of crowd animations, with a very light computational overhead as interactions are handled at the animation stage. To further explore means of reducing the uniformity of crowd animations caused by the small number of motions typically used, we also conducted an experiment on the perception of motion variations in large-scale virtual human crowds (250 to 1000 characters) [Adili et al. 2021]. We relied on a statistics-based method for creating motion variations, also exploring whether they can be used to increase the visual variety and realism of such large-scale crowds. Surprisingly, our results suggest that even for large-scale virtual crowds, it is very difficult for viewers to perceive that the same motions are used multiple times, even when only two motions (one male and one female) are used to animate the whole crowd.

Perception of Virtual Crowds in Immersive Virtual Environments

To further our knowledge of human behaviours, we leverage VR to immerse participants in immersive virtual crowds and study their behaviours in complex situations that would sometimes be difficult to replicate in real-life. In Chapter 4, we present two main contributions on this topic, which are the result of the PhD of Florian Berton, whom we co-supervised with Julien Pettré and Anne-Hélène Olivier.

First, we performed three experimental studies to **understand pedestrians' gaze activity** when interacting with other people, as a mean of studying their interaction neighbourhood, i.e., which persons/characters within the crowd influence their motion. This is particularly important to understand, as a better definition of this neighbourhood could lead to improvements in crowd simulations, as most of the current agent-based crowd simulation models highly rely on this concept. We therefore first explored the impact of VR on gaze activity in a pairwise collision avoidance task, by carrying out an experiment in a real and a virtual environment with four different VR setups [Berton et al. 2019]. We then extended our study by exploring gaze activity during navigation in a crowded street by comparing real and virtual crowded situations, before investigating the **impact of crowd density on gaze activity** in VR to gain new insights on the interaction neighbourhood [Berton et al. 2020b].

Second, as we are relying on VR to study users' behaviour, we are also interested in **improving user immersion when interacting with virtual characters**. Our hypothesis is that by improving immersion, we are also increasing the probability that users will behave in the immersive situations as they would in real life. To this end, we explored the **use of wearable haptics to render contacts with virtual characters** when navigating in a virtual crowd [Berton et al. 2020a]. We showed that providing haptic feedback improves the overall realism of the interaction, as participants more actively avoid collisions.

Avatars

In Chapter 5, we then present our last contributions which aim at creating avatars that are better embodied and more interactive. Most of them focus on furthering our understanding of the sense of virtual embodiment, and are the result of the PhDs of Rebecca Fribourg and Diane Dewez, whom we co-supervised with Ferran Argelaguet and Anatole Lécuyer.

First, we conducted several experiments studying the **factors influencing the sense of embodiment**. In our initial contribution [Argelaguet et al. 2016], we explored the influence of the realism of the virtual hand representation, and showed that realism influences differently the sense of agency and the sense of ownership. Following studies then focused on full-body avatars, and evaluated the **relative importance** of external factors (i.e., related to the virtual environment and the avatar) commonly found in the literature to influence the sense

of embodiment (namely the avatar’s appearance, the avatar’s control, and the user’s point of view) [Fribourg et al. 2020a], as well as the influence of internal factors (i.e., related to the users themselves), namely personality traits and body awareness of participants [Dewez et al. 2019]. Simultaneously, we explored **methodologies for measuring the sense of embodiment**. We demonstrated that the commonly used introduction of a threat does not alter users’ sense of embodiment [Fribourg et al. 2020b; Fribourg et al. 2021a]. We also proposed to rely on a novel methodology (subjective matching technique) to better understand the **inter-relations between factors** influencing the sense of embodiment [Fribourg et al. 2020a].

Concurrently, we studied **inter-relations between interaction techniques and the sense of embodiment**, e.g., to explore whether some techniques to control an avatar might negatively affect the sense of embodiment. First, we compared three state-of-the-art VR locomotion techniques (real walking, walking in place, or steering) [Dewez et al. 2020], and found that participants had a comparable sense of embodiment with all techniques. More recently, we also explored whether a dual avatar representation might be beneficial during anisomorphic manipulation [Dewez et al. 2022], to compensate for the mismatch between the performed and the displayed action which can negatively impact the sense of embodiment. While we found that users preferred having a single representation, embodiment scores were however similar between single and dual representations. We also explored the idea of designing “Avatar-Friendly” techniques, i.e., techniques which take the user’s avatar into account in the design process to preserve both user performance and sense of embodiment [Dewez et al. 2021].

Finally, we also proposed **novel paradigms of interacting through the avatar**. We proposed a novel experience of controlling a six-finger virtual hand [Hoyet et al. 2016a], and demonstrated that participants responded positively to the possibility of controlling the virtual hand despite the structural difference. We also recently proposed the novel concept of “virtual co-embodiment”, which enables users to share their virtual avatar with another entity (e.g., another user) [Fribourg et al. 2021b], and demonstrated that the users’ level of control over such a shared avatar, as well as their prior knowledge of their actions to perform, have an influence on their sense of agency.

Perspectives

Following the four chapters related to our contributions to the creation of more realistic virtual humans, Chapter 6 concludes this manuscript by providing general perspectives and future works than we plan to conduct.

2

Perception of Individual Virtual Characters

CONTEXT AND OBJECTIVES

When animating virtual humans, one of the first objectives is that they appear, move and behave similarly to a real human. For instance, we need characters that display motions that are adapted to their characteristics (e.g., gender, age, morphology), that are able to express and communicate intentions as well as emotions to the viewer, that act and behave according to their surroundings (e.g., environment, other characters, events), etc. However, because of the complexity of human motion, creating automatically such characters is still challenging, and is therefore mostly performed manually in large-scale productions.

To tackle this challenge, a number of animation techniques have been proposed over the last 30 years. Such approaches can be roughly divided into procedural, physics-based or example-based approaches. Among these approaches, example-based methods are probably the most commonly used today, as they facilitate capturing the naturalness of biological human motion. For instance, parts of or entire motion examples can be combined to create new complex sequences using Motion Graphs [Kovar et al. 2002; Min and Chai 2012], Finite State Machines [Wang and Bodenheimer 2008], or even Motion Matching [Buttner 2015], a recent state-of-the-art solution now currently used as an industrial standard. A number of approaches also rely on learning statistical models, especially when dealing with variations or style [Hsu et al. 2005; Lau et al. 2009], or interpolate or extrapolate existing styles in low dimensional spaces [Urtasun et al. 2004; Min et al. 2010]. More recently, a current trend started to rely on deep learning in the context of character animation, in order to learn human motion manifolds from data using compact models (e.g., [Holden et al. 2016; Holden et al. 2020; Starke et al. 2021]). However, none of these approaches account for

how users perceive resulting motions, which might lead to unnatural results. Evaluating the visual realism of these approaches is therefore highly important to provide the best trade-off between computation costs and visual realism.

Before, and concurrently to, the development of animation techniques, perception of biological body motion has been a very active field of research. In his pioneer work, Johansson [1973] developed a special stimulus known as the “point-light walker”, providing a simple visual representation leaving only motion cues from the main parts of the body visible while removing most of the structural information. Such stimuli were for instance used to demonstrate that enough information is available in the motion signal for recognition of gender [Pollick et al. 2005], of a particular person [Cutting and Kozlowski 1977], or one’s own walking pattern [Beardsworth and Buckner 1981]. Realistic virtual characters have since been demonstrated to provide more visual information about how motion is perceived [Chaminade et al. 2007], and are therefore now extensively used to explore the perception of biological human motion on virtual characters (e.g., [Hodgins et al. 2010; McDonnell et al. 2012; Zell et al. 2015; Durupinar et al. 2016]).

With the growing need for virtual humans in interactive applications, understanding the factors that affect our perception of such characters is becoming even more crucial. This chapter presents our research activities on this topic. It focuses on the perception of ‘individual’ human characters, i.e., where interactions with the environment or other characters do not influence the movement of characters. This complementary aspect is developed separately in Chapter 3. We have therefore conducted over the years a number of experiments to explore which factors influence our perception of virtual characters. Our contributions on this topic can be divided into two main categories.

- We conducted a number of experiments exploring **how viewers perceive individual characteristics of virtual characters**. In particular, we are interested in being able to create virtual characters that display specific individual characteristics, visually identifiable by viewers, which is then important for creating visual variety in large groups of characters. For instance, we explored the factors that make human motion recognizable and appealing [Hoyet et al. 2013], by exploring how distinctive and attractive were the walking, jogging and dancing captured motions of 30 different actors. We also demonstrated that the Walk Ratio, and invariant biomechanical parameter specific to each individual and constant for different speeds, is identifiable by viewers when animating characters from motion capture data [Niay et al. 2020]. As we are also interested in creating visually expressive characters, i.e., able to convey some wanted information to the viewer, we also explored how emotions were perceived by viewers from virtual character motions, e.g., by investigating the effects of displayed emotions on gender judgments [Zibrek et al. 2013; Zibrek et al. 2015].

- Concurrently, we also worked on **evaluating and proposing perception-driven animation techniques**, where the resulting animations are probed in terms of visual plausibility. More precisely, we explored the perception of footsliding artefacts [Pražák et al. 2011], which are commonly encountered when animating virtual characters. In our experiments, we determined minimal perceivable footsliding thresholds, and showed that participants can perceive even very low levels of footsliding, especially when environment cues highlight the artifacts. We also demonstrated that viewers always preferred when perceivable footsliding levels were corrected, despite the potential side effects of the correction method, but that a simple approach of lengthening limbs was usually preferred to a more complex state-of-the-art approach. We also explored novel ways of simultaneously capturing body and finger motions to produce more natural animations [Hoyet et al. 2012b]: we investigated the perception of hand animations computed from a number of reduced marker sets, and proposed a number of guidelines for balancing visual realism with the complexity of simultaneously capturing body and finger motions.

The remaining of this chapter details two of these contributions. Focus 1 presents an example of our experimental processes, focusing on the question of how distinctive and attractive individual motions are perceived by viewers [Hoyet et al. 2013]. Then, Focus 2 presents some work where we investigated the perception of finger animations computed from reduced marker sets [Hoyet et al. 2012b].

FOCUS 1: DISTINCTIVENESS AND ATTRACTIVENESS OF VIRTUAL HUMAN MOTIONS



Figure 2.1: Example frames from distinctive (leftmost pairs) and non-distinctive walking, jogging and dancing motion clips from our experiments [Hoyet et al. 2013].

In this work, we were interested in understanding the factors that make human motion recognizable and appealing, which is of great value in industries where creating a variety of appealing virtual characters with realistic motion is required.

Typically, animating large groups of characters involves using a limited number of motion capture examples over a larger number of characters, because of computational and production costs. Such a limitation in the number of motions introduces uniformity that can be detrimental to the overall visual realism of the scene. For instance, the style of a person's motion can be quite distinctive, and therefore easily recognized when applied to one or more characters (e.g., in crowds). Similarly, it can also be undesirable to use motions that might be unappealing or unattractive to some or all of the audience. For this reason, understanding how individual motions are perceived by viewers, such as in terms of how distinctive or attractive one's motion is, is therefore highly important to later on animate larger groups of characters.

Research Questions

To better understand how to create natural and sufficiently varied motions given a limited repertoire of human motions, we were interested in the following questions. *How distinctive and attractive are the motions of different actors when shown on exactly the same realistic virtual body?* As mentioned above, this question is particularly important for the overall visual realism when animating large groups of characters. Most of the research on the recognition and attractiveness of biological motion was performed using simple representations, e.g., point-light displays [Johansson 1973; Troje 2002]. While providing very useful insights for perception, it is also valuable to know how realistic human motion is perceived in more ecologically valid situations, such as on a realistic 3D character. Furthermore, the stimuli used often contain information about body shape which can highly influence how distinctive or attractive a person is. We are therefore also interested in exploring the effects of removing cues related to body shape that could affect recognition and perceived attractiveness as much as possible. *Are more distinctive actors less attractive? Is an average motion the least easy to recognise, and more attractive than the other motions, as for faces?* Previous studies on the perception of faces have demonstrated that distinctiveness of faces correlates negatively with facial attractiveness, and that distinctive faces can be perceived to be less attractive than average and/or symmetric faces [Rhodes 2006]. As creating attractive motions can be valuable, we are interested in exploring whether such results extend to synthetic average body motions. *Are all gaits equally easy to recognise? Is a person equally attractive when performing different gaits?* It has previously been found that humans find it difficult to distinguish between the motions of multiple walking people [McDonnell et al. 2008; Pražák et al. 2011], but it is not clear if this is true for other gaits and actions apart from walking. Also, while there have been some studies with simple stimuli on the attractiveness of human motion [Grammer et al. 2003; Johnson and Tassinary 2005; Johnson and Tassinary 2007], the perception of realistic virtual characters performing a variety of different human motions has not been performed. *Are there cultural or gender differences with respect to recognition?* With the

global reach of the games, movie and other industries that deploy virtual characters in their products, it is important that such characters are appealing to all audiences. It is therefore important to take cultural and gender issues into account when exploring such questions.

Methods

To understand these issues, we ran a set of experiments to explore the distinctiveness and attractiveness of virtual humans walking, jogging and dancing. Motions were captured from 15 male and 15 female Caucasian European actors, and retargeted to a single realistic female or male model (Figure 2.1), thereby removing any cues as to body shape that could affect recognition and perceived attractiveness. Actors were reasonably similar in age and body shape, to minimize retargeting errors. Walking and jogging motions were recorded after training actors to walk or jog at a specific gait frequency using a metronome, to prevent the gait frequency from influencing the distinctiveness of each actor. Dancing motions were recorded by asking all the actors to perform the same dance from a 30s video clip. We also took particular care to avoid introducing any motion artefacts, and tried to interfere with the original motion as little as possible even though some motion processing was required for the experiments (e.g., creating seamless cyclified locomotions).

All our experiments therefore included 15 female and 15 male walking, jogging, and dancing motions, as well as one additional male and female average motion for walking and jogging. These average motions were created by averaging over the corresponding 15 male and 15 female actors' motions. Despite that all actors were performing the same dance sequence, the amount of variability due to each individual's dancing style made it however impossible to create natural average dancing motions, which we therefore did not explore in this work. For all our experiments, male and female actors were then displayed in separate blocks, in counterbalanced order. Furthermore, to test for cultural effects, for all experiments 50% of participants were Caucasian European and 50% were Korean.

First, we explored the **distinctiveness** of the actors' different gaits, by determining how well participants could remember whether each actor was present or absent from a group of three others. Taking inspiration from previous work in shape recognition [Fugard et al. 2011], participants first watched a 5s clip of three different actors walking, jogging or dancing (Figure 2.2.a). Participants then viewed a single actor performing the same gait for up to ten seconds, and were asked to indicate whether the single actor was present or absent in the previous group of three. The target actor was present or absent in an equal number of trials, while the remaining characters were randomly selected amongst the other actors. The "present or absent" task helped us to answer the question: "*is the person distinctive enough to be remembered?*" Similar experimental designs are common in the field of shape recognition [Fugard et al. 2011], and we chose this task to avoid simple



Figure 2.2: a) Examples of the distinctiveness stimuli where 3 different actors were presented on screen first (left), then a single actor (right). Participants indicated whether the single actor was present or absent in the group of three. b) Example of the Cross-gait stimuli where participants were asked to rate how likely it was that the motions of the characters were from the same actor.

matching between motions, preferring instead to pose a true signal detection/recognition challenge. We then probed distinctiveness in terms of the sensitivity of each participant to the presence or absence of each actor performing each gait, by computing their corresponding d' -prime value using Signal Detection Theory, which is commonly used in psychophysical studies to measure sensitivity to a signal.

Then, we performed a **cross-gait analysis** experiment, to explore whether it was possible to recognize when the same actor was performing pairs of side-by-side gaits: walk/jog, jog/dance, or walk/dance. To evaluate whether cross-gait recognition is affected by how distinctive actors are for each gait, we selected four groups of four actors (2M, 2F) for each gait combination using the results of the distinctiveness experiment: actors who are distinctive in both gaits; distinctive in one gait, non-distinctive in the other; vice-versa; and non-distinctive in both gaits. Participants saw two characters on screen, displaying one motion of each combination (see Figure 2.2.b), either performing motions from the same actor or from different actors. Each actor pair was presented for a maximum of 10s, and participants were asked to rate how likely it was that the motions were from the same actor, using a Likert scale ranging from 1 (very unlikely) to 7 (very likely).

Finally, we were interested in determining the perceived **attractiveness** of the motions of our thirty actors and their average motions, and to explore how these ratings related to their distinctiveness. To this end, participants were presented with a single character on screen, walking, jogging or dancing, and were asked to rate the attractiveness of the character's motion on a Likert scale from 1 (very unattractive) to 7 (very attractive).

Main Results

In terms of distinctiveness, our first experiment showed that **motions are not equally as easy to recognize depending on the gait**: dancing motions were easiest to recognise, followed by jogging and finally walking. While previous research has demonstrated that cloned walking motions can be difficult to detect in groups of characters, this result suggests

that viewers' ability to distinguish such motion clones might depend on how distinctive, and therefore easier to recognize, the type of motion is. Also, we found that female motions were on average more distinctive than male ones, except for dancing for which they were equally recognizable to male's. However, we also found that **distinctiveness in one gait does not transfer to other gaits**, as it does not predict how recognizable the same actor is when performing a different motion. The exceptions to this is that **average motions were always amongst the least distinctive**.

The question of whether the perceived distinctiveness affected the perceived similarity of an actor's different gaits was further explored in the result of the cross-gait experiment, where we found that it was very difficult to tell whether the same actor was performing different motions or not, suggesting that **style characteristics do not transfer well across gait**. There was some limited evidence that some motion characteristics may transfer between walking and jogging, but only when the walk was distinctive but the jog non-distinctive (for all four actors in this category). However, further investigation would be required to explore in more details what features might be transferred, and why they occur in this particular condition. Nevertheless, the implications of the general results for industry could be useful, in that it may not be necessary to capture the motion of as many actors for group or crowd scenes, as long as multiple different gaits of the same actor are being simultaneously displayed.

In terms of attractiveness, we found that **there was a negative correlation between attractiveness and distinctiveness for walking and jogging**, but not for dancing. As predicted, **average motions were perceived to be most attractive** and always amongst the least distinctive, which mirrors previous results observed on the perception of average faces. This is encouraging for application areas where the time for capturing and processing motions is severely limited, and yet appealing characters is very important for user engagement. An average motion could potentially be used far more frequently, especially if it could be parameterized in some way to create style variations. The developer could be assured that such motions would then be more appealing to the target audience.

Regarding cultural effects, we found some small differences, suggesting that cultural and/or familiarity effects might be at play when identifying actors, which is little studied in the context of creating characters appealing to all audiences.

Conclusion

In this work, we explored the distinctiveness and attractiveness of human motions displayed on virtual characters, which showed that different actors, and different gaits, are perceived differently by viewers. Understanding such effects is a useful contribution for creating realistic and engaging virtual scenarios. For instance, it can be interesting to be able to create and use distinctive motions, e.g., for a hero or villain character in a game or movie, as well as to be

able create less distinctive motions, as motions that are too easily recognized might detract from the perceived realism, e.g., when displaying an animated crowd. Similarly, attractive characters could facilitate the engagement of a user with a movie, game or interactive experience. Such results also raise novel interesting questions to further explore these effects, as well as to take them into account in novel animation techniques, e.g., by identifying which motion properties contribute most to the attractiveness or distinctiveness effects, or by exploring how body shape and motion interact to affect our perceptions.

FOCUS 2: PERCEPTION OF FINGER MOTION




Figure 2.3: Using Inverse Kinematics with 8 markers per hand enables us to simultaneously capture full-body and hand motion while retaining most of the fingers’ perceived information. This figure depicts side-by-side comparisons of real movies and computer generated animations using our approach presented in [Hoyet et al. 2012b].

In this work, we were interested in another aspect of creating visually plausible character motions, concerned with our ability to acquire and include subtle animation details, and more specifically finger motions.

Gestures and finger motions play an important role in our everyday actions: we use our hands to interact with the environment as well as to convey and emphasize information in conversations. Adding such details to computer generated animations can therefore help to increase the naturalness, credibility and appeal of virtual characters. To this end, finger motions are typically captured using datagloves, optical systems, or video-based (RGB or RGB-D) approaches. However, because of a number of constraints limiting the capture area (e.g., wires, small optical markers), and therefore the actions that can be performed, hands and bodies are often captured during separate sessions and later spliced together [Majkowska et al. 2006; Jörg et al. 2012].

However, humans rely on synchronization cues when perceiving communication and gestures [McNeill 2005; Giorgolo and Verstraten 2008]. Breaking this synchrony, either between auditory and visual information or between different visual channels (hands, arms, body or facial movements), can reduce the credibility of virtual characters, especially as such

anomalies are easily detectable on virtual characters [Jörg et al. 2010]. While this problem has been partially mitigated with the development of novel wireless datagloves since the publication of this work, the ability of capturing simultaneously body and finger animations in a unified manner using a single system still has a number of advantages. For instance, capturing multiple channels simultaneously ensures that the synchrony of the actions is preserved, while synchrony can be impaired when using different capturing devices, e.g., because of latency or synchronization issues. Similarly, equipping actors with several different devices can be more tedious and cumbersome, as well as more technically complex (e.g., separate calibrations, running several software concurrently, etc.).

In the case of optical systems, which are probably still the systems most used in communities relying on virtual humans, capturing simultaneously fingers and bodies can sometimes be extremely challenging, especially for actions requiring a large capture space or when several actors are involved. Capturing fingers requires moving the cameras closer to the actor to increase the resolution of the projected image of small markers (< 6mm), thereby drastically reducing the capture space. Simultaneously, it introduces about 20 additional markers per hand (vs. approximately 50 larger – 14 to 20mm – markers for full-body capture), leading to numerous occlusions, labeling errors, and a major increase of the manual workload on animators. Because of these constraints, hand animations are often ignored or greatly simplified during busy production schedules for games or movies.

Research Questions

To overcome these limitations, we were interested in exploring the following questions. *Is it possible to use standard animation techniques to animate hands from reduced marker sets?* Relying on standard animation techniques and reducing the number of markers required to capture plausible hand movements has several potential benefits. It increases the distance between markers, which allows the use of larger markers to simultaneously capture hands and bodies when the number and resolution of cameras are constrained. It also greatly reduces the number of occlusions and labeling errors, thereby drastically reducing manual post-processing time. *What is the influence of not capturing all the fingers individually?* In a lot of everyday motions, especially talking, hand movements do not always exhibit independent motions between fingers. We were therefore interested in exploring whether we could drive hand animations by capturing only the index and pinky fingers, and linearly interpolating the joint angles of the middle and ring fingers. *Is there a benefit in using a well selected static hand pose when finger animation is not available?* As game companies often use static hand poses selected from a database instead of fully animating characters' hands, we also wished to evaluate whether such static hand poses could be perceived to be the same as a high quality hand animation in some specific situations.

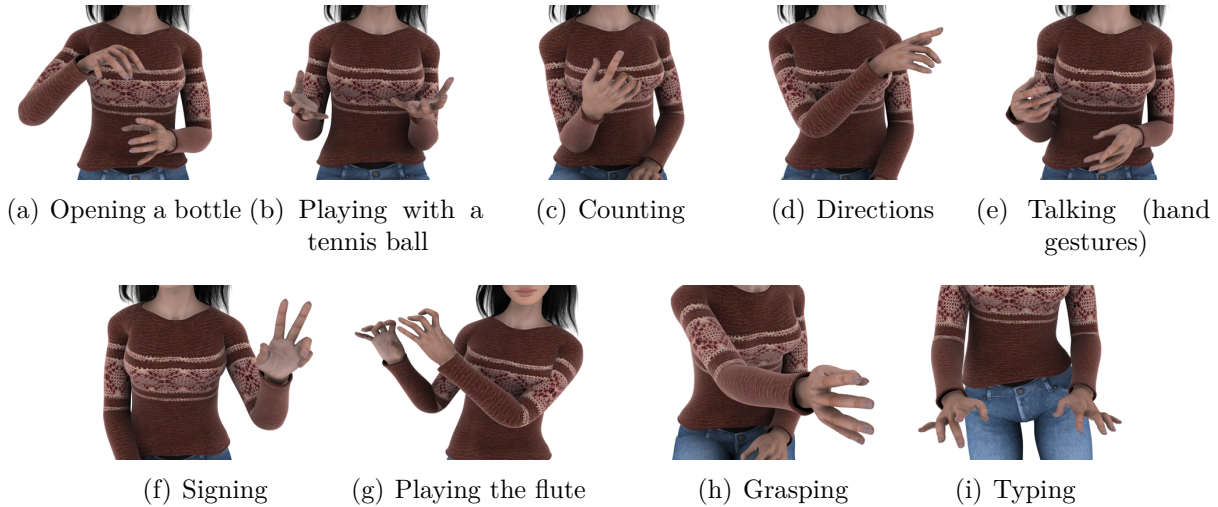


Figure 2.4: Scenarios captured for the experiments presented in [Hoyet et al. 2012b].

Methods

To explore these questions, we conducted two experiments designed to evaluate the perceived fidelity of finger animations created from a range of marker sets using traditional Inverse Kinematics (IK) methods, compared to the corresponding high quality animations generated using Forwards Kinematics (FK) with a full marker set. Our aim was to find an optimal trade-off between perceived fidelity and the number of markers used to capture finger motion.

Before conducting our experiments, we recorded simultaneously the full-body and finger movements of one non-professional female actor. Capturing the small markers positioned on the fingers ($4mm$) required us to carefully position our motion capture cameras around a small area ($\approx 1.5m \times 1.5m$). To cover a large range of finger motions, we captured nine scenarios (see Figure 2.4), including different levels of finger coordination and velocity, four of the six functional grasps for daily living [Edwards et al. 2002], as well as motion sequences involving both hands where each finger was moved independently.

To animate the virtual character’s hands, we used a skeleton similar to those often used in the literature, consisting of 4 DoF (2+1+1) for each of finger and the thumb. We however used a higher quality skeleton for the creation of the gold standard animation (GS), replacing the 2 DoF (flexion and abduction) joint with a 3 DoF joint for the finger’s Metacarpophalangeal (MCP) joint and the thumb’s Carpometacarpal (CMC) joint, as we found that it produced gold standard animations capturing smaller and subtler tilting due to the movements of the metacarpal bones. Gold standard finger motions were then generated using Forward Kinematics to drive each of the fingers and the thumb, using their corresponding four markers.

We also designed different models that mapped decreasing numbers of markers to the virtual hands (Figure 2.5). These simpler models are based on IK with only two markers




Figure 2.5: The different sets of marker used in our experiments. x F- y T stands for x markers used for the animation of the four Fingers and y markers used for the animation of the Thumb.

used to drive the animation of one given finger/thumb, complemented with biomechanical data to estimate each phalange’s relative length. Because of the illposedness of the IK problem, we also relied on a joint angle dependency discovered early in hand animation between the Proximal Interphalangeal (PIP) and the Distal Interphalangeal (DIP) joints to remove the DIP joint DoF [Rijpkema and Girard 1991]. Furthermore, in some models the middle and ring fingers were not driven by markers, in which case their joint configurations were computing by linear interpolation using the joint angles of the index and pinky fingers. For each captured motion, we also manually selected a static hand configuration, from the corresponding GS motion, that conveyed information about the desired action.

Two successive experiments were then conducted to evaluate the perceived fidelity of the finger animations created using these reduced marker sets or using a static hand pose. In our experiments, participants viewed pairs of movies, where the first movie of each pair was always the reference GS animation and the second movie displayed an animation created using one of our methods. Participants were instructed that the first movie was always the GS animation, in order to evaluate the faithfulness of the second movie, and were asked to answer if the second movie was the same as or different from the first movie. This protocol enables us to directly compare two movies and to ask participants to give a binary answer (*same* or *different* in our case). To test the difficulty of the task, participant accuracy, and to avoid a response bias, we also included one repetition of the GS animations against themselves, to which we expected participants to respond as “same”.

Main Results

First, we found that **finger animations whose motions of the four fingers were computed using IK were considered to be perceptually similar to the GS animations**, in most cases. This was true for all the models 8F-4T (69% similar to the GS), 6F-4T (69%) and 6F-2T (65%). In comparison, participants considered the GS to be the “same” as itself only 81% of the time (first experiment) and only 77% of the time (second experiment),

which shows the difficulty of the task. Also, **when independence between fingers is not important, simpler models perform relatively well**. In particular, only capturing the thumb, index and pinky fingers (4F-2T and 4F-4T models) produced reasonable results in our examples of talking, grasping an object or playing with a tennis ball.

Interestingly, while the thumb is much more complex than the other fingers from a biomechanical point of view, we found that **using 2 or 4 markers for the reconstruction of the thumb did not produce significantly different results**. It is also possible that this may have been due to the saliency of the fingers, rather than the quality of the thumb, which would require future studies to evaluate the relative saliency of these different areas of the hand.

Finally, **using a static pose can produce natural animations in certain circumstances**. We found that the absence of finger movements using a static hand pose can still lead to a high perceptual similarity to captured animations when these animations do not display a lot of finger movement, such as in the case of direction, grasping and signing. This was surprising as we felt that participants would notice the absence of subtle secondary motions present in the GS.

Conclusion

In this work, we showed that people are not highly sensitive to all the subtle details of finger animation. Using IK with a small number of markers produced animations that were perceived to be very similar to those generated using FK with a larger number of markers. From our observations, we compiled a list of guidelines that we believe will help designing finger motion capture sessions:

- For the majority of cases, we recommend using the simple 8 marker hand model (6 for the fingers, 2 for the thumb), which produces sufficiently high quality motions. This is particularly useful when simultaneously capturing full-body and finger movements in large capture areas, as it allows the use of large markers.
- When finger curvature is highly important, IK may flatten the fingers too much. In this case, we recommend using FK with a full set of markers in a small capture area.
- When independence between fingers is not important and processing time is limited, capturing only the thumb, index and pinky fingers (2 markers each) will produce reasonable results.
- If the fingers are only displaying secondary motion, a well chosen static pose may be adequate.

CHAPTER CONCLUSION

This chapter presented some of our contributions related to the perception of virtual character motions, with the goal of understanding, and taking into account, how viewers perceive motion characteristics to create more believable and plausible virtual characters. While the presented research explored specific aspects of the problems encountered when animating virtual characters, we have applied similar methodologies and approaches to a larger variety of questions and problems, including the perception of emotions on virtual characters [Zibrek et al. 2013; Zibrek et al. 2015], multi-modal animation such as exploring the effects of synchronization of body and faces on character expressivity [Ennis et al. 2013], perceptual evaluations of motion artefacts such as footskating [Pražák et al. 2011], etc.

However, despite the unquestionable advances made in the field over the last decades, the question as to what makes virtual human motions visually plausible still remains. Perhaps most notably our required efforts in this regard can be observed in animations resulting from physics-based approaches, which are still without question unmistakable with biological human motion, despite the incredible advances that have been made in this branch of character animation in the last decade. For this reason, most state-of-the-art approaches still rely on motion capture or manual artist labour to create expressive characters. Identifying the characteristics and principles that make virtual human motions expressive, and visually plausible, would therefore be extremely valuable to bring virtual characters to a next step of realism. Such knowledge could be for instance related to the definition of metrics for visual plausibility, which could supersede traditional character animation metrics which rely on purely numerical comparisons (e.g., numerical errors, jerk). Despite a few attempts in this direction [Tang et al. 2008; Durupinar 2021], we are still lacking relevant perceptual metrics to evaluate how plausible animations of virtual characters are.

Also extremely important in this direction of visual plausibility is our ability to synthesize motions that are adapted to the characteristics of each character. Most applications rely on a limited number of animations, without ensuring that these motions are indeed adapted to the characters, e.g., that a motion would be appropriate for a young dynamic character, or a tall and heavier older character. This problem also raises the question of the characteristics that govern intra- and inter-individual variations representative of different individuals (e.g., morphology, sex, age). An interesting aspect to identifying such characteristics is also to explore motion parameters that have been studied and found relevant in other fields. For instance, relying on biomechanical parameters of human motion, and validating that these parameters indeed have an effect on the visual realism of virtual characters is extremely valuable to create plausible animations. This is one direction that we are exploring in the PhD of Benjamin Niay, who demonstrated that an invariant parameter of human walking

named the Walk Ratio (step length to step frequency ratio), which was shown to be specific to each individual and constant for different speeds, is indeed perceived by viewers and relevant for character animation [Niay et al. 2020]. We are currently exploring how this new parameter might be used to drive the animation of walking virtual characters. In particular, one of our goal is to create variety in the motion of different individuals depending on their personal characteristics, as well as to provide animators with novel means to create natural walking animations through the use of simple parameters.

Concurrently to improvements in character animation techniques, the visual appearance of virtual characters has also drastically progressed in recent years (e.g., MetaHuman characters¹). However, mismatches appear if the quality of virtual character’s motions or behaviours are not on par with their appearance. For instance, virtual characters might have extremely realistic faces, but still elicit the same eeriness if their face is not animated, or if their eyes do not move. Including more subtle motions and behaviours, such as fingers, faces, eyes, is therefore becoming more and more crucial. However, capturing simultaneously all these different information is still technically complex, and requires numerous dedicated hardware and software. Understanding the relative importance of such features on the plausibility of virtual characters is therefore extremely important, as is the ability of providing novel means of combining, synchronising, or generating automatically such features when they are not directly captured.

Finally, character interactions is another important aspect that should be accounted for when creating plausible animations (e.g., characters grasping objects naturally, avoiding other characters, following social norms) and are explored specifically in the next chapter.

Related Publications

[Hoyet et al. 2010b] Ludovic Hoyet, Franck Multon, Anatole Lécuyer, and Taku Komura (2010b). “Can We Distinguish Biological Motions of Virtual Humans? Perceptual Study with Captured Motions of Weight Lifting”. In: *Proceedings of the 17th ACM Symposium on Virtual Reality Software and Technology*. VRST ’10, pp. 87–90. DOI: [10.1145/1889863.1889878](https://doi.org/10.1145/1889863.1889878)

[Pražák et al. 2011] Martin Pražák, Ludovic Hoyet, and Carol O’Sullivan (2011). “Perceptual Evaluation of Footskate Cleanup”. In: *Proceedings of the 2011 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. SCA ’11, pp. 287–294. DOI: [10.1145/2019406.2019444](https://doi.org/10.1145/2019406.2019444)

[Hoyet et al. 2012b] Ludovic Hoyet, Kenneth Ryall, Rachel McDonnell, and Carol O’Sullivan (2012b). “Sleight of Hand: Perception of Finger Motion from Reduced Marker Sets”. In: *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*. I3D ’12, pp. 79–86. DOI: [10.1145/2159616.2159630](https://doi.org/10.1145/2159616.2159630)

[Ennis et al. 2013] Cathy Ennis, Ludovic Hoyet, Arjan Egges, and Rachel McDonnell (2013). “Emotion Capture: Emotionally Expressive Characters for Games”. In: *Proceedings of Motion on Games*. MIG

¹<https://metahuman.unrealengine.com/>

'13, pp. 53–60. DOI: [10.1145/2522628.2522633](https://doi.org/10.1145/2522628.2522633)

[Zibrek et al. 2013] Katja Zibrek, Ludovic Hoyet, Kerstin Ruhland, and Rachel McDonnell (2013). “Evaluating the Effect of Emotion on Gender Recognition in Virtual Humans”. In: *Proceedings of the ACM Symposium on Applied Perception*. SAP '13, pp. 45–49. DOI: [10.1145/2492494.2492510](https://doi.org/10.1145/2492494.2492510)

[Hoyet et al. 2013] Ludovic Hoyet, Kenneth Ryall, Katja Zibrek, Hwangpil Park, Jehée Lee, Jessica Hodgins, and Carol O’Sullivan (2013). “Evaluating the Distinctiveness and Attractiveness of Human Motions on Realistic Virtual Bodies”. In: *ACM Trans. Graph.* 32.6. DOI: [10.1145/2508363.2508367](https://doi.org/10.1145/2508363.2508367)

[Zibrek et al. 2015] Katja Zibrek, Ludovic Hoyet, Kerstin Ruhland, and Rachel McDonnell (2015). “Exploring the Effect of Motion Type and Emotions on the Perception of Gender in Virtual Humans”. In: *ACM Trans. Appl. Percept.* 12.3. DOI: [10.1145/2767130](https://doi.org/10.1145/2767130)

[Ennis et al. 2015] Cathy Ennis, Ludovic Hoyet, and Carol O’Sullivan (2015). “Eye-tracktive: Measuring Attention to Body Parts when Judging Human Motions”. In: *EG 2015 - Short Papers*. DOI: [10.2312/egsh.20151009](https://doi.org/10.2312/egsh.20151009)

[Kiiski et al. 2015] Hanni Kiiski, Ludovic Hoyet, Andy T. Woods, Carol O’Sullivan, and Fiona N. Newell (2015). “Strutting Hero, Sneaking Villain: Utilizing Body Motion Cues to Predict the Intentions of Others”. In: *ACM Trans. Appl. Percept.* 13.1. DOI: [10.1145/2791293](https://doi.org/10.1145/2791293)

[Niay et al. 2020] Benjamin Niay, Anne-Hélène Olivier, Katja Zibrek, Julien Pettré, and Ludovic Hoyet (2020). “Walk Ratio: Perception of an Invariant Parameter of Human Walk on Virtual Characters”. In: *ACM Symposium on Applied Perception 2020*. SAP '20. DOI: [10.1145/3385955.3407926](https://doi.org/10.1145/3385955.3407926)

3

Perception of Interacting Virtual Characters

CONTEXT AND OBJECTIVES

While creating visually plausible individual motions is a first important step to create expressive virtual characters, humans are social beings that are meant to interact with each other. Interactions are therefore important to create compelling simulations involving virtual characters. Such interactions can be of various nature, e.g., reacting to what another character says, avoiding another character walking past, tackling a character in a sports game, tripping on a object on the floor, etc. Thus, identifying and understanding the factors that influence the perception of such interactions is important to create believable simulations.

This is especially relevant for human interactions because of the complexity of human motion, as simulating interactions is a highly complex problem still often leading to visual anomalies. Previous research has therefore focused on exploring approaches combining motion clips to generate new large-scale multi-character environments [Lee et al. 2006; Shum et al. 2008], on computing descriptors of interactions between characters [Ho et al. 2010; Jin et al. 2018], or rely on machine learning to learn captured interactions [Starke et al. 2019; Starke et al. 2020]. A number of approaches also explored character interactions from a physics-based perspective [Liu et al. 2006; Bergamin et al. 2019]. These are however only a few examples over a large body of literature exploring these topics.

In parallel to these works, a number of approaches have also explored how such interactions are perceived, as eradicating any useless computational expense is always beneficial in interactive applications (e.g., video games) given the limited resources available. Perceptual studies are therefore beneficial to evaluate the importance of simplifications made to the simulation or animation techniques on the visual quality of the results, in order to find

optimal trade-offs between complexity and visual accuracy. For instance, the question of the visual fidelity of collisions between simple objects was explored in the context of collision detection [O’Sullivan et al. 2003; Reitsma and O’Sullivan 2009], where angular, momentum and spatio-temporal distortions were evaluated in terms of their effect on the perception of 3D rigid body collisions. Yeh et al. [2009] explored some of these results further to tune a physics-based rigid-body simulation system. Similarly, Kulpa et al. [2011] demonstrated that collision avoidance algorithms can be progressively discarded for background characters, therefore saving on expensive computational resources.

This chapter presents our research activities exploring how we perceive such interactions by studying the factors that affect their visual realism, both in terms of naturalness or perceived anomalies, with the goal of producing more natural interactions between virtual characters. Our contributions on this topic are twofold.

- We studied the factors that affect the **perception of causality in human interactions**, i.e, the impression that the action of one individual causes an appropriate reaction in another person or object. To understand how such causal events are perceived when they involve virtual humans, we first investigated the effects of timing, force, and angular anomalies on the perceived plausibility of two-character interactions [Hoyet et al. 2012a]. We then explored similar questions in the context of the perception of throwing animations where the throwing motion of the character and the trajectory of the thrown object were edited simultaneously or separately [Vicovaro et al. 2012; Vicovaro et al. 2014]. Overall, the results of these studies provide insights to aid in choosing the minimum number of motions to use and in understanding the effects of motion editing strategies commonly used in interactive applications.
- We also explored the **perception of interactions in the context of crowd animations**, where the quality of interactions between characters is paramount to creating visually compelling virtual crowds. First, we proposed a novel way for virtual characters to give the impression that they can achieve complex interactions with other characters, by displaying secondary shoulder motions when passing at close distances [Hoyet et al. 2016b]. More specifically, we explored the value of adding shoulder motions on the perceived visual quality of crowd animations, in terms of perceived collisions and animation naturalness. We also evaluated the perceived motion variety in the case of large-scale crowds (from 250 to 1000 characters), i.e., whether observers are able to identify virtual crowds in which motions are reused several times (motion clones). Surprisingly, we found that even for large-scale virtual crowds, it is very difficult for viewers to perceive that the same motions are used multiple times, even when only two motions (one male, one female) are used to animate the whole crowd [Adili et al. 2021].

The remaining of this chapter details two of these contributions. Focus 1 presents an investigation of the factors influencing the plausibility of two-character interactions [Hoyet et al. 2012a]. Then, Focus 2 presents two experiments exploring the value of adding shoulder motion interactions on the perceived visual quality of crowd animations [Hoyet et al. 2016b].

FOCUS 1: PERCEIVING CAUSALITY IN VIRTUAL INTERACTIONS



Figure 3.1: A typical virtual interaction depicting a causal event: one character pushes another, who reacts appropriately. Timing errors, force mismatches and angular distortions can affect the perceived realism of the scene [Hoyet et al. 2012a].

In this work, we were interested in the realism of physical interactions between virtual characters. While virtual humans have reached impressive levels of realism, with captured motions used extensively to provide very natural animations, disturbing artifacts such as interpenetrations and footsliding can occur as soon as interactions between virtual humans and the environment are needed. More specifically, the animation of characters who actually come into physical contact with each other and react appropriately (as illustrated in Figure 3.1), e.g., when touched, pushed or struck, is a difficult problem to solve, especially due to the large variety of possible physical reactions to such events.

In most situations, it is not possible to capture the motions of interacting characters simultaneously, thereby necessitating significant motion editing and animation by hand to correctly match up animations. It is also not feasible to capture reactions to all the possible physical interactions that might happen on a character, such as being pushed from all the directions, with a ranges of physical forces, as well as all the possible reactions that a character might display (e.g., falling, stepping and recovering, bracing itself). Alternatively, animation budget constraints can call for extensive reuse of a set of generic motions, resulting in a limited set of reactions that can often look implausible when linked together [Therrien and Bernard 2008; Laidacker and Barbeau 2011].

Research Questions

To understand how to create more visually plausible physical interactions, we were interested in exploring the following questions. *What is the minimum number of motion examples that are required to produce a range of natural physical interactions?* As it is typically not possible to capture the whole range of possible physical reactions, what is for instance the range of forces that viewers can discriminate and which should therefore be captured? *To what extent do physical distortions reduce the realism of character interactions?* It has been shown that the realism of dynamic events can be affected by different types of distortions, especially temporal, momentum and angular errors [Michotte 1963; O’Sullivan et al. 2003; Reitsma and O’Sullivan 2009]. For instance, for rigid body interactions, delays of 60ms were found to be perceptible 50% of the time [Reitsma and O’Sullivan 2009]. We can therefore wonder whether participants would be less sensitive to errors in character interactions than in the rigid body case, perhaps due to the increased complexity of human motion and perceived animacy. Also, as a selection of generic motions are often ‘mixed and matched’ (e.g., using the action from one actor and the reaction from another), *how acceptable are motion mismatches and what are the minimum differences that need to be captured when mismatching animations?*

Methods

To explore these questions, we studied the factors that affect the perception of causality in human interactions, i.e, the impression that the action of one individual (the *source*) causes an appropriate reaction in another person (the *target*). The perception of causal events has been studied by many researchers in psychology, and it has been shown by Michotte [1963] and in many subsequent studies that by manipulating timing and physical interactions during causal interactions, the meaning or realism of such events can be altered. Building on previous works on the perception of causality in interactions between simple objects [O’Sullivan et al. 2003; Reitsma and O’Sullivan 2009], our goal was therefore to understand how such causal events are perceived when they involve virtual humans.

We therefore investigated the effect of a number of anomalies that commonly appear when dealing with character interactions, using a motion capture set of two actors pushing each other at five different levels of force. After computing an estimation of the impulse force exerted by the source character on the target in each motion capture sequence, and validating that both actors were able to consciously push at five distinct force levels, we were interested in evaluating whether these different force levels were perceptible to a human viewer. We therefore first explored in a set of baseline experiments how the force of each interaction was perceived, and whether more dynamic information was available from the




Figure 3.2: Examples of physical interactions used in our experiments [Hoyet et al. 2012a]. a) Timing was altered to evaluate the effect of timing errors on the perception of causality during two-character interactions, and pushes and reactions from different sequences were mismatched to evaluate the effect of force mismatches. b) The 7 push directions (blue) and corresponding reactions (red) used to evaluate the effect of angular distortions, with examples of pushes from different directions.

motion of one of the two characters, where participants rated on a 5-point Likert scale the force level of interactions displaying either both characters, or only one of them.

To quantify the tolerance to physical errors in human interactions, we then conducted a number of “causality” experiments, investigating the effects of timing, force, and angular anomalies on the perceived plausibility of interactions. First, we tested participants on a range of **Timing Errors** from 0ms to 450ms, displaying either early or late reactions, for three different force levels (very light, medium, very heavy) which had been perceptually validated from the baseline experiments. Second, we tested participants on a range of **Force Mismatches**, combining source and target motions from different sequences, with varying levels of forces (very light, light, medium, heavy, very heavy). We hypothesized that the greater the force mismatch between the push and reaction force levels, the more inappropriate an interaction would appear, and that over-reactions (where the reaction is stronger than the actual push) would be less disturbing than under-reactions (as for rigid body interactions [O’Sullivan et al. 2003]). Third, we evaluated the perception of angular distortions for a push, depending on the direction in which it was applied, by testing source Push angles from 4 directions (0° , 22° , 45° and 67°) as well as pushes from both the left and right side of the target. For each direction, we generated 4 **Angular Distortions** of the target reaction with steps of 0° (i.e., matching source and target directions), 22° , 45° and 67° . Figure 3.2 illustrates some of the stimuli used in our experiments.

Main Results

First, we found that **participants could easily distinguish between different interaction forces**, even with only partial information (Figure 3.3, left). This was particularly true when both characters or only the target character were presented to participants. While animations displaying only the source character were consistently considered to look lighter,




Figure 3.3: Some results of the experiments in [Hoyet et al. 2012a]. Left: Average user rating of the force level of the interaction, when both or only one character was presented. Middle: Percentage of answers where the timing offset is considered to be correct, for different force levels. Right: Percentage of answers where the given mismatch between source and target force levels is considered appropriate.

differences were however small and within half a force level. Both source and target character motions therefore provide reliable information about interaction forces, showing that care must be taken when capturing and editing both character’s motions.

We also found that **participants were sensitive to all three types of anomalous interactions**. Timing errors of over 150ms were acceptable less than 50% of the time (Figure 3.3, middle), with early or late reactions being equally perceptible. Compared to perceptual timing thresholds in simple rigid body collisions where a delay of 60ms was acceptable 50% of the time [O’Sullivan et al. 2003], **users were more accepting of timing anomalies in human interactions**. This may be because of the added complexity of human motions compared to simple objects, and the fact that the characters are perceived as sentient beings who can act under their own volition. Tolerance of timing errors was however much lower for heavy and medium interactions than for light ones, possibly because of the level of dynamics involved.

Participants could also perceive force mismatches, though over-reactions were more acceptable than under-reactions. In our experiment, a difference of one step on our scale from 1 to 5 (very light to very heavy) was the maximum acceptable force level mismatch (Figure 3.3, right), showing that the synthesis of new interactions will only be plausible if the force level of both characters are closely matched. For interactive applications where a reaction is selected from a motion database, it may therefore be necessary to ensure that at least one reaction within this force step would be available for all the possible upcoming force magnitudes.

Angular distortions when a character reacts to a pushing force reduce the acceptability of the interactions. This should therefore be taken into account when designing a database of responsive motions, as angular distortions are likely to happen in interactive applications because of the mis-alignment of characters at the time of interaction.

Conclusion

In this work, we explored the perception of causality in character interactions, which help to better understand how viewers perceive interactions between virtual characters. Our results demonstrate that care must be taken when choosing the force level and the angular distortion of interactions, while simultaneously timing appropriately the reaction, or else the plausibility of an animation will be adversely affected.

Our results provide some first steps towards more natural interactions between virtual characters. However, more complex factors can also influence real interactions, such as an emotional response to a push, or multi-modal cues such as sound or touch. These contextual parameters could have a significant effect on the perception of the physical interaction, and should be further explored in the future.

FOCUS 2: PERCEPTUAL EFFECT OF SHOULDER MOTIONS ON CROWD ANIMATIONS



Figure 3.4: Crowd simulators typically do not account for character interactions at the animation level (red circle), while secondary motions such as shoulder motions (green and blue circles) can prevent spectators from perceiving slight residual collisions and globally increase the perceived level of naturalness [Hoyet et al. 2016b].

In this work, we explored the question of the realism of interactions resulting from character animations in crowds. More specifically, our goal was to improve the perceived visual quality of crowd animations by investigating the value of adding shoulder motions (i.e., secondary motions) to characters passing at close distances.

Typical crowd engine pipelines animate characters according to a two-step process [Thalmann 2007]. First, a crowd simulator generates the characters' global 2D displacement trajectories in the environment. Second, an animation system transforms these global trajectories into full-body motions. This two-step decomposition is interesting for computational reasons, as crowd simulators raise quadratic complexity issues by nature. For the sake of simplicity, simulation models are often limited to 2D moving circles with 3 DoF, i.e., two translations and a rotation. The complete set of internal trajectories (30 to 60 DoF per

character) is then considered at the animation step only, where characters are processed independently. This two-step process avoids combining the complexity of crowd simulators with the dimensionality of character kinematic models. However, this decomposition has drawbacks. The notion of interactions between characters is considered at the simulation level, and is lost at the animation level. Body animations are therefore not influenced by the presence of neighbours, only global trajectories are. Finally, discrepancies between the shapes of the 2D simulated agents and the animated geometrical character 3D models may create artefacts, such as residual collisions.

Research Questions

In real-world situations, shoulder rotations are often used by people to navigate efficiently in crowds, especially when the crowd density increases. As our objective is to create animated virtual crowds that look more natural, we were interested in exploring the following questions. *Can shoulder motions be used to mask collisions in pedestrian interactions?* Because of the two-step decomposition of crowd simulators, residual collisions between the 3D character models are often observed in virtual scenes. We believed that displaying shoulder motions might prevent viewers from perceiving some of these residual collisions, as the characters would be giving the impression that they achieve complex interactions. *Can shoulder motions be used to improve the naturalness of animations?* Simultaneously, we believed that displaying such secondary motions would make characters passing at close distances visually more natural, by adding interaction-dependent secondary motions, and therefore have a positive effect on the perceived quality of animations. *Is it possible to get such improvements only by changing the animation step to avoid costly effects on computational performances?* While there would be value in understanding the perceptual effects of using such secondary shoulder animations, the benefits would be limited if it required to invalidate the whole crowd animation pipeline. We however believed that these results could be obtained without significant complexification, by adding shoulder motions only at the animation stage based on simple triggering rules. This focus also demonstrates a typical experimental approach that we try to apply when possible to our research questions: understanding first the local effects (in that case focusing on isolated two-character interactions) before demonstrating the benefits on a larger-scale (in that case demonstrating that adding secondary motions in character interactions has also a significant impact on the visual quality of crowd animations with a very light impact on the computational cost of the whole animation pipeline).




Figure 3.5: We conducted 2 experiments to explore the influence of shoulder motions in character interactions. First, we focused on the local effects, by exploring the influence of the characters’ interpersonal distance and of the presence of shoulder motions on perceived collisions and naturalness in two-character interactions. Then, we evaluated the benefits of introducing shoulder motions in the situation of crowded scenes.

Methods

To explore these questions, we proposed to rely on the 2-step process illustrated in Figure 3.5. First, we focused on the local effect of these secondary motions and conducted a perceptual experiment considering two characters walking in opposite directions and passing by each other at close distances (Figure 3.5, top). We were interested in whether shoulder motions could be used to mask collisions in paired interactions, as well as how the crossing distance between two walkers would influence the perception of collisions. We therefore used 3 **Shoulder Motion** conditions: none, one or both characters displaying shoulder motions (resp. labelled Shoulder Motion: 0, 1 and 2); and 11 **Interpersonal Distances** ranging from 0.2m to 0.7m by 0.05m steps, computed as the distance between the root of the characters at the time of interaction. Finally, while interpersonal distance is a relevant and easy-to-control parameter for crowd simulation, shoulder motions naturally reduce the distance where collisions start to occur, because of the modification of the interpersonal volume between characters. We therefore also explored the perception of collisions regarding the actual **Interpenetration Volume**, which was computed a posteriori of the experiment as the average intersection volume of the characters’ mesh over the frames displaying interpenetrations. The effect of these independent variables was then evaluated in terms of **Perceived Collisions** and **Naturalness Ratings**.

In a second step, we evaluated the benefits of introducing shoulder motions in the situation of crowded scenes, where shoulder motions are diluted into much more visually complex animations (Figure 3.5, bottom). A large difference with the previously studied situation is that only a subset of the visible pairs of characters pass at close distances and actually require to display reactions. As a result, the addition of shoulder motions

makes a more subtle visual difference in the context of crowds than in the one of isolated characters. The question guiding the design of this experiment was then about evaluating what remains of the previously observed positive effects, which were particularly strong, now that shoulder motions are diluted into a more complex animation. We therefore generated crowd animations (using a gradient-based approach [Dutra et al. 2017]) and incorporated shoulder motions for characters brushing past each other or having residual collisions. We generated two different **Scenarios**: A) a single character walking against a unidirectional crowd flow, and B) two large groups of characters walking in a bidirectional flow. For each scenario, we also controlled the **Density** of characters on screen, (0.5, 0.7, 1.0, 1.6 and 2.8 *character/m²*), as well as the probability for characters to display shoulder avoidance reactions when an upcoming close interaction was detected (0%, 33%, 66% and 100%), which we called **Trigger Level**. The effect of these independent variables was then evaluated in terms of **Collision Ratings** and **Naturalness Ratings**.

Main Results

First, our results showed that **shoulder motions improve how viewers perceive interactions between two characters**. It decreases the overall probability of perceiving collisions, decreases the range of distances between characters where collisions are perceived (i.e., characters can pass each other at closer distances), as well as the residual volume of interpenetration tolerated in interactions where collisions are not detected. They also improve the naturalness of interactions, which was found to be negatively correlated with the proportion of perceived collisions. Some of these results are illustrated in Figures 3.6.

Our results also suggest that **shoulder animations could be used with smaller 2D circle radii during simulation**. Crowd simulation algorithms often approximate characters using 2D circles, whose radius directly influences the minimum interpersonal distance between characters, such as used in our experiments. Wide circles overestimate the characters' shape to prevent any residual collisions, but also prevent simulations from reaching high densities, while narrow ones result in numerous residual collisions between the final animated meshes. Instead, displaying shoulder motions enables characters to pass each other at closer distances while reducing the amount of residual collisions perceived by viewers and increasing the naturalness of interactions, and could therefore be used with smaller circles. Of course, these results hold up to a certain minimum interpersonal distance, where collisions become too obvious and cannot be masked anymore, even with shoulder motions.

While interpersonal distance is usually used to control interactions, collisions between characters are actually characterised by the volume of interpenetration between the characters' mesh. We found that **shoulder motions manage to hide larger residual collisions in two-character interactions**, as they increase the volume of interpenetration tolerated

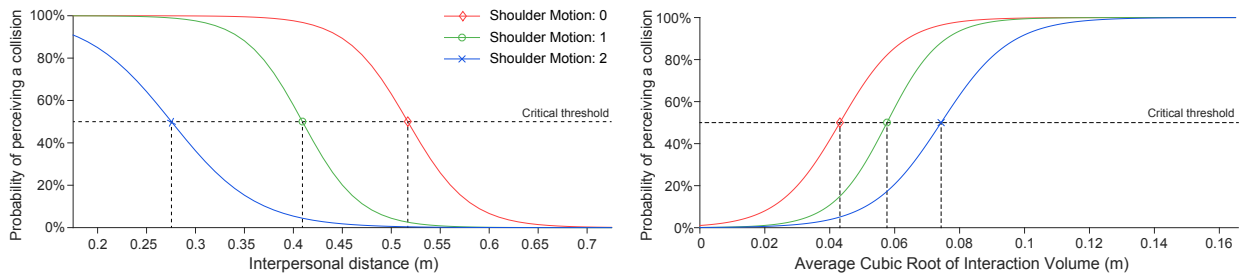


Figure 3.6: Psychometric curves depicting for each Shoulder Motion condition the probability of perceiving a collision depending on (Left) the distance between characters’ roots and (Right) the average cubic root of the interaction volume between characters.

by viewers. However, interpenetration volumes are still too computationally expensive to be used in interactive applications, even if more relevant from a perceptual point of view. Evaluating if selected parts of the characters’ body are more specifically used by viewers to detect interpenetrations (e.g., are users looking more at the lower or upper body to detect collisions?) could help simplifying such computations.

When studied in a larger context, our results also demonstrated that such **secondary shoulder motions improved the perceived naturalness of crowd animations**. This is a remarkable result since, in comparison with the first study involving only two characters, shoulder motions contribute much less to the overall animation: shoulder motions are triggered sparsely in space and time and concern few characters in proportion to all the visible ones in the crowd. Moreover, adding such secondary motions introduced low computational overheads, and did not fundamentally change the nature of the animation technique used (based on finite state machines). **Some factors however seem to limit the benefits of displaying shoulder motions in crowd animation**, since shoulder motions had no effect on naturalness in these cases. This was particularly the case for density, which directly influences the fact that more or less characters pass each other at a close distance, hence influencing the need for such secondary motions.

Interestingly, our results showed that **shoulder motions did not affect collision ratings in crowd animations**. Given the density of some of the simulations, and therefore the amount of resulting interpenetrations, it is possible that shoulder motions may not have been sufficient to counterbalance the amount of perceived collisions. However, this was also influenced by the displayed scenario, and we hypothesize that viewers could have more easily perceived the local interactions around the single character in the unidirectional flow scenario, while larger collisions could more easily be missed in the bidirectional flow scenario. Our results also suggest that viewers may evaluate indirectly the quantity of collisions that should be happening in global situations such as Scenario B, rather than looking directly at the local interactions between characters. This could explain why shoulder motions did not decrease perceived collisions, if information picked by viewers from the crowd density was stronger than the information picked from such secondary interaction-dependent motions.

Conclusion

In this work we studied the effect of adding shoulder motions on the animation of crowd characters. First, we demonstrated that shoulder motions significantly improve the visual quality of interactions between two characters. This improvement is revealed through our measurements of the perception of residual collisions and the evaluation of motion naturalness. Then, we considered animated crowds and demonstrated that the effect on the naturalness of animations remains significant, while the effect on the perception of residual collisions disappeared. Nevertheless, the trade-off between the cost of adding shoulder motions and the computation cost of such additions is extremely positive.

More generally, our work highlights the benefit of adding secondary motions to procedurally animated characters. Previous techniques allow to introduce motion variety when animating moving characters (e.g., by modulating template procedural animations), or to animate still characters during interactions (e.g., chatting characters). Our work opens a new path: we add variations to the motions of walking characters by adding interaction-dependent secondary motions at very low computational overheads. Such additions do not only improve how characters move, but also how they move with respect to their direct neighbours.

CHAPTER CONCLUSION

The work presented in this chapter explored the question of the perception of interactions between virtual characters, as creating natural interactions between characters is still an open challenge because of the complexity of human motion. Overall, the results of our experiments provide insights to guide the creation of such natural interactions, providing for instance guidelines to account for inter-relations between interaction forces and reaction timings, to appropriately combine animations from different sequences to create novel interactions, to plan motion capture sessions and ensure that sufficient motions are available to create interactions, as well as to propose novel ways to display more natural interactions in crowd animations.

While our contributions have so far focused on “physical” interactions, we believe that creating more expressive characters will require to handle more types of interactions in the future. For instance, in the context of crowd animations presented in our last focus, characters should be able to automatically display a larger range of human reactions, such as forms of social graces for characters politely giving way, expressions of surprise for characters close to colliding, eye-contacts for characters negotiating avoidance, attracting the attention of other characters at a distance, etc. Such interactions will require to go beyond purely “physical” interactions. In particular, in the context of moving towards immersive simulations as will be presented in the next chapter, we would also expect characters to react socially

more naturally to the actions of nearby users, such as reacting to inappropriate behaviours, or to unexpected events that would be made by users.

Creating more plausible interactions will also require to overcome the limitations of traditional animation approaches. As previously mentioned, animating large numbers of interacting characters usually relies on a 2-step process: global trajectories are generated by a crowd simulator (i.e., where to move), whereas full-body motions are generated by animation engines (i.e., how to move there). While this two-step decomposition is interesting for computational reasons, it also leads to potential discrepancies between the simulated trajectories and the body animations. This is a typical problem when animating crowd simulations even without considering secondary interactions, as available body animations in combination with standard animation techniques usually lead to either motion artifacts (e.g., footsliding, incorrect motion speed vs. trajectory speed) or to animations that drift from the simulated trajectories. It also means that the notion of interactions between characters is considered at the simulation level, and is lost at the animation level, therefore body animations are typically not influenced by the presence of neighbours. E.g., where we would like to observe characters side-stepping to avoid an upcoming character, we typically only see characters sliding away. On the basis of some of the results we obtained, we believe that a promising direction for future work is to store some useful information on ongoing interactions at the simulation stage, and to re-use it at the animation stage. Indeed, crowd simulators already compute various types of spatio-temporal information about ongoing character interactions (e.g., distances, relative positions or velocities, time to collision). This data is used to compute how each character's trajectory is influenced by neighbours but discarded at the animation stage, whereas it could also be used to trigger specific character animations. When possible, it would also seem more appropriate to shift paradigm and to move away from this 2-step decomposition, e.g., to explore novel ways to solve interactions directly in the motion space in a way similar to approaches such as Motion Matching [Buttner 2015].

Similarly, over the past two decades, a number of approaches have investigated the automatic synthesis of character interactions. Amongst these approaches is the regular trend of relying on physics-based methods [Zordan et al. 2005; Liu et al. 2006; Coros et al. 2010; Xie et al. 2020]. While these approaches rely on the laws of physics to synthesise motions, we also believe that hybrid approaches would be valuable to the community to improve the overall visual realism of such approaches. Indeed, while physics-based approaches are widely explored in our research community, they are still hardly used in industrial processes because of the gap in visual quality compared to data-driven or manual approaches. We believe that some of the methodologies that we have laid down in this chapter, e.g. to study causal events between virtual humans, could be useful to conduct further research into the realism of physics-based interactions. For instance, physics-based approaches are commonly

used in combination with reference motions, and one could wonder whether it would be visually more natural to add or remove forces from an existing example, e.g., would it be better to make light motions heavier than the opposite? Another class of approaches that is quickly rising in the community is to rely on Deep Learning, with the goal of designing models that will learn the important features driving interactions. Such approaches have been successfully applied to learn and synthesise interactions between a character and its environment [Starke et al. 2019], with other characters [Starke et al. 2020], or even to build physically responsive characters [Bergamin et al. 2019].

This last class of methods also raises the question of how to have access to interaction examples, especially in the context of Deep Learning which typically relies on large quantity of data. As capturing ground truth interactions can be extremely difficult, sometimes even impossible despite advances in Computer Vision approaches (e.g., imagine capturing the full-body animations of hundreds of people in a crowd), it also seems important to explore novel ways of obtaining ground truth interactions, which should also not be limited only to physical interactions but also cover the larger range of potential human interactions and reactions. As we will see in the next chapter, Virtual Reality seems to be a promising means of acquiring such data, as we are now able to immerse users in more and more realistic situations.

Related Publications

[Hoyet et al. 2010a] Ludovic Hoyet, Franck Multon, Taku Komura, and Anatole Lécuyer (2010a). “Perception Based Real-Time Dynamic Adaptation of Human Motions”. In: *Motion in Games*, pp. 266–277

[Hoyet et al. 2012a] Ludovic Hoyet, Rachel McDonnell, and Carol O’Sullivan (2012a). “Push It Real: Perceiving Causality in Virtual Interactions”. In: *ACM Trans. Graph.* 31.4. DOI: [10.1145/2185520.2185586](https://doi.org/10.1145/2185520.2185586)

[Vicovaro et al. 2012] Michele Vicovaro, Ludovic Hoyet, Luigi Burigana, and Carol O’Sullivan (2012). “Evaluating the Plausibility of Edited Throwing Animations”. In: *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. SCA ’12, pp. 175–182

[Vicovaro et al. 2014] Michele Vicovaro, Ludovic Hoyet, Luigi Burigana, and Carol O’Sullivan (2014). “Perceptual Evaluation of Motion Editing for Realistic Throwing Animations”. In: *ACM Trans. Appl. Percept.* 11.2. DOI: [10.1145/2617916](https://doi.org/10.1145/2617916)

[Hoyet et al. 2016b] Ludovic Hoyet, Anne-Helene Olivier, Richard Kulpa, and Julien Pettré (2016b). “Perceptual Effect of Shoulder Motions on Crowd Animations”. In: *ACM Trans. Graph.* 35.4. DOI: [10.1145/2897824.2925931](https://doi.org/10.1145/2897824.2925931)

[Adili et al. 2021] Robin Adili, Benjamin Niay, Katja Zibrek, Anne-Hélène Olivier, Julien Pettré, and Ludovic Hoyet (2021). “Perception of Motion Variations in Large-Scale Virtual Human Crowds”. In: *Motion, Interaction and Games*. MIG ’21. DOI: [10.1145/3487983.3488288](https://doi.org/10.1145/3487983.3488288)

4

Perception of Virtual Crowds in Immersive Environments

CONTEXT AND OBJECTIVES

In many ways, the works presented in the previous chapters all contribute to the creation of lively immersive environments. Approaches presented in Chapter 2 enable us to create characters that display individually plausible motions, while approaches presented in Chapter 3 enable us to create characters that interact more naturally with each other, with their environment, as well as with users immersed in the environment. Combining such approaches therefore contributes to creating more believable and interactive virtual crowds.

Nevertheless, designing believable virtual crowds is highly dependent on modeling individual and collective human behaviours, which in turn requires to study and understand such behaviours. For instance, it is important to understand the strategies that are used by pedestrians to avoid collisions with others, as well as to understand how they make their decision about which pedestrian to avoid, in order to faithfully reproduce such behaviours in virtual situations. Such behaviours have been widely explored in real conditions, e.g., investigating collision avoidance in pairwise [Olivier et al. 2012; Olivier et al. 2013] or in multiple pedestrian interactions [Dicks et al. 2016; Meerhoff et al. 2018b], exploring the effect of crowd density on walking speed [Seyfried et al. 2005], etc. However such studies were performed in real conditions, therefore involving dozens of human participants in a real environment, and are therefore difficult to replicate and standardize, thus limiting the factors that can be studied.

Instead, immersing participants in a virtual environment using VR has been considered to be a promising tool for exploring more complex situations in an experimentally controlled

environment, e.g., involving large number of persons, or more ecological situations. It was for instance used to compare pairwise interactions between real and virtual situations [Olivier et al. 2018], to explore when pedestrians choose to go through or around groups of characters [Bruneau et al. 2015], or to explore exit choices during evacuations [Ríos and Pelechano 2020]. It was also used for the study of proxemics in relation to virtual characters, i.e., the user response to the virtual character’s proximity, e.g., to evaluate physiological arousal [Llobera et al. 2010], or the impact of the appearance of virtual characters on the avoidance behaviour of participants [Mousas et al. 2021]. Overall the results of these studies demonstrate that behaviours in real and virtual environments display some similarities, e.g., in trajectory shapes, avoidance strategies, but that some quantitative differences nevertheless exist, e.g., participants walking more slowly in VR and keeping larger distances to obstacles [Fink et al. 2007], hence the need for quantifying the biases potentially introduced by the use of VR.

The general idea presented in this chapter is that by immersing users in virtual situations, where they can interact with (virtual) others, it is possible to study complex human behaviours that would be difficult to evaluate in real situations. Our goal is to use insights and knowledge gained from such immersive experiences to then improve the human animation models used in the previous chapters, in particular in terms of character interaction and simulation, which will in turn enable us to create more life-like immersive experiences to study more complex behaviours. This chapter presents our research activities on this topic, which I started to be involved in with the PhD of Florian Berton, whom we co-supervised with Julien Pettré and Anne-Hélène Olivier. Our contributions on this topic are the following.

- Using VR, we explored the **relation between gaze activity and collision avoidance** during pedestrian interactions. One of our goal is to better understand the interaction neighbourhood, i.e., which persons/characters within the crowd influences one’s motion, in order to improve how influencing agents are selected in crowd simulations. However, exploring such questions requires us to first investigate the biases on gaze activity that might be introduced by the use of VR. We first explored such biases in a pairwise collision avoidance situation [Berton et al. 2019], by comparing participants’ locomotion and gaze behaviour in a real environment with the same virtual situation performed using different VR setups (including a CAVE, a screen and a Head-Mounted Display). We then evaluated these biases when walking amongst people, by comparing eye-gaze activity while walking in a real street or its virtual replica [Berton et al. 2020b]. Understanding these biases enabled us to further explore the question of the interaction neighbourhood by investigating the impact of crowd density on gaze activity [Berton et al. 2020b].

- We also explored whether users’ behaviours while navigating within virtual crowds can be influenced by **rendering contacts with virtual characters using wearable haptics**. As we are relying on VR to study users’ behaviour in other investigations, our goal was to explore novel ways of improving user immersion when interacting with virtual characters. We therefore performed an experiment where participants navigated in a crowded virtual environment while experiencing or not haptic feedback of their collisions with virtual characters [Berton et al. 2020a], and showed that providing haptic feedback improves the overall realism of the interaction, as participants more actively avoid collisions.

The remaining of this chapter details two of these contributions. Focus 1 presents a study exploring eye-gaze activity within a crowd, to better understand the interaction neighborhood [Berton et al. 2020b]. In this study, we first evaluated potential biases introduced by the use of VR on the visual activity when walking among people, before evaluating in VR the influence of crowd density on gaze activity. Then, Focus 2 presents a study exploring whether rendering contacts with virtual characters using wearable haptics during navigation in a immersive crowded environment influences users’ behaviours [Berton et al. 2020a].

FOCUS 1: GAZE ACTIVITY IN VIRTUAL CROWDS



Figure 4.1: To analyze eye-gaze activity within a crowd, we designed two experiments where participants physically walked both in a real and virtual street populated with other walkers, while we measured their eye-gaze activity (red circle) [Berton et al. 2020b]. We evaluated the effect of virtual reality on eye-gaze activity by comparing real and virtual conditions (left and middle-right), then investigated the effect of crowd density in virtual reality (right).

In this work, we were interested in understanding how eye-gaze activity is influenced by the number of people we are interacting with in a crowd, as a mean of better understanding how people interact and avoid colliding with other pedestrians.

Modeling virtual crowds requires to design numerical models of local interactions which define by whom and how each individual (or agent) is influenced by the movement of others. The question of the who is also known as the interaction neighbourhood [Warren 2018]. Several models have been proposed for this neighbourhood, such as a fixed number of nearest agents [Wolinski et al. 2016], or any agents closer than a distance threshold [Helbing and

Molnár 1995], but these solutions were arbitrarily designed based on the study of trajectories. However, it was recently demonstrated that a combined analysis of gaze and trajectory data is meaningful for exploring these questions, e.g., that agents with a high risk of collision are more gazed at [Meerhoff et al. 2018a]. Thus, it is important to develop experimental protocols seeking to simultaneously study gaze and motion data. As such experiments remain difficult to perform in real conditions, VR therefore presents unique opportunities for such studies.

Research Questions

Following a previous study in which we investigated the influence of the choice of VR setups on gaze behaviour during collision avoidance between two walkers [Berton et al. 2019], we were interested in further exploring eye-gaze activity when interacting with a crowd. In particular, we were interested in exploring the following questions. As gaze seems to be indicative of how people take into account other individuals to navigate in crowds [Meerhoff et al. 2018a], *how is the eye-gaze activity influenced by the density of the crowd?* We were wondering whether the frequency with which people observe other individuals increases with density, enabling them to take a larger number of neighbours into account to adjust their trajectory, or, on the contrary, whether the number of neighbours taken into account is constant over time, while people pay more attention to those presenting the greatest risk of collision. We hypothesised that the duration of fixations would decrease as the amount of visual information increases with character density. We hypothesised that the amplitude of the saccades and the area covered by the gaze would also decrease as the result of participants' gaze being more focused towards the center of the visual field where pedestrians with the highest risk of collision typically are. *Is the eye-gaze activity of a human walking a street biased in VR conditions compared to real ones?* We showed in previous work [Berton et al. 2019] that eye-gaze activity showed qualitative similarities between virtual and real situations involving an interaction between two walkers (e.g., number and duration of fixations, ratio of what participants looked at in the scene, similar eye-head patterns), despite observing a few differences between the studied VR setups (HMD, Cave or Desktop Screen). However, it is unclear how such results would generalize to more complex scenarios involving larger numbers of pedestrians, as well as visually more complex scenes. We hypothesised that gaze spatial distribution and amplitude of eye saccades would be different between virtual and real situations, in part because of the reduced field of view of using a HMD, while the duration of gaze fixations would be similar as participants should take similar visual information in virtual and real situations.

Methods

To answer these questions, we conducted two VR experiments precisely controlling the visual neighbourhood of an immersed walker, while simultaneously measuring his/her movement and eye-gaze activity. In the first experiment, we estimated the **biases induced by the use of VR on eye-gaze activity**. We asked participants to walk in a real street, and recorded both their eye-gaze activity and their visual environment using Tobii pro eye-tracking glasses. We then reproduced the same situation in VR using a digital replica of the street, while reproducing as qualitatively as possible the virtual conditions in terms of virtual characters encountered (as seen by participants in the real situation estimated using video tracking). Participants performed 4 trials in each condition, as well as 2 initial training trials to get accommodated to VR. In a second experiment, we evaluated the **influence of crowd density on eye-gaze activity**. We therefore generated scenarios with 6 different conditions of densities $d \in [2, 5, 10, 14, 18, 24]$ (corresponding to d virtual characters every 15 meters at the start of the trial), and asked participants to navigate the virtual street populated with different densities of virtual characters. For each density, participants performed 4 repetitions. In our experiments, virtual characters were driven by RVO [Berg et al. 2008], an open-source crowd simulator often used in videogames, whose computational performances enable to have multiple agents avoiding collisions with other obstacles without impacting the framerate, which is crucial for VR experiments.

Both experiments were conducted in the context of walking in a busy street, which presented the advantage of corresponding to a daily-life situation, with no ambiguity on how to realize it: participants were asked to navigate in the street while avoiding collisions with pedestrians, and to perform multiple round trips between two specific locations (separated by 20m). Having a clear and simple task was important, as the nature of the task has a direct impact on the eye-gaze activity [Yarbus 2013]. Figure 4.1 provides examples of the two experiments. In both experiments, eye-gaze activity was analysed by comparing eye-gaze fixations (average duration and amplitude) and eye-gaze spatial distribution (peak of fixations and coverage).

Main Results

The first experiment enabled us to **identify differences and similarities in the walker's eye-gaze activity between the real and virtual situations**. In particular, participants' gaze in the virtual conditions was more intensely focused on the center of the visual field and less spread on the horizontal axis, while simultaneously showing on average a smaller amplitude of saccades and longer duration of fixations. There was however no difference in terms of visual coverage, while we expected the area covered by the gaze to be smaller in the

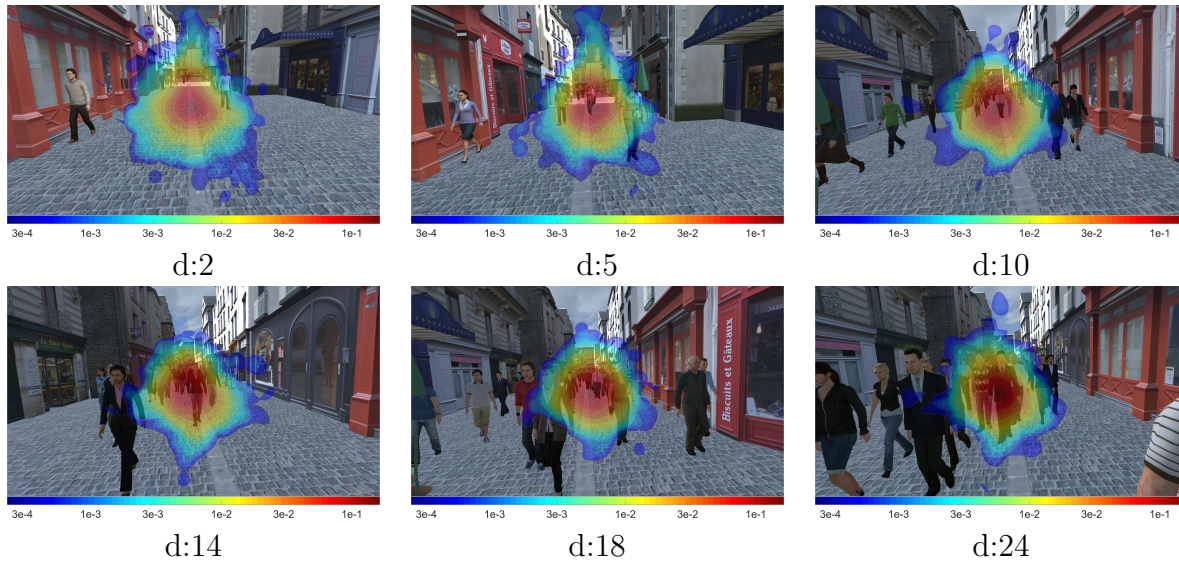


Figure 4.2: Virtual street with the different crowd densities, where the fixation area (log-transformed) is displayed on top of the image. The coverage was found to be significantly larger when navigating in a crowd with a density of 2 than for any other condition, while the average amplitude of saccades was found to be significantly higher for densities of 2 and 5 than any other.

virtual conditions because of a reduced field of view compared to real conditions. However, it is also possible that the performed task, i.e., walking straight in a street, may not have required the exploration of peripheral areas, especially as we found that participants showed a similar eye-gaze behaviour in relation to the goal of the task between real and virtual conditions, as they displayed a centered distribution for gaze location toward their goal. The average duration of fixations was also significantly longer in virtual conditions, which we hypothesized might have been caused by the level of detail in the digital replica of the street. In spite of the high quality of the digital scene, there are always missing details, e.g., birds, or sounds, that could have attracted the visual attention of our participants in the real conditions, and provoked faster fixations.

Then, results from the second experiment highlighted an **influence of density on gaze deployment**. In particular, we found that as density increases participants tend to look more at the center of their visual field, while the average amplitude of saccades decreases. These results suggest that **participants' gaze might be more focused towards pedestrians with the highest risk of collision**, which are typically the ones in front of them (i.e., in the center of their visual field). However, **participants took the same amount of time to search for visual information regardless of crowd density**. From the perspective of crowd modelling, our result that density has no significant effect on the duration of fixations lets us think that the number of interaction neighbours tends to remain constant with density. For a similar situation such as an opposite crowd in a street, we thus recommend to work with a constant number of neighbours in simulations. Concerning the selection features

to use to select neighbour agents, we recommend ordering them by risk of collision. Even though coverage did not significantly change with density, our results reveal that gaze tends to refocus around the visual center when density increases: this is the area where characters present the highest risk of collision in this bidirectional traffic condition.

Conclusion

In this work, we studied eye-gaze activity when interacting with a crowd of pedestrians. Overall, these results provide valuable insights to understand which aspects of eye-gaze activity are qualitatively similar between real and virtual situations, to further explore, using VR, aspects about navigating in a crowd which are difficult or not possible to evaluate in real situations. We thus believe that VR is a valid tool to study such activity, while experimenters should remain aware that some differences exist, e.g., that elements in real environments may distract walkers' gaze more. This work also opens new perspectives for future research, such as exploring whether these results generalize to other types of crowds (e.g., crossing flows, crowds in a more open space), the importance of peripheral areas of the field of vision depending on type of crowds, or whether some characters are observed more often than others.

FOCUS 2: CROWD NAVIGATION IN VR: EXPLORING HAPTIC RENDERING OF COLLISIONS



Figure 4.3: Our objective is to understand whether and to what extent providing haptic rendering of collisions during navigation through a virtual crowd (left) makes users behave more realistically [Berton et al. 2020a]. Whenever a collision occurs (center), armbands worn on the arms locally vibrate to render this contact (right). We carried out an experiment with 23 participants, testing both subjective and objective metrics regarding the user's path planning, body motion, kinetic energy, senses of presence and embodiment.

In this work, we were interested in exploring the simulation of physical contacts when navigating in a virtual crowd in VR, as illustrated in Figure 4.3, with the overall goal of improving user immersion.

As previously mentioned, the general idea presented in this chapter is that immersing users in virtual situations, where they can interact with virtual characters, enables us to study complex human behaviours that would be difficult to evaluate in real situations. However, when running such experiments there is also a high chance that participants will behave differently from how they behave in real life. The causes of such mismatches can be multiple, but are often linked to the limited sense of realism that VR can sometimes provide. In this respect, significant efforts have been put in the evaluation of the realism of behaviours in VR [Cirio et al. 2013; Olivier et al. 2018], in the understanding of the visual information required [Lynch et al. 2018] or in the development of highly-realistic VR environments and characters [Achenbach et al. 2017].

However, despite these efforts, most VR experiences still lack of any haptic sensation, which is of course of paramount importance when studying crowd behaviour and interactions. Haptic rendering was for instance demonstrated to be necessary for participants to display optimal adaptive behaviour (i.e., avoiding collision) when passing through a narrow door [Mestre et al. 2016], as well as to potentially improve presence and embodiment when a virtual human “bumps” into a person [Krogmeier et al. 2019]. More generally, haptic rendering can also be used as a mean of increasing communication between virtual humans and users [Boucaud et al. 2021]. In the case of virtual crowds, if we are unable to render the sensation of bumping into virtual characters when navigating in a crowded environment, participants might stop avoiding collisions, leading to data that does not capture well how humans truly behave. For this reason, studies of collective behaviour in VR are often limited to cases considering distant interactions only [Rio and Warren 2014; Rio et al. 2018], so as not to require any haptic feedback.

Research Questions

In this work, our objective was to investigate whether and to what extent the rendering of contacts with virtual characters influences users’ behaviour during navigation in dense crowds, as well as limits the occurrence of certain well-known artifacts, such as when the user’s virtual avatar interpenetrates other virtual characters. In particular, we were interested in exploring the following questions. *Will the use of haptic rendering influence participants’ body motions while navigating in the virtual crowd?* We hypothesized that it would make participants aware of collisions, and therefore make them navigate more carefully in the crowd, while displaying different body motions, e.g., turning shoulders to avoid collisions with virtual characters. We expected that this would lead to both fewer collisions and smaller volumes of interpenetration when a collision occurs. *Will there be some after-effect due to haptic rendering after it is disabled?* We expected that participants would remain more aware and careful about collisions even after we disabled haptic rendering, and would therefore continue

to display different body motions to avoid collisions with virtual characters. *Will haptic rendering improve immersion of participants in VR?* As we are interested in improving the overall immersion in such virtual experiences, we believed that haptic rendering would make users become more aware of their virtual body dimensions in space with respect to neighbour virtual characters, and therefore positively influence immersion. Immersion was explored through the sense of presence (i.e., the sense of being there) as well as through the sense of embodiment (i.e., the sense that the virtual body is as if it was one’s own biological body), which are often studied in the literature. In particular, exploring the sense of embodiment will be the common denominator of the work to be presented in Chapter 5 .

Methods

To investigate these questions, we conducted an experiment where participants navigated in a crowded virtual train station without, then with, and then again without haptic feedback of their collisions with virtual characters. For the purpose of the experiment, they were equipped with an Xsens motion capture system, a Pimax 5K+ HMD (chosen for its wide field of view, important in these situations of close proximity with other characters), four vibrotactile armbands (one on each arm and forearm) to provide sensations of contacts to the user’s arms, as well as a VR-ready backpack computer.

Participants were immersed in a digital reproduction of the Moscow metro station “Mayakovskaya”, amongst a virtual static crowd (see Figure 4.3, center). A total of 8 different configurations of the scene were prepared in advance and used in the experiment, with varying character positions, orientations, and animations. Characters were positioned according to a Poisson distribution, leading to an average density of 1.47 ± 0.06 character/m² and ensuring that a gap of 0.60 m on average existed between each character for participants to go through. The crowd was composed of standing virtual characters animated with varied idle animations to prevent the exact same animation clip to be used for several characters.

Participants were asked to walk through the crowd “as if they were walking in a real train station” so as to reach the board displaying train schedules on the other side of the virtual scene. They were then asked to read aloud the track number of the next train displayed on the board before coming back to their initial position. They were physically walking in the real room, while their position and movements were used to animate their avatar.

The experiment consisted of 3 blocks of 8 trials, in the following order: *NoHaptic1*, *Haptic*, and *NoHaptic2*. The *Haptic* block corresponded to performing the task with haptic rendering of contacts, while the *NoHaptic* blocks did not involve any haptic rendering of contacts. The purpose of the successive blocks was 1) to measure a baseline of participants’ reactions (*NoHaptic1*), 2) to investigate whether introducing haptic rendering influenced participants’ behaviour while navigating in a crowd (*Haptic*), and 3) to measure potential



Figure 4.4: Left: shoulder rotations (angle α_{SA}) are defined between the participants’ shoulder-to-shoulder axis and the segment connecting the two virtual characters. Right: the volume of interpenetration is computed for each frame iteratively using voxel spaces of decreasing dimensions. The final interpenetration volume is displayed in purple.

after-effects (*NoHaptic2*). At the end of each block, participants were asked to answer embodiment and presence questionnaires. Analysis of participants’ behaviours was then evaluated in terms of trajectory through the crowd (global motion), average walking speed, shoulder rotations (local body motion), number of collisions, maximum volume of interpenetration, presence and embodiment.

Main Results

First, our results demonstrated that **haptic rendering had an effect on local body motions**. Participants rotated more their shoulders when traversing the gaps between virtual characters during the *Haptic* block than during the *NoHaptic1* block. Our results suggest that participants might have tried to reduce the risk of collision with virtual characters when they experienced haptic rendering. This is consistent with previous work for participants passing through a virtual half-open door with or without haptic rendering [Mestre et al. 2016]. This is also similar to the results presented in Focus 2 of Chapter 3, where we demonstrated that shoulder motions enable to reduce both the actual and the perceived collisions between virtual characters. We also found that **participants walked at a slower speed on average when they experienced haptic rendering of collisions**, suggesting that participants moved more cautiously, which also resulted in fewer and less intense collisions. Haptic rendering did not however change the way participants selected their path through the crowd.

Interestingly, **some of the behaviours demonstrated by participants persisted even after disabling haptic rendering of collisions**: they continued to rotate more their shoulders and to display fewer collisions on average. However, this after-effect did not equally influence all measurements, as walking speed increased again. It is therefore possible that participants “calibrated” during the experiment, as they became more familiar with the environment, the task to be performed, and their virtual representation. However, to provide a more definitive conclusion on the role of the haptic after-effect would require to

add a control group with no haptic rendering throughout the whole experiment, which could be explored in the future.

Finally, **we did not find any significant change in terms of user’s perceived senses of embodiment and presence when experiencing haptic feedback**. This result was quite surprising, given that we found significant effects in other measurements, suggesting that participants took different actions when provided with haptic sensations of contact. An explanation for this result could lie in the fact that users already registered high embodiment and presence levels even when they did not experience haptic feedback, therefore quickly reaching a ceiling effect and leaving little room for improvement afterwards.

Another possibility is that vibrotactile feedback is not suited to improve immersion while rendering collisions in crowds, although there are several examples of this type of feedback being used to render similar events [Devigne et al. 2020]. Future works should also consider the question of the plausibility [Slater 2009] of such haptic feedback in a situation where participants can bump into virtual agents. For instance, unplausible haptic feedback may indeed bring participants close to the uncanny valley of haptics introduced by Berger et al. [2018], and thus would have an impact on the presence and embodiment of participants. Finally, a last explanation could be the location and number of our haptic devices. Employing a higher number of bracelets spread throughout the body might better render the target contact sensations.

Conclusion

In this work, we conducted an experiment to evaluate the effects and the after-effects of haptic rendering in a highly crowded environment. Results showed that providing haptic feedback impacted the way participants moved through the virtual crowd, even after it was disabled, but quite surprisingly did not impact the sense of presence or embodiment. Overall, these results demonstrate that a combination of visual and haptic feedback improves the overall realism of the experience, as participants show more realistic behaviours by being more cautious about not touching virtual characters. This also suggests that using haptic rendering would be beneficial while studying human behaviour that may lead to contacts in immersive situations.

CHAPTER CONCLUSION

This chapter presented some of our contributions related to the perception of virtual crowds in immersive situations. One of the purpose of our work is to improve current crowd and animation techniques by better understanding human behaviour within a crowd. Traditional approaches in modeling human behaviours or motions typically rely on observing, analysing

or abstracting real data to derive model parameters. As the complexity of the studied situations increases, it becomes more and more difficult to obtain such data. For instance, consider the logistics of organising crowd experiments involving simultaneously hundreds of participants in a single location. Now, add the additional constraints involved in capturing their body movements, as well as their facial expressions, eye-gaze activity, etc., to study how one goes through a dense crowd, where one looks at, or how one reacts to a sudden event.

VR therefore provides a powerful tool to overcome these limitations, as it enables to study multi-modal factors in controlled experimental situations by immersing a limited number of participants in a simulated virtual environment. Such experiments have for instance enabled us to study the *interaction neighbourhood* used in crowd simulators, i.e., which pedestrians influence ones' walking trajectory, based on people's eye-gaze activity while interacting with virtual crowds. We are also currently exploring how VR can be used to acquire high-quality eye-gaze activity, that will be used to validate a novel saliency-based model for animating the eye-head movements of autonomous virtual characters. Similarly, we are also interested in how motion characteristics can influence human behaviours in such contexts, e.g., how motion attractiveness influences proxemics in dyadic interactions [Zibrek et al. 2020], or how character behaviours such as making or averting eye contact [Raimbaud et al. 2022] can help attracting user's attention.

Our goal is therefore to explore these questions further in the near future. As there is a clear trend for digital content and applications to move towards immersive applications, our goal is also to progressively make approaches studied Chapters 2 & 3 converge with those presented in this chapter. This seems particularly important to explore more subtle and complex situations, as well as to better understand how we perceive character motions or interactions in immersive contexts, e.g., such as in terms of physical or social interactions.

Despite the potential of VR to study such questions, it is however important to keep in mind that VR also induces numerous biases. For instance, it has been demonstrated that participants frequently report shorter distances in virtual than in real environments [Renner et al. 2013]. In terms of interactions with the environment or virtual characters, they also often walk more slowly, (e.g., [Agethen et al. 2018]), keep higher distances to virtual objects (e.g., [Gérin-Lajoie et al. 2008; Argelaguet-Sanz et al. 2015]), and display differences in affordance judgments, e.g., whether they can step over a pole [Lin et al. 2015] or go through an aperture [Bhargava et al. 2020]. Similarly, eye-gaze activity was found to differ on some aspects, in particular in terms of eye-head coordination [Pfeil et al. 2018; Berton et al. 2019], which could be due to limitations in the field of view, screen resolution, or weight of commercially available HMDs. Such differences were also found to depend on the type of VR setup used [Berton et al. 2019], e.g., when using a CAVE or a HMD. Finally, it is also important to keep in mind while designing such experiments that the overall realism of the

virtual environment can have a strong impact on these biases. For instance, details such as sounds or small objects might attract the visual attention of people in real situations, and their absence in virtual environments might lead to behavioural differences. Similarly, limitations in terms of the virtual character models and motions available might affect the perceived variety and naturalness of the crowd, and therefore our behaviours towards these virtual characters. Overall, it is more than likely that factors that affect negatively the perception of virtual characters in non-immersive situations (such as presented in the previous chapters) are actually amplified in such life-like immersive situations, and should therefore be carefully taken into consideration. Hence the importance of this virtuous cycle where immersive experiments lead to improvements in character and crowd animation techniques, which in turn will lead to higher-quality immersive experiences.

Finally, VR also opens up new potentialities and enables us to propose novel unexplored paradigms, in particular where the user can be brought into the simulation or the creative process. For instance, we have recently started to explore with the PhD of Tairan Yin whether it is possible to collect a valuable crowd dataset using VR. The idea is that by immersing only one single user, successively embodying several virtual agents, he/she progressively builds the overall crowd behaviour by him/herself. We call this concept the One-Man-Crowd paradigm [Yin et al. 2022], and are exploring this line of research by comparing such synthetic crowd trajectories with real trajectories, in terms of emerging collective behaviours (e.g., propagation waves, line formations). While our results are still limited to unilateral interactions (i.e., pre-recorded characters cannot interact with the user), we are now exploring potential extensions of this work to include bilateral interactions, e.g., by considering iterative processes or by enabling previously recorded characters to potentially react to the new interaction, as well as to include a few users to act numerous characters. Nevertheless, this approach already shows an extremely valuable potential in proposing novel ways of acquiring complex crowd datasets while involving only a limited number of users, which would be valuable to several communities interested in crowd analysis or simulation. Similarly, we are exploring with the PhD of Vincenzo Abichequer Sangalli novel manners of augmenting virtual characters' motion abilities by letting users act examples of interactions directly with the virtual characters. Bringing the user in the simulation and creative process also raises questions in terms of user immersion, in particular when interacting physically and socially with other characters. In this context, how the users are represented in the virtual environment through their "avatar", and how this virtual representation affects how they interact with the environment is paramount to creating engaging experiences, which we explore in detail in the next chapter through the study of the sense of embodiment.

Related Publications

[Berton et al. 2019] Florian Berton, Anne-Hélène Olivier, Julien Bruneau, Ludovic Hoyet, and Julien Pettre (2019). “Studying Gaze Behaviour during Collision Avoidance with a Virtual Walker: Influence of the Virtual Reality Setup”. In: *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 717–725. DOI: [10.1109/VR.2019.8798204](https://doi.org/10.1109/VR.2019.8798204)

[Berton et al. 2020b] Florian Berton, Ludovic Hoyet, Anne-Hélène Olivier, Julien Bruneau, Olivier Le Meur, and Julien Pettre (2020b). “Eye-Gaze Activity in Crowds: Impact of Virtual Reality and Density”. In: *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 322–331. DOI: [10.1109/VR46266.2020.00052](https://doi.org/10.1109/VR46266.2020.00052)

[Berton et al. 2020a] Florian Berton, Fabien Grzeskowiak, Alexandre Bonneau, Alberto Jovane, Marco Aggravi, Ludovic Hoyet, Anne-Helene Olivier, Claudio Pacchierotti, and Julien Pettre (2020a). “Crowd Navigation in VR: exploring haptic rendering of collisions”. In: *IEEE Transactions on Visualization and Computer Graphics*. DOI: [10.1109/TVCG.2020.3041341](https://doi.org/10.1109/TVCG.2020.3041341)

[Zibrek et al. 2020] Katja Zibrek, Benjamin Niay, Anne-Hélène Olivier, Ludovic Hoyet, Julien Pettre, and Rachel McDonnell (2020). “The Effect of Gender and Attractiveness of Motion on Proximity in Virtual Reality”. In: *ACM Trans. Appl. Percept.* 17.4. DOI: [10.1145/3419985](https://doi.org/10.1145/3419985)

[Raimbaud et al. 2022] Pierre Raimbaud, Alberto Jovane, Katja Zibrek, Claudio Pacchierotti, Marc Christie, Ludovic Hoyet, Julien Pettré, and Anne-Hélène Olivier (2022). “The stare in-the-crowd effect in virtual reality”. In: *To Appear in IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*

[Yin et al. 2022] Tairan Yin, Ludovic Hoyet, Marc Christie, Marie-Paule Cani, and Julien Pettré (2022). “The One-Man-Crowd: Single User Generation of Crowd Motions Using Virtual Reality”. In: *To Appear in IEEE VR - Journal Track / IEEE Transactions on Visualization and Computer Graphics*

5

Avatars: Perception of Self Virtual Characters in Immersion

CONTEXT AND OBJECTIVES

The recent massive development of consumer-grade VR products has reinvigorated the need for self virtual characters, also called avatars. These avatars refer to the users' virtual representation and are typically used in HMD-based applications to compensate for the fact that users do not see their real body. Avatars have been shown to enhance perception of the virtual environment [McManus et al. 2011], as well as user performance [Rybarczyk et al. 2014] and experience [Yee and Bailenson 2007]. However, their usage typically raises a number of questions, such as: how to ensure that users are properly embodied in their avatar? How to animate avatars so that their movements correspond to the user's movements? What representation should be used to maximise the user's experience? Understanding what are the factors that affect users' perception of their avatar is therefore paramount to create more engaging experiences, which we explore in this chapter through the psychological construct of virtual embodiment.

Exploring users' sense of embodiment boils down to understanding the level to which an embodied avatar is considered to be one's own (virtual) body in the virtual experience. According to Kilteni et al. [2012], the sense of virtual embodiment is typically decomposed into three main dimensions: the sense of self-location (i.e., the feeling of being located inside the avatar's body), the sense of agency (i.e., the feeling of being in control of the avatar and having an influence on the virtual environment) and the sense of body ownership (i.e. the feeling of owning the avatar's body). The general idea is that design choices, such as in terms of avatar appearance or interaction capabilities, as well as individual characteristics, have

potential relative influences on the sub-components of the sense of embodiment. Following the pioneer Rubber Hand Illusion (RHI) experiment of Botvinick and Cohen's [1998], embodiment was quickly explored by the VR community, first by transposing the RHI in VR [Slater et al. 2008; Yuan and Steed 2010]. Numerous factors potentially affecting virtual embodiment have been explored, e.g., visuotactile and visuomotor synchronicity [Kokkinara and Slater 2014], avatar appearance coherency [Maselli and Slater 2013] and realism [Gorisse et al. 2019], visuo-tactile mismatches and self-contacts [Bovet et al. 2018]. It has also been demonstrated that it is possible to embody avatars highly different from one's own body, e.g., opposite-gender [Slater et al. 2010b] or child avatars [Banakou et al. 2013].

In this context, we have conducted a significant line of research on the topic of virtual embodiment, framed within the context of the Inria Research Challenge Avatar, where our overall goal is to push further the limits of perception and interaction through our avatars, to obtain avatars that are better embodied and more interactive. Most of the contributions resulted from the PhDs of Rebecca Fribourg and Diane Dewez, whom we co-supervised with Ferran Argelaguet and Anatole Lécuyer. These contributions can be divided into the following categories.

- We conducted a number of experiments **exploring the factors influencing the sense of embodiment**. For instance, some of our results showed that the visual realism of a virtual hand influences both the sense of agency and the sense of ownership, where more abstract representations elicited a stronger sense of agency and more realistic ones a stronger sense of ownership [Argelaguet et al. 2016]. We also evaluated the **relative importance** of three factors commonly found in the literature to influence the sense of embodiment (namely: the avatar's appearance, the avatar's degree of control, and the user's point of view), and found that point of view and control were consistently considered to be of higher importance than appearance [Fribourg et al. 2020a]. From the observation that embodiment scores can highly vary between users, we also explored whether some user characteristics (personality traits and body awareness) might explain such differences [Dewez et al. 2019]. Our results suggest that different dimensions of the Locus of Control influence the senses of ownership and agency.
- We **explored methodologies for measuring the sense of embodiment**, which is typically assessed in the literature through a combination of subjective (e.g., questionnaires, interviews) and objective (e.g., behavioural, physiological) measures. For instance, user responses to virtual threats are often used as an objective measure of the sense of embodiment. We demonstrated that the introduction of a threat does not alter users' sense of embodiment but can change their behaviour after the threat occurrence [Fribourg et al. 2020b; Fribourg et al. 2021a]. To understand inter-relations

between factors influencing the sense of embodiment, we also recently proposed to rely on a subjective matching technique, a novel methodology which had never been used in the context of the sense of embodiment [Fribourg et al. 2020a].

- We studied **inter-relations between interaction techniques and the sense of embodiment**, as some techniques might not be compatible with using full-body avatars. E.g., the avatar might hide important information, or the technique itself might negatively affect the sense of embodiment. First, we compared three state-of-the-art VR locomotion techniques (real walking, walking in place, or steering) [Dewez et al. 2020], and found that participants had a comparable sense of embodiment with all techniques. More recently, we also explored whether a dual avatar representation might be beneficial during anisomorphic manipulation [Dewez et al. 2022], to compensate for the mismatch between the performed and the displayed action which can negatively impact the sense of embodiment. While we found that users preferred having a single representation, embodiment scores were however similar between single and dual representations. We also explored the idea of designing “Avatar-Friendly” techniques, i.e., techniques which take the user’s avatar into account in the design process to preserve both user performance and sense of embodiment [Dewez et al. 2021].
- We proposed two **novel paradigms of interacting through the avatar**. We first evaluated the possibility of creating a Six-Finger Illusion in VR [Hoyet et al. 2016a], and demonstrated that a) participants responded positively to the possibility of controlling the virtual hand despite the structural difference, and b) accepted the virtual hand and individual digits as their own to some extent. We also recently proposed the novel concept of “virtual co-embodiment”, which enables a user to share their virtual avatar with another entity (e.g., another user) [Fribourg et al. 2021b]. We demonstrated that the users’ level of control over such a shared avatar, as well as their prior knowledge of their actions to perform, have an influence on the sense of agency (e.g., participants significantly overestimate their sense of agency when they can anticipate the motion of the avatar).

The remaining of this chapter details two of these contributions. Focus 1 presents our study on the relative preference between appearance, control and point of view, which relied on a subjective matching technique [Fribourg et al. 2020a]. Then, Focus 2 details an exploratory study aiming at identifying whether personality traits or body awareness might influence the sense of embodiment towards a virtual avatar [Dewez et al. 2019].

FOCUS 1: AVATAR AND SENSE OF EMBODIMENT: STUDYING THE RELATIVE PREFERENCE BETWEEN APPEARANCE, CONTROL AND POINT OF VIEW



Figure 5.1: The four tasks implemented in the subjective matching experiment [Fribourg et al. 2020a] exploring the relative preference between appearance, control and point of view and their effect on the sense of embodiment. From left to right: Punching, Soccer, Fitness and Walking tasks, with the avatar’s appearance at the maximum level of realism.

In the past years, many studies have tried to better understand how users perceive their avatar in VR by evaluating their sense of embodiment. More precisely, they focused on three subcomponents [Kilteni et al. 2012]: the sense of self-location, the sense of ownership and the sense of agency. From those researches emerged different “factors of influence” towards these three subcomponents, e.g., the avatar’s appearance [Argelaguet et al. 2016] or the user’s point of view [Galvan Debarba et al. 2017; Gorisse et al. 2017]. However, despite the worthwhile highlights brought by these studies, the inter-relations between the factors influencing the sense of embodiment remain uncertain. Indeed, if we start to better understand the influence of isolated factors on the sense of embodiment, we still have little information regarding the relative contribution of each factor towards the sense of embodiment, or regarding the user’s preference for a factor over another while being embodied in an avatar.

Research Questions

In this work, we aimed at better understanding the inter-relations between three factors commonly found in the literature to influence the sense of embodiment, namely the avatar’s visual appearance, the avatar’s control, and the user point of view. More specifically, we were interested in investigating the following questions. *Is there a dominant contribution between the factors of influence towards the sense of embodiment?* Common studies exploring the influence of factors towards the sense of embodiment tend to study these factors in isolation, i.e., they focus on one factor at a time and measure its influence on the sense of embodiment. Indeed, exploring potential inter-relations raises challenges in terms of experimental protocols




Figure 5.2: Levels of the Appearance factor. From left to right: Abstract, Stickman, Dummy, Opposite realistic, Neutral realistic and Personalized realistic avatars.

due to the numerous amount of possible factor combinations. We therefore aimed at better understanding these inter-relations by conducting a subjective matching experiment, which had already been used in the context of Presence [Slater et al. 2010a; Skarbez et al. 2017] and enables to evaluate a relatively high number of factor combinations. *Do users have preferences between those factors when it comes to enhance their sense of embodiment towards an avatar?* While it is important to understand these inter-relations, we also believe that it would be valuable for designers to better understand whether some of these factors should be prioritized in the creation of virtual avatars, in order to enable them to create virtual scenes enhancing the users' sense of embodiment. *What is the relative influence on the sense of embodiment of different levels of a specific factor?* The various factors influencing the sense of virtual embodiment typically have different levels, e.g., from abstract to realistic for appearance. However, little is actually known about the relative differences between these levels in terms of their influence on the sense of embodiment. For instance, does a stickman character elicit a significantly lower sense of embodiment compared to a dummy character? Or do animations driven by Inverse Kinematics (IK) elicit similar or lower levels of embodiment than animations driven by a motion capture system?

Methods

To provide insights to these questions, we conducted two experiments exploring user preference and perception of these three factors (i.e., visual appearance, control, and user point of view). For each factor, we identified several levels with an initial pre-supposed ranking. Based on past studies, we chose for the *Appearance* levels: an Abstract avatar, a Stickman, a Dummy, and three realistic avatars selected by each participant to be their opposite, neutral, or personalized avatar (Figure 5.2); for the *Control* levels: automatic pre-recorded animations, user-triggered pre-recorded animations, IK-based animations and motion capture; for the *Point of View* levels: third-person and first-person point of views. The main requirements for choosing the factors and levels were to ensure good coverage of potential implementations of an avatar according to each factor, as well as allowing the combination of levels between factors. For instance, we did not separate Appearance into texture and shape as realistic textures would

hardly be combinable with abstract geometrical representations. Similarly, we did not include finger animation since it could not consistently be combined with all the appearance levels.

The **baseline experiment** had the objective to create an ordered list for the levels within each factor (e.g., ranking between the different degrees of realism for an avatar appearance, ranging from abstract to personalised avatars). For each factor, participants experienced all levels while performing a task and had to rank the preference for each level on a scale from 0 to 100. The task consisted in recreating a yoga posture in front of a mirror.

The **subjective matching experiment** used a subset of levels for each factor, based on the results obtained in the baseline experiment, to explore how participants combined the different factor levels to reach a given level of sense of embodiment. The experiment consisted in having participants experiencing an “optimal” configuration of an avatar and then “recreate” the experienced sense of embodiment by iteratively increasing, one level at a time, one factor, starting from a “minimal” configuration. The final matched configuration, named accepted configuration, should match the same sense of embodiment experienced with the “optimal” configuration. The initial “optimal” configuration was supposed to elicit a high sense of embodiment as it considered a personalized realistic avatar, full-body motion capture and a first-person point of view, while the “minimal” configuration consisted in a minimal avatar, with automatic animations and a third-person point of view. In addition, to assess the potential impact of users’ actions while being embodied in an avatar, the subjective matching experiment considered four different tasks which covered four actions that can be done in a virtual environment: a) an interaction with the upper-body, b) an interaction with the lower-body, c) mimicking the actions of another virtual character full-body motions, or d) a constrained walking task (Figure 5.1).

Analysis was performed in terms of evaluating the scores given to each level in the baseline experiment. In the subjective matching experiment, we analysed both Accepted Configurations (i.e., the configuration in which participants reported a sense of embodiment similar to the one experienced with the “optimal” configuration), and the set transitions chosen by participants between configurations (i.e., the order of improvements made by participants to go from one configuration to another, analysed through the elaboration of a Markov chain for each of the four tasks).

Main Results

First, the baseline experiment enabled us to create a **monotonic ranking of the studied levels of each factor**, summarized hereafter. For *Appearance*, the Abstract avatar was the least preferred, then both Stickman and Dummy, then Neutral, and the Personalized avatar was the most preferred among all. Interestingly, the scores of the Opposite avatar did not permit us to separate it either from the Neutral avatar or from the Stickman &

Dummy. For *Control*, the results mostly highlighted that Motion Capture was preferred to all the other levels. For *Point of View*, the first-person PoV was preferred to the third-person PoV. Overall, these results are in line with previous works showing similar levels of body ownership towards avatars with different levels of anthropomorphism [Lugrin et al. 2015], higher scores for personalized avatars [Waltemate et al. 2018], as well as a preference for a first-person PoV [Gorisse et al. 2017], but clarified the relative ratings between levels. Levels which were not different enough within a factor were then discarded for the subjective matching experiment, to avoid introducing a bias towards selecting one factor in priority over another because of the lack of a difference between levels.

Then, the subjective matching experiment highlighted a number of interesting results. Maybe most important, our results showed that **Point of View and Control clearly appeared as the preferred factors when improving the configuration of the avatar**. This was primarily reflected in the first choice made by most of the participants in all tasks, as they chose to increase first either their level of Control or their Point of View at least 90% of the time (see Figure 5.3). This result was also visible both from the most likely paths chosen by participants, built from the set of transitions, where increases of the appearance level typically happened late, and from the analysis of the most accepted configurations, where realistic characters only appeared in the most accepted configurations for the Punching task. Our results thus suggest that for all tasks except Punching a low level of appearance (stickman) was enough to match the level of sense of embodiment felt in the optimal configuration. This is a rather intriguing result since avatar appearance is a factor widely studied in VR and known to have a strong impact on the sense of body ownership (e.g., [Argelaguet et al. 2016; Latoschik et al. 2017]). This also suggests that the type of task can have an important effect on the participants' expectations.

Moreover, we also found that **an increased preference for Control was visible in the Soccer task**, as it was increased twice more often than point of view in that task only. In all the other tasks, the most likely choice was to improve the appearance, after having improved control and point of view once (i.e., reach the [Stickman, IK and first-person PoV] configuration), while for Soccer the tendency was to increase Control at its maximum level (i.e., reach the [Abstract avatar, motion capture and first-person PoV] configuration). In our opinion, this was a rather intriguing result since the major improvement made from IK to motion capture is the gain in precision regarding the position and orientation of middle parts of the body (knees, elbows, etc.), which is not visible with the abstract avatar. While we do not have an explanation for why the control was that much improved in that task since the main change between the two levels of control should not have been visible, this result reinforces the importance of considering the potential influence of the task characteristics, especially as it might relate to performance. For instance, participants were instructed

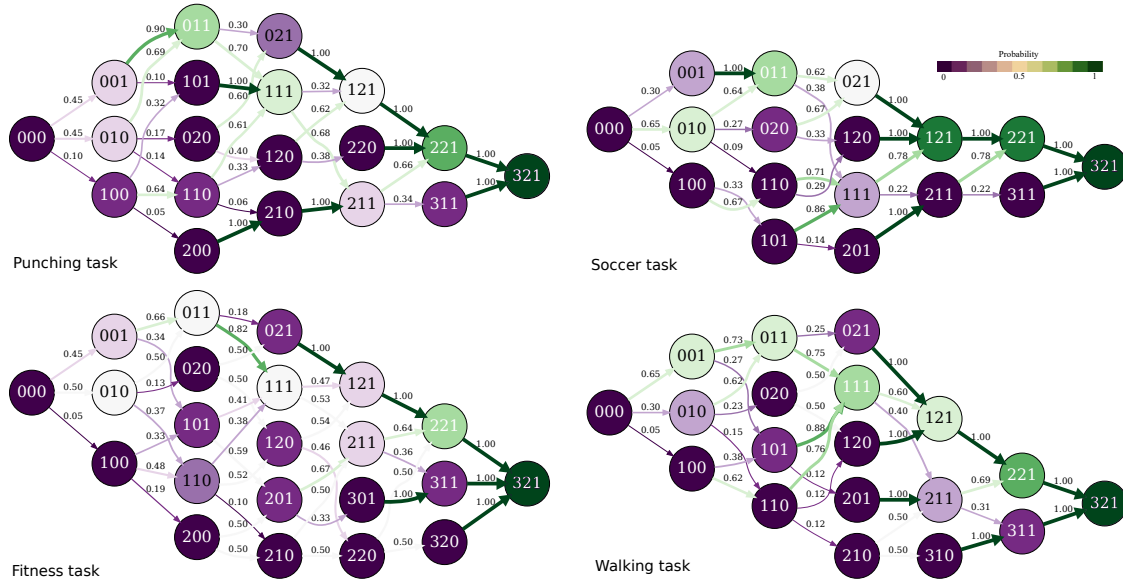


Figure 5.3: Analysis of participants’ choices of transitions in the subjective matching experiment [Fribourg et al. 2020a], using markov chains representing the transition matrix probability for each task. The configuration corresponding to each node is indicated by the ijk notation, where i is the level of Appearance (0: Abstract avatar; 1: Stickman; 2: Neutral realistic; 3: Personalized realistic) j is level of Control (0: Automatically launched animation; 1: Inverse Kinematics; 2: Motion Capture), and k is level of Point of View (0: Third- and 1: First-person Point of View). The color of a node represents the probability that the node is reached. The color and the thickness of the edges represent the transition probability from a given node.

that the objective of the task was only to kick the ball, and not to score a goal, but they might still have interpreted it a success or failure, and therefore associated the increased level of control with an increased chance of scoring a goal.

Conclusion

In this work, we conducted two experiments exploring user preference and perception of three factors commonly found in the literature to influence the sense of embodiment in VR. Our results show that appearance of the avatar was given less importance than control or point of view. We also found that when it comes to virtual embodiment users do not necessarily need to reach the optimal avatar configuration to feel a fulfilling sense of embodiment, suggesting that designers may not always need to provide high-end graphics avatars but should provide a high degree of control. Third, we showed that the accepted configurations can vary depending on the task performed, stressing the importance of this aspect for future studies and applications.

FOCUS 2: INFLUENCE OF PERSONALITY TRAITS AND BODY AWARENESS ON THE SENSE OF EMBODIMENT



Figure 5.4: We conducted an experiment exploring whether internal factors (personality traits and body awareness) might influence the sense of embodiment towards a virtual avatar [Dewez et al. 2019]. Participants first performed an exposure phase in which they executed a number of visuomotor tasks (left & center), after which a virtual character entered and stabbed the participants' virtual hand (right). After the experiment, we measured participants' sense of embodiment to evaluate the influence of internal factors on the acceptance of the virtual body.

In this work, we were interested in exploring another category of factors that can influence the sense of embodiment, namely internal factors (i.e., factors related to the user him or herself). More specifically, we report an exploratory study aiming at identifying whether personality traits or body awareness might cause either a resistance or a predisposition to feel a sense of embodiment towards a virtual avatar.

Most studies on the sense of embodiment explored what we might call external factors, i.e., factors related to the avatar or virtual environment themselves (such as in the previous Focus). In comparison, internal factors related to the users' personal state (e.g. personality or background experiences) have hardly been explored in the context of the sense of embodiment in VR. Such personality traits were previously explored in relation to the sense of presence showing, e.g., positive correlations between agreeableness and spatial presence [Sacau et al. 2005], and only more recently in relation to the sense of embodiment, showing a correlation between the sense of agency and the Locus of Control [Jeunet et al. 2018]. The majority of the works addressing such internal factors have otherwise mainly focused on users' sense of embodiment in the physical world.

Our observations while performing experiments over the years is that we typically observe a range of embodiment levels across participants, with some people easily believing in the virtual embodiment illusion, while others are on the contrary totally refractory. For this reason, we were interested in investigating whether some internal factors would have an influence on the sense of embodiment in VR, which would prove valuable for understanding how to maximize the sense of embodiment, or to adapt the VR experience to the user.

Research Questions

While numerous internal factors could be studied, we focused in this first exploratory study on the following selection, which were shown to either be related to the sense of presence in past studies, or to some sub-components of the sense of embodiment. Our goal was therefore to explore the following questions.

Are some of the users' personality traits correlated with their sense of embodiment in VR? The Big Five Personality traits taxonomy is a common way of describing one's personality using five dimensions (Openness to experience, Conscientiousness, Extraversion, Agreeableness and Neuroticism). For instance, agreeableness was found to be positively associated with spatial presence [Sacau et al. 2005]; openness, extraversion and neuroticism were also found to be positively correlated to immersive tendency [Weibel et al. 2010], which contributes to the sense of presence [Ling et al. 2013]. However, it is important to mention that a number of studies have found contradictory results. For the sense of embodiment, the majority of the works addressing the effect of personality traits have mainly focused on the RHI in the physical world, except for the recent work of Jeunet et al. [2018] which showed that the feeling of agency is linked to an internal Locus of Control (LoC). LoC is another set of personality traits, describing the degree to which people believe that they have control over the outcome of events in their lives (internal LoC), as opposed to external forces beyond our control (external LoC). LoC was also demonstrated to have an influence on the sense of presence [Murray et al. 2007; Wallach et al. 2010]. Our goal was therefore to explore whether some of these personality traits can have an influence on the sense of embodiment in VR. In the case of the LoC, we also wanted to confirm whether an internal Loc was indeed positively correlated with the sense of agency [Jeunet et al. 2018], and to explore whether it is also positively correlated with the sense of agency, since some studies found that the sense of ownership and the sense of agency are based on similar processes and can strengthen each other [Kalckert and Ehrsson 2012; Dummer et al. 2009].

Can shorter personality questionnaires be used as a simple way to assess users' sense of embodiment? Some of the questionnaires used to assess personality traits exist with various levels of complexity. For instance, Big Five questionnaires include the 10-item TIPI [Gosling et al. 2003], the 44-item BFI [John et al. 1991], and the 240-items NEO PI-R [Costa and McCrae 2008]. Therefore, we were also wondering whether there would be value in using shorter questionnaires to quickly assess personality traits prior to a VR experience, in the case where they would show similar correlations with the sense of embodiment than more complex questionnaires. Such a short evaluation could be a valuable tool in the future in attempts to maximise the users' sense of embodiment according to their personal characteristics.

Does the way we perceive our real body could have an influence on the perception of our virtual body? Body awareness is a cognitive ability that makes us aware of our body

processes, and which demonstrated contradictory results in the case of the RHI in the physical world [David et al. 2014; Tsakiris et al. 2006]. Furthermore, to our knowledge, no studies were conducted to investigate its influence on the sense of embodiment in VR.

Methods

To explore these questions, we conducted an experiment involving 123 participants, in which they were embodied in a gender-matched generic avatar. We chose to use generic avatars to obtain both low and high embodiment ratings across participants, as using personalised avatars might have led to high embodiment ratings for most participants [Waltemate et al. 2018] which would have prevented us from exploring the influence of individual traits on embodiment. Participants saw the virtual environment through an HTC Vive PRO HMD, while their hand movements were tracked using the Vive controllers. The FinalIK plugin was used to animate the participants' avatar with Inversed Kinematics and to provide visuomotor feedback, based on the participants' head (HMD) and hands (controllers) movements. The avatar was calibrated according to participants height and arm span.

The main experimental task was a visuomotor task involving the upper-body in order to elicit the sense of embodiment over the avatar. Participants sat in front of a real table and saw a similar co-located virtual table, while being immersed in the virtual environment from a first-person perspective (Figure 5.4, left). Participants were instructed to reproduce 2D trajectories that were displayed on a virtual screen in front of them (see example in Figure 5.4, left), with their left or right hand according to instructions. The trajectories were chosen to be relatively simple (number eight, circle, triangle, etc), to avoid a high cognitive load which could have distracted participants from their avatar and the environment. The task lasted two minutes during which participants saw their avatar moving synchronously according to their movements.

After the achievement of the task, a virtual character entered the room and stabbed the virtual hand with a knife. We recorded participants' hand reaction to the threat. This is a situation typically used in embodiment studies to provide an objective measure of the sense of embodiment, where participants who feel embodied in their avatar are more likely to remove their hand [Ehrsson et al. 2007].

Participants were then asked to fill in a number of questionnaires, always using the Likert scales proposed in the corresponding questionnaire. First, questionnaires on the sense of embodiment (19 items) [Gonzalez-Franco and Peck 2018] and on the sense of presence (6 items) [Usoh et al. 2000]). Participants then answered four additional questionnaires, chosen to explore aspects we believed could influence embodiment, while ensuring that the duration for answering all these questionnaires was reasonable. These included two well-known Big Five questionnaires: the 44-item Big Five Inventory (BFI) [John et al. 1991], providing

a good trade-off between length and completeness [John, Srivastava, et al. 1999], as well as the Ten Item Personality Inventory (TIPI) [Gosling et al. 2003], to evaluate whether a shorter questionnaire would lead to similar correlations with the sense of embodiment. Then the 24-item IPC scale [Levenson 1981], which determines LoC according to three dimensions: Internal, Powerful others and Chance. These dimensions typically mean that someone with an external LoC will tend to think that everything happens because of fate (chance type of locus) or powerful people (powerful others type of locus), while someone with an internal LoC will tend to think that he/she can change events with his own will and actions. Finally, the 18-item Body Awareness Questionnaire (BAQ) [Shields et al. 1989]. On average, participants took 20 minutes to answer all the questionnaires. Pearson correlations were then computed to explore potential links between the different components of embodiment and the other questionnaires.

Main Results

First, we observed from our experiment that **embodiment scores showed a high variability across our 123 participants**, ranging from low to high levels of embodiment: body ownership scores were average overall, agency and self-location scores were slightly higher, and both external appearance and threat perception scores were particularly low. On average participants felt in control of the avatar's movements, but did not really have the feeling that the avatar looked like them, which is what we expected as we chose to use a generic avatar. The latter was particularly true for female participants, which seemed to be related to the fact that the generic female avatar had less average physical characteristics regarding the population of the experiment than the male avatar. This potentially suggests that the avatar's visual resemblance also influenced these results, as well as highlights the importance of avatar personalisation as shown by previous studies [Waltemate et al. 2018; Fribourg et al. 2020a]. Overall, these results confirm some of our previous observations that some people easily believe in the virtual embodiment illusion, while others are less receptive.

In conducting this study, we expected some of the studied internal factors to influence the sense of embodiment. Our results suggest that **Big Five personality traits and body awareness do not explain for differences in the level of embodiment** of different participants. We only found a positive correlation between neuroticism and the reaction to threat (both in embodiment questionnaires and in actual behavioural responses). Even though body ownership scores tended to be low for people with high body awareness, the results were not significant.

Interestingly, our results demonstrated that **the Locus of Control is correlated to some extent with components of the sense of embodiment**: an internal LoC was found to be correlated with agency, while an external LoC seems to be linked to body

ownership. Therefore, it seems that people who feel a higher control on happening events tend to also experience a higher control of their virtual body, and might therefore feel more responsible of the avatar's movements. In contrast, body ownership seems to be influenced by external dimensions of the LoC. People with an external LoC tend to think that things happening to them depend mostly on the influence of other people or chance, and might feel more easily embodied in a virtual representation.

Finally, **our results showed similar correlations of personality traits with presence than in previous studies**, namely a correlation with agreeableness, as well as with an internal Locus of Control. Our goal was to validate that our experimental setup provided a similar basis than previous studies, to simultaneously strengthen the value of any results found in terms of influence of personality traits on the sense of embodiment.

Conclusion

This experiment on the sense of embodiment, in which 123 participants took part, is to our knowledge the first VR experiment measuring embodiment as well as several personality traits and body awareness. Given the amount of personality and cognitive models as well as questionnaires in the literature, we decided to focus on some of the most common models (i.e., Big Five, Locus of Control, Body Awareness). Our results confirmed that the locus of control provides some information about the sense of embodiment, but that Big Five personality traits and body awareness do not seem to be influencing factors. Given the interest in understanding how users' individual characteristics can influence the sense of embodiment, such as to maximise their sense of embodiment or to adapt their virtual experience, we believe that this work is a valuable first step in this direction.

CHAPTER CONCLUSION

This chapter presented some of our contributions related to the perception of avatars, i.e., the users' representation in immersive situations, which we explore through the psychological construct of virtual embodiment. Our global objective is to improve avatars in VR by a) enhancing our understanding of how users perceive their avatar and which factors influence their sense of embodiment, and b) exploring novel interaction techniques and representations that draw from these insights to increase users' sense of embodiment. Most of the research we performed on this topic was achieved thanks to the hard work of Rebecca Fribourg and Diane Dewez during their PhDs, which were framed within the Inria Research Challenge Avatar.

Overall, our results show that the sense of embodiment is a complex psychological process, which can be influenced by numerous factors, both internal and external to the

user. While our work has contributed to better understanding the influence of some of these factors, as well as their inter-relations, much is still left to explore on these topics. Over the studies we performed, maybe one of our most important result is the importance of the level of control that users expect to have in such experiences. This is something that we observed in several studies, e.g., in terms of preference for increasing first control over appearance [Fribourg et al. 2020a] or in terms of potential mismatch between the type of control available and the realism of the visual representation [Argelaguet et al. 2016], even though we also demonstrated that it is possible to make people overestimate their actual level of control in some circumstances [Fribourg et al. 2021b].

This thirst for Control is directly linked to the question of how we animate these avatars, which is highly related to Chapters 2 & 3. Our observations is that most current work still rely on simple animation techniques, e.g., Inverse Kinematics using the head and hand positions (3-point IK), sometimes including feet (5-point IK) and additional pelvis information (6-point IK). In a minority of cases, higher quality motion capture systems are used (e.g., Xsens, Vicon, Optitrack), while we have observed a clear preference for motion capture based animation when asking participants [Fribourg et al. 2020a]. As some participants reported, in several of our studies using IK, to be disturbed by animation artifacts (in particular related to leg movements [Dewez et al. 2019]) we believe that higher quality animation approaches specially tailored for VR purposes will be necessary in the future. The challenge in designing such approaches is to be able to display natural motions based on information from few consumer-level interaction devices (typically hand-held controllers, HMD, and trackers). While some recent approaches are going in this direction, either by proposing upper-body VR-tailored IK approaches based on heuristics [Parger et al. 2018] or by relying on deep learning models to predict lower-body poses from head, hands, and pelvis positions [Yang et al. 2021], we are still a long way from being able to generate high-quality motions for avatars in VR, with approaches designed with virtual embodiment in mind.

Considering virtual embodiment is however something that we have already started to explore in the context of interactions in VR, by exploring how to design interaction techniques compatible with avatars. We call such techniques “Avatar-Friendly” techniques [Dewez et al. 2021], and define them as *techniques which take the user’s avatar into account in the design process to preserve both user performance and the sense of embodiment*. As most interaction techniques were not designed with avatars in mind, our goal is therefore to go beyond the standard process typically based on user performance to also account for the actual perception of the avatar as well. For instance, different anisomorphic manipulation techniques have been proposed to enable users to interact remotely or more finely in VR, and which therefore distort users’ motions. As such techniques were not originally designed to be used with avatars, they can negatively impact the sense of embodiment as they create a mismatch

between the performed and displayed actions. For this reason, we explored whether a dual avatar representation might make such interaction techniques more “avatar-friendly” [Dewez et al. 2022]. While we found in that specific case that users preferred having a single representation, despite similar embodiment scores between single and dual representations, we believe that there is value in this process of designing interaction techniques aiming at preserving both user performance and the sense of embodiment. This work started during the PhD of Diane Dewez, and is now being further explored in the PhD of Maé Mavromatis, in collaboration with Ferran Argelaguet and Anatole Lécuyer.

Another major challenge in evaluating users’ perceptions of their avatar lies in the actual measurement of their level of embodiment. Despite their clear limitations, the research community still commonly relies on subjective questionnaires. While several attempts have been made recently to have more “standard” questionnaires [Gonzalez-Franco and Peck 2018; Roth and Latoschik 2020; Peck and Gonzalez-Franco 2021], because of differences specific to each experiment the individual elements of these questionnaires can vary significantly from one study to another. Also, these questionnaires typically provide discrete measurements of the sense of embodiment: in most cases a single occurrence at the end of the experiment or of each specific condition. These limitations demonstrate the need for novel, if possible less subjective, ways of measuring the sense of embodiment. In this sense, exploring novel methodologies to evaluate how users perceive their avatars, as we explored with the subjective matching technique to study inter-relations between factors [Fribourg et al. 2020a], seems to be a promising direction. Similarly, other types of objective measures other than the response to a virtual threat would be valuable, as they could also provide the ability to measure a more “continuous” level of embodiment, potentially exploring ways of measuring when the sense of embodiment is impaired in a way similar to what has been done with Breaks in Presence [Slater and Steed 2000]. Finally, as previously mentioned in the Focus 2, we still do not understand well how individual differences influence our abilities to feel embodied in avatars in VR. While several studies seem to agree that the Locus of Control does provide some information, it would also be interesting to push this direction of research further, to potentially be able to know in advance how embodied a given user will be in a specific experience. This could enable designers to tailor the experience to maximize each user’s sense of embodiment (e.g., through automatic preselection of a visual appearance or style of the avatar).

To conclude this chapter, we would also like to point that while most of our work explores virtual situations replicating real-life ones, the ability of embodying a virtual avatar also opens up novel possibilities far beyond what we can experience in real-life. We explore some of these aspects, such as to understand how users perceive an avatar with structural differences by embodying a six-digit realistic-looking virtual hand [Hoyet et al. 2016a], how users feel in control of a shared virtual avatar they are embodied in [Fribourg et al. 2021b], or whether

we can increase the sense of embodiment using dual body representations in anisomorphic manipulations [Dewez et al. 2022]. We believe that exploring these questions is also extremely important, as results from such original experiments have the potential to unfold a new range of applications in the fields of VR-based training, education, or even reeducation.

Related Publications

[Hoyet et al. 2016a] Ludovic Hoyet, Ferran Argelaguet, Corentin Nicole, and Anatole Lécuyer (2016a). ““Wow! I Have Six Fingers!”: Would You Accept Structural Changes of Your Hand in VR?”. In: *Frontiers in Robotics and AI* 3, p. 27. DOI: [10.3389/frobt.2016.00027](https://doi.org/10.3389/frobt.2016.00027)

[Argelaguet et al. 2016] Ferran Argelaguet, Ludovic Hoyet, Michael Trico, and Anatole Lécuyer (2016). “The role of interaction in virtual embodiment: Effects of the virtual hand representation”. In: *2016 IEEE Virtual Reality (VR)*, pp. 3–10. DOI: [10.1109/VR.2016.7504682](https://doi.org/10.1109/VR.2016.7504682)

[Fribourg et al. 2018] Rebecca Fribourg, Ferran Argelaguet, Ludovic Hoyet, and Anatole Lécuyer (2018). “Studying the Sense of Embodiment in VR Shared Experiences”. In: *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 273–280. DOI: [10.1109/VR.2018.8448293](https://doi.org/10.1109/VR.2018.8448293)

[Dewez et al. 2019] Diane Dewez, Rebecca Fribourg, Ferran Argelaguet, Ludovic Hoyet, Daniel Mestre, Mel Slater, and Anatole Lécuyer (2019). “Influence of Personality Traits and Body Awareness on the Sense of Embodiment in Virtual Reality”. In: *2019 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 123–134. DOI: [10.1109/ISMAR.2019.00-12](https://doi.org/10.1109/ISMAR.2019.00-12)

[Dewez et al. 2020] Diane Dewez, Ludovic Hoyet, Anatole Lécuyer, and Ferran Argelaguet (2020). “Studying the Inter-Relation Between Locomotion Techniques and Embodiment in Virtual Reality”. In: *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 452–461. DOI: [10.1109/ISMAR50242.2020.00070](https://doi.org/10.1109/ISMAR50242.2020.00070)

[Fribourg et al. 2020a] Rebecca Fribourg, Ferran Argelaguet, Anatole Lécuyer, and Ludovic Hoyet (2020a). “Avatar and Sense of Embodiment: Studying the Relative Preference Between Appearance, Control and Point of View”. In: *IEEE Transactions on Visualization and Computer Graphics* 26.5, pp. 2062–2072. DOI: [10.1109/TVCG.2020.2973077](https://doi.org/10.1109/TVCG.2020.2973077)

[Fribourg et al. 2020b] Rebecca Fribourg, Evan Blanpied, Ludovic Hoyet, Anatole Lécuyer, and Ferran Argelaguet (2020b). “Influence of Threat Occurrence and Repeatability on the Sense of Embodiment and Threat Response in VR”. in: *ICAT-EGVE 2020 - International Conference on Artificial Reality and Telexistence and Eurographics Symposium on Virtual Environments*. DOI: [10.2312/egve.20201262](https://doi.org/10.2312/egve.20201262)

[Dewez et al. 2021] Diane Dewez, Ludovic Hoyet, Anatole Lécuyer, and Ferran Argelaguet Sanz (2021). “Towards “Avatar-Friendly” 3D Manipulation Techniques: Bridging the Gap Between Sense of Embodiment and Interaction in Virtual Reality”. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. DOI: [10.1145/3411764.3445379](https://doi.org/10.1145/3411764.3445379)

[Fribourg et al. 2021a] Rebecca Fribourg, Evan Blanpied, Ludovic Hoyet, Anatole Lécuyer, and Ferran Argelaguet (2021a). “Does virtual threat harm VR experience?: Impact of threat occurrence and repeatability on virtual embodiment and threat response”. In: *Computers & Graphics* 100, pp. 125–136. DOI: [10.1016/j.cag.2021.07.017](https://doi.org/10.1016/j.cag.2021.07.017)

[Fribourg et al. 2021b] Rebecca Fribourg, Nami Ogawa, Ludovic Hoyet, Ferran Argelaguet, Takuji

Chapter 5. Avatars: Perception of Self Virtual Characters in Immersion

Narumi, Michitaka Hirose, and Anatole Lécuyer (2021b). “Virtual Co-Embodiment: Evaluation of the Sense of Agency While Sharing the Control of a Virtual Body Among Two Individuals”. In: *IEEE Transactions on Visualization and Computer Graphics* 27.10, pp. 4023–4038. DOI: [10.1109/TVCG.2020.2999197](https://doi.org/10.1109/TVCG.2020.2999197)

[Dewez et al. 2022] Diane Dewez, Ludovic Hoyet, Anatole Lécuyer, and Ferran Argelaguet Sanz (2022). “Do You Need Another Hand? Investigating Dual Body Representations During Anisomorphic 3D Manipulation”. In: *To Appear in IEEE VR - Journal Track / IEEE Transactions on Visualization and Computer Graphics*

“Once a scientist attacks a problem which he knows to have an answer, his entire attitude is changed. He is already some fifty per cent of his way toward that answer.”

— Norbert Wiener, *The Human Use Of Human Beings: Cybernetics And Society*

6

Perspectives

This manuscript presented an overview of the work we performed over approximately the last ten years on the topic of creating visually plausible virtual humans. Our research was largely driven from Computer Graphics and Applied Perception paradigms, with applications for the movies and video-games industries. Hence the importance in our research questions of being able to balance visual quality with computational and economical constraints, i.e., being able to identify how much can be computed interactively, to estimate the visual plausibility of the resulting animations, and to provide insights on how to create plausible motions from the viewer’s point of view. Over this period, we believe that our most worthy achievements are the following.

- We performed numerous perceptual studies exploring a variety of questions related to the perception of virtual character motions, most of them directly related to the topics presented in the different chapters of this manuscript. In total, approximately 1700 participants took part in more than 35 different experiments. Even though these numbers do not represent unique participants, they demonstrate some of the energy and time required to explore the topics presented here.
- Amongst these studies, many of them were performed with the target of creating life-like virtual characters, a target which evolved over the years towards creating life-like animated crowds in which users can immerse themselves using VR. This objective is extremely challenging as it requires expertise in many areas, such as computer animation, crowd simulation, virtual reality, biomechanics, evaluation methods (often inspired from the fields of experimental psychology and psychophysics), etc. The complementary expertise of all the persons involved in our research projects enable

us to contribute to all the different areas involved in creating such immersive life-like virtual crowds. For instance, we explore novel methods for creating motion variety both at the individual and collective levels, where one of our main objective is to be able to create characters with perceivable individual characteristics while simultaneously maintaining a visual coherence at the collective level. We also contribute to novel crowd simulation approaches, as well as to exploring novel ways of understanding human behaviour by immersing people using VR. This puts us in an original position of being able to study increasingly more complex and ecological situations, while simultaneously exploiting the results of such immersive studies to further improve animation techniques.

- In a relatively short time, we have managed to position ourselves at the international level on the topic of avatars. By relying on our combined expertise on computer animation techniques and virtual reality, we believe that we have significantly contributed to improving our understanding of some of the factors that influence the sense of virtual embodiment, and are now able to explore in a unified framework research questions related to both avatar interaction and avatar embodiment.

The remaining of this chapter will now detail some of the questions we are planning to explore in the near future. These research directions (summarized in Figure 6.1) are driven by a number of ongoing research projects and PhDs, and will mostly continue to be organised around the four objectives presented in Chapter 1. In short, our objective is to explore a) novel animation techniques for creating perceptual motion variety for individuals and groups, b) novel ways of designing character interactions from user demonstrations in VR, and c) novel animation and interaction techniques specifically designed for avatars to account for both performance and sense of embodiment.

MOTION VARIETY

We believe that we are still lacking in **efficient animation techniques** that would enable us to easily create motions that are adapted to the **individual characteristics** of virtual characters, and that are necessary for helping in creating **variety in animated crowds**. As mentioned in Chapter 2, this means that in most cases animating virtual characters still relies on a small number of generic motions, that are used with little consideration of the actual characteristics of the animated character. For these reasons, we want to further explore two main directions concerning motion variety in the near future.

First, we believe that identifying the parameters that affect how viewers perceive individual characteristics is particularly relevant to **create motions with characteristics that are**

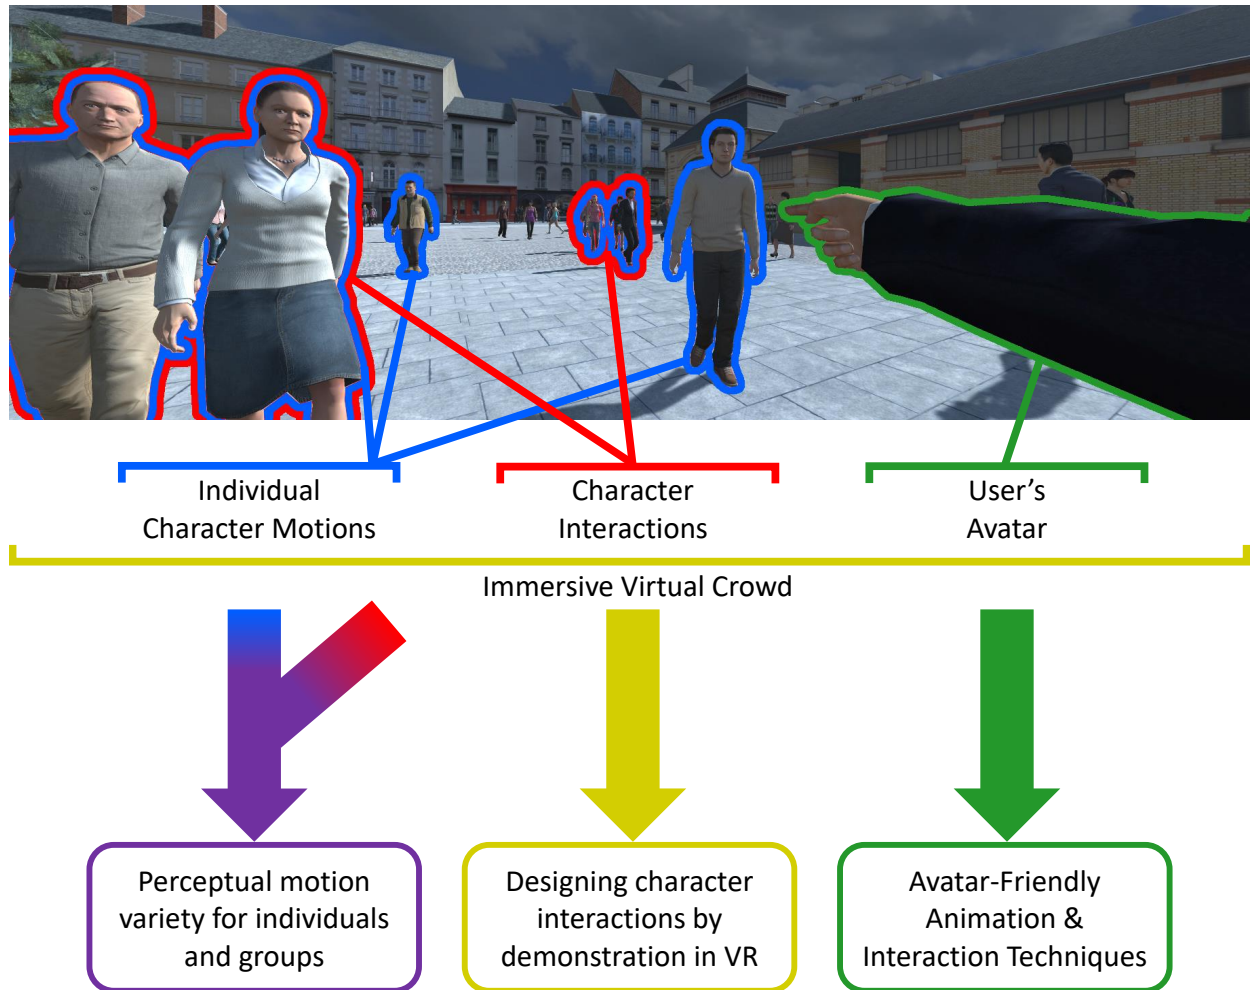


Figure 6.1: The three principal research directions that we are planning to explore in the near future, positioned in the general framework of the research presented in this manuscript.

perceived and identified as such by viewers. If we consider the ‘simpler’ case of locomotion, we are interested in making characters distinguishable by the way they walk (PhD of Benjamin Niay, ANR JCJC Per²) depending on their individual characteristics (e.g., walking style, morphology, physiology, age), just like real humans are. For instance, we are interested in exploring how to create motions that are perceived as corresponding to the motion of someone tall, small, heavy, young, old, etc. This is particularly challenging as it requires to acquire a considerable dataset of human motions representative of a large range of individual characteristics, before identifying which motion features are perceptually relevant for variety using statistical and data mining approaches. Then, the challenge becomes to be able to synthesise natural motion variations based on these perceptual features, while maintaining both the visual realism of the resulting motion, as well as identifiable individual characteristics. We are also currently exploring the use of machine learning approaches to synthesize character motions usable in industry-standard productions (PhD

Lucas Mourot, CIFRE InterDigital), e.g., by creating more detailed or varied animations, by synthesizing additional joints not observed during data acquisition, or by “upsampling” existing animations. Our idea is that it should be possible to learn a detailed, compact, and expressive representation of human motion, so that such additional details can be inferred from lower-quality animations with the goal of efficiently introducing motion variety.

Second, we are also interested in the question of the **perception of variety in animated crowds**, as we found somehow contradictory results in previous studies: we found that the introduction of a few shoulder animations when characters walked pass each other at close distances improved the overall naturalness of crowd animations [Hoyet et al. 2016b], possibly because of the task-dependent variety we introduced, but also that viewers could hardly differentiate crowd animations with only two unique motions from those involving more unique animations [Adili et al. 2021]. As we also found that there are clear differences in how visually distinctive some individual motions typically are [Hoyet et al. 2013], we are therefore interested in exploring further the question of whether the perception of motion variety is more related to task-dependent than individual-dependent variety. In other words, is it more important to display characters performing a variety of actions, e.g., avoiding in different ways other characters, talking on the phone, hurrying through the crowd, than displaying a variety of motions adapted to the characteristics of the users, e.g., motions of both old and young characters. We are also interested in exploring this question in the context of **immersive virtual crowds** (Robin Adili, ANR JCJC Per²). To our knowledge, all the previous studies on this topic have focused on crowd animations presented from a top-down or canonical viewpoint, as viewed from an bird’s-eye observer. However, being immersed in a virtual crowd means that viewers will be interacting with virtual characters from a viewpoint similar to real-life situations. We believe that the question of motion variety might become more crucial in such cases, as idiosyncrasies might become more perceivable as we repeatedly interact with characters replicating the same motions. As viewers navigate through the virtual crowd, the temporal repetition of motion variations might therefore increase, as viewers would come into contact with the characters animated with the same motion in a temporal more than spatial manner.

ANIMATION BY DEMONSTRATION USING VR

In the near future, we are also planning to explore how to further **bring users into the design of virtual characters** by enabling them to **demonstrate interactions using VR**. Today, designing how characters behave and move is typically done by relying on behaviour models, potentially associated with animation data for displaying full-body motions. These models are often manually scripted by designers (e.g., using Behaviour trees or Finite State

Machines), or abstracted from real-world data which can sometimes be difficult to acquire. For instance, creating a locomotion controller requires to capture a number of motion clips and to organize them in such a way that they can be controlled to make the character walk with the desired direction and velocity. However, capturing and organizing motion data is time consuming, and in more complex cases identifying in advance which motion clips might be required can be difficult. It can therefore easily result in missing animation clips to cover the whole animation space required for a specific application, which in turn requires to organise novel motion capture sessions, etc. Similarly, collision avoidance models are often calibrated using crowd datasets, i.e., real-world trajectories from numerous people moving together in a same location, which involves costs, logistical, ethical and technical constraints. We are therefore interested in exploring novel ways of letting users demonstrate relevant interactions by interacting directly with virtual characters using VR, in order to facilitate the collection of human motions and behaviours. To explore these questions, we project to explore the two following directions.

First, we plan to explore new ways to endow virtual characters with novel behaviours by letting users **demonstrate behaviours by interacting directly with virtual characters in immersive situations** (PhD of Vincenzo Abichequer Sangalli, H2020 ITN CLIPE). To illustrate, let us consider the case of a virtual character without any avoidance ability, i.e., who does not know how to avoid other characters or users. Our objective is to let users directly interact with the character, so that they demonstrate to the virtual character how to react in a specific situation to avoid colliding with others. In turn, the character will start to learn a number of interaction abilities, that will enable him/her to more naturally interact with other characters and users. A similar approach could be followed to endow other abilities to the virtual characters, e.g., by demonstrating how one would interact with specific elements of the virtual environment such as in the case of a train station (i.e., identifying how to interact with a ticket machine vs. opening a door). This problem however raises a number of important challenging questions. In particular, one of the main challenge is to appropriately identify which are the relevant parameters describing a specific situation. In our previous example, this would consist in identifying automatically the parameters informing that the virtual character is in a future colliding situation. While there is a large related literature in the case of collision avoidance, there is also the question of how to generalize the approach to more complex examples. Once these relevant parameters are identified, we can also wonder how to appropriately introduce variations in the interaction abilities learned by virtual characters, so that they do not repeat exactly the same motion every time they encounter the same situation (which also creates a nice link with the previous research direction). It also raises the question of how users would identify when and what type of motions they should demonstrate. One solution that we have in mind is to explore an

iterative process where users identify that the character is not behaving appropriately, and could therefore switch to a ‘demonstrator’ mode where they directly demonstrate potential past reactions to the character. This is something that could be performed during the design process, or that could open the possibility of collaboratively elaborating complex human behaviours. This is clearly a difficult problem, nevertheless we believe that such an approach can also be extremely valuable in the creative process, and potentially help designers and animators in creating natural interactions more easily by being immersed in a virtual environment and directly interacting with virtual characters.

Second, we want to explore leveraging VR to acquire varied **synthetic datasets of character interactions** (PhD of Tairan Yin, H2020 ITN CLIPE), involving a number of persons demonstrating interactions in various situations. In many cases, such real-world datasets are crucial to further our understanding of human motions and behaviours, and are used to design, train, calibrate and evaluate models. However, they can sometimes be complex to acquire because of costs, logistical, ethical and technical issues, especially when numerous persons are required. A typical example would be the crowd datasets mentioned above, but could be generalized to any dataset of interactions between persons including body motions, trajectories, sound, etc. We are therefore interested in exploring novel ways of acquiring such complex datasets using VR. We have recently performed our first significant step in this direction by proposing the first prototype of the One-Man-Crowd paradigm [Yin et al. 2022], in which a single user is immersed in virtual scenarios where he/she successively acts each individual member of a virtual crowd. By relying on virtual reality and motion capture techniques, we are able to record the past trajectories and body movements of the user, to display them on virtual characters, and to enable the user to progressively build the overall crowd behaviour by him/herself. We also demonstrated that natural collective behaviours can emerge from virtual crowd data generated using our approach, even though the variety in behaviours is lower than in real situations. While the first results are extremely promising, they also raised a number of questions which we want to further explore. For instance, a major limitation of the OMC paradigm is the presence of a *variety bias*, i.e., some specific individual behaviours are further amplified as a single participant is acting the motions of the whole crowd, instead of being averaged out when each person in the crowd is a unique individual. We are therefore interested in exploring whether this variety bias can be limited, either by finding a way to compute specific interaction characteristics from numerous single-user examples, or by exploring different ways of immersing a limited number N of users, where N would be as small as possible. The other major limitation is that OMC leads to unilateral interactions by nature, since only the immersed participant can react to virtual agents, but not conversely, which we call the *asynchrony bias*. This raises a number of interesting challenges to potentially compensate this limitation, for which we want to

explore several solutions: i) acting each agent individually for a specified time period, then for the following time period and so on; ii) exploring the possibility of acting iteratively a number of agents until convergence of the interaction (e.g., in a dual interaction between agents A and B, let the user controlling A, then B to unilaterally interact with A, and then A again to interact unilaterally with B, etc.); iii) using standard animation techniques to correct previously recorded motions when the user’s action would violate the previously recorded situations. While such possible solutions might enable us to generalise the use of OMC for other types of scenarios, they also raise additional important questions, such as whether alterations would reduce the realism of the generated dataset, whether iterative interactions would converge, etc. Finally, generating such synthetic datasets using VR also raises the general question of their validity and usability in designing, training or calibrating models. While natural characteristics and behaviours can emerge from these situation, it is still unclear whether their usage would be on par with real-world datasets, or to which extend they are able to capture the subtleties that can be found in real-world situations. In any case, the general objective is to explore a) the potential of such approaches to generate synthetic datasets of character interactions, and their validity in comparison to real ones, and b) to generalize their application beyond crowd datasets to explore increasingly more complex and ecological scenarios involving interactions between persons.

AVATAR-FRIENDLY ANIMATION & INTERACTION TECHNIQUES

The last principal research direction which we will be exploring is related to avatars. Despite the advances made in designing and controlling these avatars, as well as in understanding the factors that affect users’ abilities to feel embodied in them (i.e., their sense of virtual embodiment), we are still lacking in **animation and interaction techniques specifically designed for avatars**. For this reason, we are particularly interested in exploring the idea of designing “**Avatar-Friendly**” techniques, which as defined in Chapter 5 refer to *techniques which take the user’s avatar into account in the design process to preserve both performance and the sense of embodiment*. In particular, we are interested in exploring two principal research directions.

First, we are interested in **adapting or proposing interaction techniques compatible with avatars** (PhD of Maé Mavromatis, Inria Challenge Avatar). To interact with virtual environments, a large number of interaction techniques have been proposed, aiming at covering manners of interacting ranging from “natural” interaction methods (i.e., aiming to mimic real-life interactions, such as grabbing a virtual object with a virtual hand), to more “super-natural” interaction methods (i.e., exploiting non-realistic constraints made possible by VR, such as teleporting, flying, grabbing objects at a distance). However, each

interaction technique can potentially disrupt the virtual embodiment, e.g., because of its influence on the visual representation of the avatar, or because of physical inconsistencies with users' real and displayed motions. Similarly, the avatar can also hinder the interaction process, e.g., by occluding virtual objects and therefore hinder their manipulation. Our goal is therefore to further understand the causes of such disruptions, and to propose novel "avatar-friendly" interaction methods, providing effective interaction but compliant with virtual embodiment. For instance, we are interested in exploring the extend to which we can alter users' motions and visual representations when using "super-natural" interaction techniques, while simultaneously maximizing the resulting sense of embodiment.

Second, we are also interested in **adapting existing animation methods**, used to drive the avatar's motions, to **account for virtual embodiment**. As mentioned in Chapter 5, avatar animations are typically driven by traditional character animation techniques, similar to the approaches we rely on in Chapter 2. While motion capture systems provide a higher-quality control over the avatar's motions, it is unrealistic to expect such systems to be deployed in consumer-level applications in a near future. While video-based approaches might become a valuable alternative one day, most applications still rely on animations driven from users' hand-held controllers, HMD and more occasionally additional trackers (e.g., 3-point to 6-point IK). We therefore believe that there is a need to bridge the gap between motion-capture-based and IK-based avatar control, while considering their effect on virtual embodiment. We are interested in exploring recent advances in character animation techniques, e.g. Deep Learning based approaches, which might be valuable for VR and avatar-based applications. In particular, we are interested in exploring how to adapt these novel methods to tailor them for avatar applications, so that we can balance required input (i.e., how much of the user's motions need to be captured) and the resulting sense of embodiment experienced by users. In our opinion, this is a particularly interesting direction to explore given that the quality of the control of the avatar was identified to be a major factor of influence towards the sense of embodiment. In the long term, our goal is also to work towards a convergence of traditional animation approaches at large and avatar-based approach, to reach similar levels of realism by exploring the effect of introducing more subtle details, such as fingers and facial animations.

CONCLUSION

In this manuscript, we presented our contributions towards creating perception-based character animations, i.e., studying and understanding how viewers perceive virtual characters in order to make them increasingly visually more plausible. This section then presented some of the research directions that we are planning to explore in the near future, or are

Chapter 6. Perspectives

already exploring though a number of ongoing research projects and PhDs. We believe that these research directions raise numerous challenges, which when solved will have brought us important fundamental knowledge, technical models, valuable experience, as well as closer to our general objective.

Finally, I would like to thank again all the persons whom I collaborated with, and without whom all this work would not have been possible. All these contributions are the result of hours and hours of fruitful discussions, meetings, brainstorming, relying on the expertise in several complementary fields of a number of colleagues, on the support of several more senior colleagues who have guided and inspired me over the years, as well as on the dedication of the students whom I had the chance to supervise.

References

- Achenbach, Jascha, Thomas Waltemate, Marc Erich Latoschik, and Mario Botsch (2017). “Fast Generation of Realistic Virtual Humans”. In: *Proceedings of the 23rd ACM Symposium on Virtual Reality Software and Technology*. VRST '17. DOI: [10.1145/3139131.3139154](https://doi.org/10.1145/3139131.3139154).
- Adili, Robin, Benjamin Niay, Katja Zibrek, Anne-Hélène Olivier, Julien Pettré, and Ludovic Hoyet (2021). “Perception of Motion Variations in Large-Scale Virtual Human Crowds”. In: *Motion, Interaction and Games*. MIG '21. DOI: [10.1145/3487983.3488288](https://doi.org/10.1145/3487983.3488288).
- Agethen, Philipp, Viswa Subramanian Sekar, Felix Gaisbauer, Thies Pfeiffer, Michael Otto, and Enrico Rukzio (July 2018). “Behavior Analysis of Human Locomotion in the Real World and Virtual Reality for the Manufacturing Industry”. In: *ACM Trans. Appl. Percept.* 15.3. DOI: [10.1145/3230648](https://doi.org/10.1145/3230648).
- Argelaguet-Sanz, Ferran, Anne-Helene Olivier, Gerd Bruder, Julien Pettré, and Anatole Lécuyer (2015). “Virtual proxemics: Locomotion in the presence of obstacles in large immersive projection environments”. In: *2015 IEEE Virtual Reality (VR)*, pp. 75–80. DOI: [10.1109/VR.2015.7223327](https://doi.org/10.1109/VR.2015.7223327).
- Argelaguet, Ferran, Ludovic Hoyet, Michael Trico, and Anatole Lécuyer (2016). “The role of interaction in virtual embodiment: Effects of the virtual hand representation”. In: *2016 IEEE Virtual Reality (VR)*, pp. 3–10. DOI: [10.1109/VR.2016.7504682](https://doi.org/10.1109/VR.2016.7504682).
- Banakou, Domna, Raphaela Groten, and Mel Slater (2013). “Illusory ownership of a virtual child body causes overestimation of object sizes and implicit attitude changes”. In: 110.31, pp. 12846–12851. DOI: [10.1073/pnas.1306779110](https://doi.org/10.1073/pnas.1306779110).
- Beardsworth, T. and T. Buckner (1981). “The ability to recognize oneself from a video recording of ones movements without seeing ones body”. In: *Bulletin of the Psy. Soc.* 18, pp. 19–22.
- Berg, Jur van den, Sachin Patil, Jason Sewall, Dinesh Manocha, and Ming Lin (2008). “Interactive Navigation of Multiple Agents in Crowded Environments”. In: *Proceedings of the 2008 Symposium on Interactive 3D Graphics and Games*. I3D '08, pp. 139–147. DOI: [10.1145/1342250.1342272](https://doi.org/10.1145/1342250.1342272).
- Bergamin, Kevin, Simon Clavet, Daniel Holden, and James Richard Forbes (2019). “DReCon: Data-Driven Responsive Control of Physics-Based Characters”. In: *ACM Trans. Graph.* 38.6. DOI: [10.1145/3355089.3356536](https://doi.org/10.1145/3355089.3356536).
- Berger, Christopher C., Mar Gonzalez-Franco, Eyal Ofek, and Ken Hinckley (2018). “The uncanny valley of haptics”. In: *Science Robotics* 3.17. DOI: [10.1126/scirobotics.aar7010](https://doi.org/10.1126/scirobotics.aar7010).
- Berton, Florian, Fabien Grzeskowiak, Alexandre Bonneau, Alberto Jovane, Marco Aggravi, Ludovic Hoyet, Anne-Helene Olivier, Claudio Pacchierotti, and Julien Pettre (2020a).

- “Crowd Navigation in VR: exploring haptic rendering of collisions”. In: *IEEE Transactions on Visualization and Computer Graphics*. DOI: [10.1109/TVCG.2020.3041341](https://doi.org/10.1109/TVCG.2020.3041341).
- Berton, Florian, Ludovic Hoyet, Anne-Hélène Olivier, Julien Bruneau, Olivier Le Meur, and Julien Pettre (2020b). “Eye-Gaze Activity in Crowds: Impact of Virtual Reality and Density”. In: *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 322–331. DOI: [10.1109/VR46266.2020.00052](https://doi.org/10.1109/VR46266.2020.00052).
- Berton, Florian, Anne-Hélène Olivier, Julien Bruneau, Ludovic Hoyet, and Julien Pettre (2019). “Studying Gaze Behaviour during Collision Avoidance with a Virtual Walker: Influence of the Virtual Reality Setup”. In: *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 717–725. DOI: [10.1109/VR.2019.8798204](https://doi.org/10.1109/VR.2019.8798204).
- Bhargava, Ayush, Kathryn M Lucaites, Leah S Hartman, Hannah Solini, Jeffrey W Bertrand, Andrew C Robb, Christopher C Pagano, and Sabarish V Babu (2020). “Revisiting affordance perception in contemporary virtual reality”. In: *Virtual Reality* 24.4, pp. 713–724. DOI: [10.1007/s10055-020-00432-y](https://doi.org/10.1007/s10055-020-00432-y).
- Botvinick, Matthew and Jonathan Cohen (1998). “Rubber hands ‘feel’ touch that eyes see”. In: *Nature* 391.6669, pp. 756–756. DOI: [10.1038/35784](https://doi.org/10.1038/35784).
- Boucaud, Fabien, Catherine Pelachaud, and Indira Thouvenin (2021). “Decision Model for a Virtual Agent That Can Touch and Be Touched”. In: *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*. AAMAS ’21, pp. 232–241.
- Bovet, Sidney, Henrique Galvan Debarba, Bruno Herbelin, Eray Molla, and Ronan Boulic (2018). “The Critical Role of Self-Contact for Embodiment in Virtual Reality”. In: *IEEE Transactions on Visualization and Computer Graphics* 24.4, pp. 1428–1436. DOI: [10.1109/TVCG.2018.2794658](https://doi.org/10.1109/TVCG.2018.2794658).
- Bruneau, Julien, Anne-Hélène Olivier, and Julien Pettré (2015). “Going Through, Going Around: A Study on Individual Avoidance of Groups”. In: *IEEE Transactions on Visualization and Computer Graphics* 21.4, pp. 520–528. DOI: [10.1109/TVCG.2015.2391862](https://doi.org/10.1109/TVCG.2015.2391862).
- Buttner, Michael (2015). “Motion Matching-The Road to Next-Gen Animation”. In: *Nucl. ai Conference*.
- Chaminade, Thierry, Jessica Hodgins, and Mitsuo Kawato (2007). “Anthropomorphism influences perception of computer-animated characters’ actions”. In: *Social cognitive and affective neuroscience* 2.3, pp. 206–216.
- Cirio, Gabriel, Anne-Hélène Olivier, Maud Marchal, and Julien Pettré (2013). “Kinematic Evaluation of Virtual Walking Trajectories”. In: *IEEE Transactions on Visualization and Computer Graphics* 19.4, pp. 671–680. DOI: [10.1109/TVCG.2013.34](https://doi.org/10.1109/TVCG.2013.34).
- Coros, Stelian, Philippe Beaudoin, and Michiel van de Panne (2010). “Generalized Biped Walking Control”. In: *ACM Trans. Graph.* 29.4. DOI: [10.1145/1778765.1781156](https://doi.org/10.1145/1778765.1781156).
- Costa, Paul and Robert McCrae (2008). *The Revised Neo Personality Inventory (neo-pi-r)*. Sage Publications, Inc, pp. 179–198. DOI: [10.4135/9781849200479.n9](https://doi.org/10.4135/9781849200479.n9).
- Cutting, James E and Lynn T Kozlowski (1977). “Recognizing friends by their walk: Gait perception without familiarity cues”. In: *Bulletin of the Psychonomic society* 9.5, pp. 353–356. DOI: [10.3758/BF03337021](https://doi.org/10.3758/BF03337021).
- David, Nicole, Francesca Fiori, and Salvatore Maria Aglioti (2014). “Susceptibility to the rubber hand illusion does not tell the whole body-awareness story”. In: *Cognitive, Affective and Behavioral Neuroscience* 14, pp. 297–306. DOI: [10.3758/s13415-013-0190-6](https://doi.org/10.3758/s13415-013-0190-6).

References

- Devigne, Louise, Marco Aggravi, Morgane Bivaud, Nathan Balix, Catalin Stefan Teodorescu, Tom Carlson, Tom Spreters, Claudio Pacchierotti, and Marie Babel (2020). “Power Wheelchair Navigation Assistance Using Wearable Vibrotactile Haptics”. In: *IEEE Transactions on Haptics* 13.1, pp. 52–58. DOI: [10.1109/TOH.2019.2963831](https://doi.org/10.1109/TOH.2019.2963831).
- Dewez, Diane, Rebecca Fribourg, Ferran Argelaguet, Ludovic Hoyet, Daniel Mestre, Mel Slater, and Anatole Lécuyer (2019). “Influence of Personality Traits and Body Awareness on the Sense of Embodiment in Virtual Reality”. In: *2019 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 123–134. DOI: [10.1109/ISMAR.2019.00-12](https://doi.org/10.1109/ISMAR.2019.00-12).
- Dewez, Diane, Ludovic Hoyet, Anatole Lécuyer, and Ferran Argelaguet Sanz (2021). “Towards “Avatar-Friendly” 3D Manipulation Techniques: Bridging the Gap Between Sense of Embodiment and Interaction in Virtual Reality”. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. DOI: [10.1145/3411764.3445379](https://doi.org/10.1145/3411764.3445379).
- (2022). “Do You Need Another Hand? Investigating Dual Body Representations During Anisomorphic 3D Manipulation”. In: *To Appear in IEEE VR - Journal Track / IEEE Transactions on Visualization and Computer Graphics*.
- Dewez, Diane, Ludovic Hoyet, Anatole Lécuyer, and Ferran Argelaguet (2020). “Studying the Inter-Relation Between Locomotion Techniques and Embodiment in Virtual Reality”. In: *2020 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pp. 452–461. DOI: [10.1109/ISMAR50242.2020.00070](https://doi.org/10.1109/ISMAR50242.2020.00070).
- Dicks, Matt, Christal Clashing, Liam O’Reilly, and Chris Mills (2016). “Perceptual-motor behaviour during a simulated pedestrian crossing”. In: *Gait & Posture* 49, pp. 241–245. DOI: [10.1016/j.gaitpost.2016.07.003](https://doi.org/10.1016/j.gaitpost.2016.07.003).
- Dummer, Timothy, Alexandra Picot-Annand, Tristan Neal, and Chris Moore (2009). “Movement and the Rubber Hand Illusion”. In: *Perception* 38.2, pp. 271–280. DOI: [10.1068/p5921](https://doi.org/10.1068/p5921).
- Durupinar, Funda (2021). “Perception of Human Motion Similarity Based on Laban Movement Analysis”. In: *ACM Symposium on Applied Perception 2021*. SAP ’21. DOI: [10.1145/3474451.3476241](https://doi.org/10.1145/3474451.3476241).
- Durupinar, Funda, Mubbasir Kapadia, Susan Deutsch, Michael Neff, and Norman I. Badler (2016). “PERFORM: Perceptual Approach for Adding OCEAN Personality to Human Motion Using Laban Movement Analysis”. In: *ACM Trans. Graph.* 36.1. DOI: [10.1145/2983620](https://doi.org/10.1145/2983620).
- Dutra, Teófilo, Ricardo Marques, Joaquim Cavalcante-Neto, Creto Vidal, and Julien Pettre (May 2017). “Gradient-based steering for vision-based crowd simulation algorithms”. In: *Computer Graphics Forum* 36. DOI: [10.1111/cgf.13130](https://doi.org/10.1111/cgf.13130).
- Edwards, S. J., D. J. Buckland, and J. D. McCoy-Powlen (2002). *Developmental & Functional Hand Grasps*. Thorofare, New Jersey: Slack Incorporated.
- Ehrsson, H. Henrik, Katja Wiech, Nikolaus Weiskopf, Raymond J. Dolan, and Richard E. Passingham (2007). “Threatening a rubber hand that you feel is yours elicits a cortical anxiety response”. In: *Proceedings of the National Academy of Sciences* 104.23, pp. 9828–9833. DOI: [10.1073/pnas.0610011104](https://doi.org/10.1073/pnas.0610011104).
- Ennis, Cathy, Ludovic Hoyet, Arjan Egges, and Rachel McDonnell (2013). “Emotion Capture: Emotionally Expressive Characters for Games”. In: *Proceedings of Motion on Games*. MIG ’13, pp. 53–60. DOI: [10.1145/2522628.2522633](https://doi.org/10.1145/2522628.2522633).

- Ennis, Cathy, Ludovic Hoyet, and Carol O’Sullivan (2015). “Eye-tracktive: Measuring Attention to Body Parts when Judging Human Motions”. In: *EG 2015 - Short Papers*. DOI: [10.2312/egsh.20151009](https://doi.org/10.2312/egsh.20151009).
- Fink, Philip W., Patrick S. Foo, and William H. Warren (2007). “Obstacle Avoidance during Walking in Real and Virtual Environments”. In: *ACM Trans. Appl. Percept.* 4.1. DOI: [10.1145/1227134.1227136](https://doi.org/10.1145/1227134.1227136).
- Fribourg, Rebecca, Ferran Argelaguet, Ludovic Hoyet, and Anatole Lécuyer (2018). “Studying the Sense of Embodiment in VR Shared Experiences”. In: *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 273–280. DOI: [10.1109/VR.2018.8448293](https://doi.org/10.1109/VR.2018.8448293).
- Fribourg, Rebecca, Ferran Argelaguet, Anatole Lécuyer, and Ludovic Hoyet (2020a). “Avatar and Sense of Embodiment: Studying the Relative Preference Between Appearance, Control and Point of View”. In: *IEEE Transactions on Visualization and Computer Graphics* 26.5, pp. 2062–2072. DOI: [10.1109/TVCG.2020.2973077](https://doi.org/10.1109/TVCG.2020.2973077).
- Fribourg, Rebecca, Evan Blanpied, Ludovic Hoyet, Anatole Lécuyer, and Ferran Argelaguet (2020b). “Influence of Threat Occurrence and Repeatability on the Sense of Embodiment and Threat Response in VR”. In: *ICAT-EGVE 2020 - International Conference on Artificial Reality and Telexistence and Eurographics Symposium on Virtual Environments*. DOI: [10.2312/egve.20201262](https://doi.org/10.2312/egve.20201262).
- (2021a). “Does virtual threat harm VR experience?: Impact of threat occurrence and repeatability on virtual embodiment and threat response”. In: *Computers & Graphics* 100, pp. 125–136. DOI: [10.1016/j.cag.2021.07.017](https://doi.org/10.1016/j.cag.2021.07.017).
- Fribourg, Rebecca, Nami Ogawa, Ludovic Hoyet, Ferran Argelaguet, Takuji Narumi, Michitaka Hirose, and Anatole Lécuyer (2021b). “Virtual Co-Embodiment: Evaluation of the Sense of Agency While Sharing the Control of a Virtual Body Among Two Individuals”. In: *IEEE Transactions on Visualization and Computer Graphics* 27.10, pp. 4023–4038. DOI: [10.1109/TVCG.2020.2999197](https://doi.org/10.1109/TVCG.2020.2999197).
- Fugard, A.J.B., N. Pfeifer, B. Mayerhofer, and G.D. Kleiter (2011). “How people interpret conditionals: shifts toward the conditional event.” In: *J Exp Psy Learn Mem Cogn* 37.3, pp. 635–48.
- Galvan Debarba, Henrique, Sidney Bovet, Roy Salomon, Olaf Blanke, Bruno Herbelin, and Ronan Boulic (2017). “Characterizing first and third person viewpoints and their alternation for embodied interaction in virtual reality”. In: *PLOS ONE* 12.12. DOI: [10.1371/journal.pone.0190109](https://doi.org/10.1371/journal.pone.0190109).
- Gérin-Lajoie, Martin, Carol L. Richards, Joyce Fung, and Bradford J. McFadyen (2008). “Characteristics of personal space during obstacle circumvention in physical and virtual environments”. In: *Gait & Posture* 27.2. DOI: <https://doi.org/10.1016/j.gaitpost.2007.03.015>.
- Giorgolo, G. and F. Verstraten (2008). “Perception of speech-and-gesture integration”. In: *Proc. of the International Conference on Auditory-Visual Speech Processing*, pp. 31–36.
- Gonzalez-Franco, Mar and Tabitha C. Peck (2018). “Avatar Embodiment. Towards a Standardized Questionnaire”. In: *Frontiers in Robotics and AI* 5, p. 74. DOI: [10.3389/frobt.2018.00074](https://doi.org/10.3389/frobt.2018.00074).
- Gorisse, Geoffrey, Olivier Christmann, Etienne Armand Amato, and Simon Richir (2017). “First- and Third-Person Perspectives in Immersive Virtual Environments: Presence and

References

- Performance Analysis of Embodied Users”. In: *Frontiers in Robotics and AI* 4, p. 33. DOI: [10.3389/frobt.2017.00033](https://doi.org/10.3389/frobt.2017.00033).
- Gorisse, Geoffrey, Olivier Christmann, Samory Houzangbe, and Simon Richir (2019). “From Robot to Virtual Doppelgänger: Impact of Visual Fidelity of Avatars Controlled in Third-Person Perspective on Embodiment and Behavior in Immersive Virtual Environments”. In: *Frontiers in Robotics and AI* 6, p. 8. DOI: [10.3389/frobt.2019.00008](https://doi.org/10.3389/frobt.2019.00008).
- Gosling, Samuel D, Peter J Rentfrow, and William B Swann (2003). “A very brief measure of the Big-Five personality domains”. In: *Journal of Research in Personality* 37.6, pp. 504–528. DOI: [https://doi.org/10.1016/S0092-6566\(03\)00046-1](https://doi.org/10.1016/S0092-6566(03)00046-1).
- Grammer, Karl, Viktoria Keki, Beate Striebel, Michaela Atzmüller, and Bernhard Fink (2003). “Bodies in Motion: A Window to the Soul”. In: *Evolutionary Aesthetics*, pp. 295–323.
- Helbing, Dirk and Péter Molnár (1995). “Social force model for pedestrian dynamics”. In: *Physical Review E* 51.5, pp. 4282–4286. DOI: [10.1103/physreve.51.4282](https://doi.org/10.1103/physreve.51.4282).
- Ho, Edmond S. L., Taku Komura, and Chiew-Lan Tai (2010). “Spatial Relationship Preserving Character Motion Adaptation”. In: *ACM Trans. Graph.* 29.4. DOI: [10.1145/1778765.1778770](https://doi.org/10.1145/1778765.1778770).
- Hodgins, Jessica, Sophie Jörg, Carol O’Sullivan, Sang Il Park, and Moshe Mahler (2010). “The Saliency of Anomalies in Animated Human Characters”. In: *ACM Trans. Appl. Percept.* 7.4. DOI: [10.1145/1823738.1823740](https://doi.org/10.1145/1823738.1823740).
- Holden, Daniel, Oussama Kanoun, Maksym Perepichka, and Tiberiu Popa (2020). “Learned Motion Matching”. In: *ACM Trans. Graph.* 39.4. DOI: [10.1145/3386569.3392440](https://doi.org/10.1145/3386569.3392440).
- Holden, Daniel, Jun Saito, and Taku Komura (2016). “A Deep Learning Framework for Character Motion Synthesis and Editing”. In: *ACM Trans. Graph.* 35.4. DOI: [10.1145/2897824.2925975](https://doi.org/10.1145/2897824.2925975).
- Hoyet, Ludovic, Ferran Argelaguet, Corentin Nicole, and Anatole Lécuyer (2016a). ““Wow! I Have Six Fingers!”: Would You Accept Structural Changes of Your Hand in VR?” In: *Frontiers in Robotics and AI* 3, p. 27. DOI: [10.3389/frobt.2016.00027](https://doi.org/10.3389/frobt.2016.00027).
- Hoyet, Ludovic, Rachel McDonnell, and Carol O’Sullivan (2012a). “Push It Real: Perceiving Causality in Virtual Interactions”. In: *ACM Trans. Graph.* 31.4. DOI: [10.1145/2185520.2185586](https://doi.org/10.1145/2185520.2185586).
- Hoyet, Ludovic, Franck Multon, Taku Komura, and Anatole Lécuyer (2010a). “Perception Based Real-Time Dynamic Adaptation of Human Motions”. In: *Motion in Games*, pp. 266–277.
- Hoyet, Ludovic, Franck Multon, Anatole Lécuyer, and Taku Komura (2010b). “Can We Distinguish Biological Motions of Virtual Humans? Perceptual Study with Captured Motions of Weight Lifting”. In: *Proceedings of the 17th ACM Symposium on Virtual Reality Software and Technology*. VRST ’10, pp. 87–90. DOI: [10.1145/1889863.1889878](https://doi.org/10.1145/1889863.1889878).
- Hoyet, Ludovic, Anne-Helene Olivier, Richard Kulpa, and Julien Pettré (2016b). “Perceptual Effect of Shoulder Motions on Crowd Animations”. In: *ACM Trans. Graph.* 35.4. DOI: [10.1145/2897824.2925931](https://doi.org/10.1145/2897824.2925931).
- Hoyet, Ludovic, Kenneth Ryall, Rachel McDonnell, and Carol O’Sullivan (2012b). “Sleight of Hand: Perception of Finger Motion from Reduced Marker Sets”. In: *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*. I3D ’12, pp. 79–86. DOI: [10.1145/2159616.2159630](https://doi.org/10.1145/2159616.2159630).

- Hoyet, Ludovic, Kenneth Ryall, Katja Zibrek, Hwangpil Park, Jehee Lee, Jessica Hodgins, and Carol O’Sullivan (2013). “Evaluating the Distinctiveness and Attractiveness of Human Motions on Realistic Virtual Bodies”. In: *ACM Trans. Graph.* 32.6. DOI: [10.1145/2508363.2508367](https://doi.org/10.1145/2508363.2508367).
- Hsu, Eugene, Kari Pulli, and Jovan Popović (2005). “Style Translation for Human Motion”. In: *ACM Trans. Graph.* 24.3, pp. 1082–1089. DOI: [10.1145/1073204.1073315](https://doi.org/10.1145/1073204.1073315).
- Jeunet, Camille, Louis Albert, Ferran Argelaguet, and Anatole Lécuyer (2018). ““Do You Feel in Control?”: Towards Novel Approaches to Characterise, Manipulate and Measure the Sense of Agency in Virtual Environments”. In: *IEEE Transactions on Visualization and Computer Graphics* 24.4, pp. 1486–1495. DOI: [10.1109/TVCG.2018.2794598](https://doi.org/10.1109/TVCG.2018.2794598).
- Jin, Taeil, Meekyoung Kim, and Sung-Hee Lee (2018). “Aura Mesh: Motion Retargeting to Preserve the Spatial Relationships between Skinned Characters”. In: *Computer Graphics Forum* 37.2, pp. 311–320. DOI: [10.1111/cgf.13363](https://doi.org/10.1111/cgf.13363).
- Johansson, Gunnar (1973). “Visual perception of biological motion and a model for its analysis”. In: *Perception & Psychophysics* 14.2, pp. 201–211. DOI: [10.3758/BF03212378](https://doi.org/10.3758/BF03212378).
- John, Oliver P, Eileen M Donahue, and Robert L Kentle (1991). “Big five inventory”. In: *Journal of Personality and Social Psychology*. DOI: [10.1037/t07550-000](https://doi.org/10.1037/t07550-000).
- John, Oliver P, Sanjay Srivastava, et al. (1999). *The Big-Five trait taxonomy: History, measurement, and theoretical perspectives*. Vol. 2. Guilford Press.
- Johnson, K. L. and L. G. Tassinary (2005). “Perceiving sex directly and indirectly: Meaning in motion and morphology”. In: *Psychological Science* 16.11, pp. 890–897.
- (2007). “Compatibility of basic social perceptions determines perceived attractiveness”. In: *Proc. of the National Academy of Sciences* 104.12, pp. 5246–5251.
- Jörg, Sophie, Jessica Hodgins, and Carol O’Sullivan (2010). “The Perception of Finger Motions”. In: *Proceedings of the 7th Symposium on Applied Perception in Graphics and Visualization*. APGV ’10, pp. 129–133. DOI: [10.1145/1836248.1836273](https://doi.org/10.1145/1836248.1836273).
- Jörg, Sophie, Jessica Hodgins, and Alla Safonova (2012). “Data-Driven Finger Motion Synthesis for Gesturing Characters”. In: *ACM Trans. Graph.* 31.6. DOI: [10.1145/2366145.2366208](https://doi.org/10.1145/2366145.2366208).
- Kalckert, Andreas and Henrik H. Ehrsson (2012). “Moving a Rubber Hand that Feels Like Your Own: A Dissociation of Ownership and Agency”. In: *Frontiers in Human Neuroscience* 6, p. 40. DOI: [10.3389/fnhum.2012.00040](https://doi.org/10.3389/fnhum.2012.00040).
- Kiiski, Hanni, Ludovic Hoyet, Andy T. Woods, Carol O’Sullivan, and Fiona N. Newell (2015). “Strutting Hero, Sneaking Villain: Utilizing Body Motion Cues to Predict the Intentions of Others”. In: *ACM Trans. Appl. Percept.* 13.1. DOI: [10.1145/2791293](https://doi.org/10.1145/2791293).
- Kilteni, Konstantina, Raphaela Groten, and Mel Slater (2012). “The Sense of Embodiment in Virtual Reality”. In: *Presence* 21.4, pp. 373–387. DOI: [10.1162/PRES_a_00124](https://doi.org/10.1162/PRES_a_00124).
- Kokkinara, Elena and Mel Slater (2014). “Measuring the effects through time of the influence of visuomotor and visuotactile synchronous stimulation on a virtual body ownership illusion”. In: *Perception* 43.1, pp. 43–58. DOI: [10.1068/p7545](https://doi.org/10.1068/p7545).
- Kovar, Lucas, Michael Gleicher, and Frédéric Pighin (2002). “Motion Graphs”. In: *ACM Trans. Graph.* 21.3, pp. 473–482. DOI: [10.1145/566654.566605](https://doi.org/10.1145/566654.566605).
- Krogmeier, Claudia, Christos Mousas, and David Whittinghill (2019). “Human, Virtual Human, Bump! A Preliminary Study on Haptic Feedback”. In: *2019 IEEE Conference*

References

- on *Virtual Reality and 3D User Interfaces (VR)*, pp. 1032–1033. DOI: [10.1109/VR.2019.8798139](https://doi.org/10.1109/VR.2019.8798139).
- Kulpa, Richard, Anne-Hélène Olivierxs, Jan Ondřej, and Julien Pettré (2011). “Imperceptible Relaxation of Collision Avoidance Constraints in Virtual Crowds”. In: *ACM Trans. Graph.* 30.6, pp. 1–10. DOI: [10.1145/2070781.2024172](https://doi.org/10.1145/2070781.2024172).
- Laidacker, Aleissia and Nicolas Barbeau (2011). “Living Crowds: AI & Animation in Assassin’s Creed: Brotherhood”. In: *Presented at Game Developers Conference, San Francisco, CA, Feb 28-Mar 4*.
- Latoschik, Marc Erich, Daniel Roth, Dominik Gall, Jascha Achenbach, Thomas Waltemate, and Mario Botsch (2017). “The Effect of Avatar Realism in Immersive Social Virtual Realities”. In: *Proceedings of the 23rd ACM Symposium on Virtual Reality Software and Technology. VRST ’17*. DOI: [10.1145/3139131.3139156](https://doi.org/10.1145/3139131.3139156).
- Lau, Manfred, Ziv Bar-Joseph, and James Kuffner (2009). “Modeling Spatial and Temporal Variation in Motion Data”. In: *ACM Trans. Graph.* 28.5, pp. 1–10. DOI: [10.1145/1618452.1618517](https://doi.org/10.1145/1618452.1618517).
- Lee, Kang Hoon, Myung Geol Choi, and Jehée Lee (2006). “Motion Patches: Building Blocks for Virtual Environments Annotated with Motion Data”. In: *ACM Trans. Graph.* 25.3, pp. 898–906. DOI: [10.1145/1141911.1141972](https://doi.org/10.1145/1141911.1141972).
- Levenson, Hanna (1981). “Differentiating among internality, powerful others, and chance”. en. In: *Research with the locus of control construct*. Ed. by H. M. Lefcourt. Vol. 1. New York: Academic Press, pp. 15–63.
- Lin, Qiufeng, John Rieser, and Bobby Bodenheimer (2015). “Affordance Judgments in HMD-Based Virtual Environments: Stepping over a Pole and Stepping off a Ledge”. In: *ACM Trans. Appl. Percept.* 12.2. DOI: [10.1145/2720020](https://doi.org/10.1145/2720020).
- Ling, Yun, Harold T. Nefs, Willem-Paul Brinkman, Chao Qu, and Ingrid Heynderickx (2013). “The relationship between individual characteristics and experienced presence”. In: *Computers in Human Behavior* 29.4, pp. 1519–1530. DOI: <https://doi.org/10.1016/j.chb.2012.12.010>.
- Liu, C. Karen, Aaron Hertzmann, and Zoran Popović (2006). “Composition of Complex Optimal Multi-Character Motions”. In: *Proceedings of the 2006 ACM SIGGRAPH/Eurographics Symposium on Computer Animation. SCA ’06*, pp. 215–222.
- Llobera, Joan, Bernhard Spanlang, Giulio Ruffini, and Mel Slater (2010). “Proxemics with Multiple Dynamic Characters in an Immersive Virtual Environment”. In: *ACM Trans. Appl. Percept.* 8.1. DOI: [10.1145/1857893.1857896](https://doi.org/10.1145/1857893.1857896).
- Lugrin, Jean-Luc, Johanna Latt, and Marc Erich Latoschik (2015). “Anthropomorphism and Illusion of Virtual Body Ownership”. In: *ICAT-EGVE 2015 - International Conference on Artificial Reality and Telexistence and Eurographics Symposium on Virtual Environments*. DOI: [10.2312/egve.20151303](https://doi.org/10.2312/egve.20151303).
- Lynch, Sean Dean, Richard Kulpa, Laurentius Antonius Meerhoff, Julien Pettré, Armel Crétual, and Anne-Hélène Olivier (2018). “Collision Avoidance Behavior between Walkers: Global and Local Motion Cues”. In: *IEEE Transactions on Visualization and Computer Graphics* 24.7, pp. 2078–2088. DOI: [10.1109/TVCG.2017.2718514](https://doi.org/10.1109/TVCG.2017.2718514).
- Majkowska, A., V. B. Zordan, and P. Faloutsos (2006). “Automatic Splicing for Hand and Body Animations”. In: *Proceedings of the 2006 ACM SIGGRAPH/Eurographics Symposium on Computer Animation. SCA ’06*, pp. 309–316.

- Maselli, Antonella and Mel Slater (2013). “The building blocks of the full body ownership illusion”. In: *Frontiers in Human Neuroscience* 7, p. 83. DOI: [10.3389/fnhum.2013.00083](https://doi.org/10.3389/fnhum.2013.00083).
- McDonnell, Rachel, Martin Breidt, and Heinrich H. Bühlhoff (2012). “Render Me Real? Investigating the Effect of Render Style on the Perception of Animated Virtual Humans”. In: *ACM Trans. Graph.* 31.4. DOI: [10.1145/2185520.2185587](https://doi.org/10.1145/2185520.2185587).
- McDonnell, Rachel, Michéal Larkin, Simon Dobbyn, Steven Collins, and Carol O’Sullivan (2008). “Clone Attack! Perception of Crowd Variety”. In: *ACM Trans. Graph.* 27.3, pp. 1–8. DOI: [10.1145/1360612.1360625](https://doi.org/10.1145/1360612.1360625).
- McManus, Erin A., Bobby Bodenheimer, Stephan Streuber, Stephan de la Rosa, Heinrich H. Bühlhoff, and Betty J. Mohler (2011). “The Influence of Avatar (Self and Character) Animations on Distance Estimation, Object Interaction and Locomotion in Immersive Virtual Environments”. In: *Proceedings of the ACM SIGGRAPH Symposium on Applied Perception in Graphics and Visualization*. APGV ’11, pp. 37–44. DOI: [10.1145/2077451.2077458](https://doi.org/10.1145/2077451.2077458).
- McNeill, David (2005). *Gesture and Thought*. University of Chicago Press.
- Meerhoff, Laurentius Antonius, Julien Bruneau, Alexandre Vu, Anne-Hélène Olivier, and Julien Pettré (2018a). “Guided by gaze: Prioritization strategy when navigating through a virtual crowd can be assessed through gaze activity”. In: *Acta Psychologica* 190, pp. 248–257. DOI: [10.1016/j.actpsy.2018.07.009](https://doi.org/10.1016/j.actpsy.2018.07.009).
- Meerhoff, Laurentius Antonius, Julien Pettré, Sean Dean Lynch, Armel Crétual, and Anne-Hélène Olivier (2018b). “Collision Avoidance With Multiple Walkers: Sequential or Simultaneous Interactions?” In: *Frontiers in Psychology* 9, p. 2354. DOI: [10.3389/fpsyg.2018.02354](https://doi.org/10.3389/fpsyg.2018.02354).
- Mestre, Daniel R., Céphise Louison, and Fabien Ferlay (2016). “The Contribution of a Virtual Self and Vibrotactile Feedback to Walking Through Virtual Apertures”. In: *Human-Computer Interaction. Interaction Platforms and Techniques*. Ed. by Masaaki Kurosu. Springer International Publishing, pp. 222–232. DOI: [10.1007/978-3-319-39516-6_21](https://doi.org/10.1007/978-3-319-39516-6_21).
- Michotte, Albert (1963). *The perception of causality*. Basic Books, New York.
- Min, Jianyuan and Jinxiang Chai (2012). “Motion Graphs++: A Compact Generative Model for Semantic Motion Analysis and Synthesis”. In: *ACM Trans. Graph.* 31.6. DOI: [10.1145/2366145.2366172](https://doi.org/10.1145/2366145.2366172).
- Min, Jianyuan, Huajun Liu, and Jinxiang Chai (2010). “Synthesis and Editing of Personalized Stylistic Human Motion”. In: *Proceedings of the 2010 ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, pp. 39–46. DOI: [10.1145/1730804.1730811](https://doi.org/10.1145/1730804.1730811).
- Mousas, Christos, Alexandros Koiliias, Banafsheh Rekabdar, Dominic Kao, and Dimitris Anastaslou (2021). “Toward Understanding the Effects of Virtual Character Appearance on Avoidance Movement Behavior”. In: *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*, pp. 40–49. DOI: [10.1109/VR50410.2021.00024](https://doi.org/10.1109/VR50410.2021.00024).
- Murray, Craig D., Jez Fox, and Steve Pettifer (2007). “Absorption, dissociation, locus of control and presence in virtual reality”. In: *Computers in Human Behavior* 23.3, pp. 1347–1354. DOI: <https://doi.org/10.1016/j.chb.2004.12.010>.
- Niay, Benjamin, Anne-Hélène Olivier, Katja Zibrek, Julien Pettré, and Ludovic Hoyet (2020). “Walk Ratio: Perception of an Invariant Parameter of Human Walk on Virtual Characters”.

References

- In: *ACM Symposium on Applied Perception 2020*. SAP '20. DOI: [10.1145/3385955.3407926](https://doi.org/10.1145/3385955.3407926).
- O'Sullivan, Carol, John Dingliana, Thanh Giang, and Mary K. Kaiser (2003). "Evaluating the Visual Fidelity of Physically Based Animations". In: *ACM Trans. Graph.* 22.3, pp. 527–536. DOI: [10.1145/882262.882303](https://doi.org/10.1145/882262.882303).
- Olivier, Anne-Hélène, Julien Bruneau, Richard Kulpa, and Julien Pettré (2018). "Walking with Virtual People: Evaluation of Locomotion Interfaces in Dynamic Environments". In: *IEEE Transactions on Visualization and Computer Graphics* 24.7, pp. 2251–2263. DOI: [10.1109/TVCG.2017.2714665](https://doi.org/10.1109/TVCG.2017.2714665).
- Olivier, Anne-Hélène, Antoine Marin, Armel Crétual, Alain Berthoz, and Julien Pettré (2013). "Collision avoidance between two walkers: Role-dependent strategies". In: *Gait & Posture* 38.4, pp. 751–756. DOI: [10.1016/j.gaitpost.2013.03.017](https://doi.org/10.1016/j.gaitpost.2013.03.017).
- Olivier, Anne-Hélène, Antoine Marin, Armel Crétual, and Julien Pettré (2012). "Minimal predicted distance: a common metric for collision avoidance during pairwise interactions between walkers." In: *Gait & posture* 36.3, pp. 399–404. DOI: [10.1016/j.gaitpost.2012.03.021](https://doi.org/10.1016/j.gaitpost.2012.03.021).
- Parger, Mathias, Joerg H. Mueller, Dieter Schmalstieg, and Markus Steinberger (2018). "Human Upper-Body Inverse Kinematics for Increased Embodiment in Consumer-Grade Virtual Reality". In: *Proceedings of the 24th ACM Symposium on Virtual Reality Software and Technology*. VRST '18. DOI: [10.1145/3281505.3281529](https://doi.org/10.1145/3281505.3281529).
- Peck, Tabitha C. and Mar Gonzalez-Franco (2021). "Avatar Embodiment. A Standardized Questionnaire". In: *Frontiers in Virtual Reality* 1, p. 44. DOI: [10.3389/frvir.2020.575943](https://doi.org/10.3389/frvir.2020.575943).
- Pfeil, Kevin, Eugene M. Taranta, Arun Kulshreshth, Pamela Wisniewski, and Joseph J. LaViola (2018). "A Comparison of Eye-Head Coordination between Virtual and Physical Realities". In: *Proceedings of the 15th ACM Symposium on Applied Perception*. SAP '18. DOI: [10.1145/3225153.3225157](https://doi.org/10.1145/3225153.3225157).
- Pollick, Frank E, Jim W Kay, Katrin Heim, and Rebecca Stringer (2005). "Gender recognition from point-light walkers." In: *Journal of Experimental Psychology: Human Perception and Performance* 31.6, pp. 1247–1265. DOI: [10.1037/0096-1523.31.6.1247](https://doi.org/10.1037/0096-1523.31.6.1247).
- Pražák, Martin, Ludovic Hoyet, and Carol O'Sullivan (2011). "Perceptual Evaluation of Footskate Cleanup". In: *Proceedings of the 2011 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. SCA '11, pp. 287–294. DOI: [10.1145/2019406.2019444](https://doi.org/10.1145/2019406.2019444).
- Raimbaud, Pierre, Alberto Jovane, Katja Zibrek, Claudio Pacchierotti, Marc Christie, Ludovic Hoyet, Julien Pettré, and Anne-Hélène Olivier (2022). "The stare in-the-crowd effect in virtual reality". In: *To Appear in IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*.
- Reitsma, Paul S. A. and Carol O'Sullivan (2009). "Effect of Scenario on Perceptual Sensitivity to Errors in Animation". In: *ACM Trans. Appl. Percept.* 6.3. DOI: [10.1145/1577755.1577758](https://doi.org/10.1145/1577755.1577758).
- Renner, Rebekka S., Boris M. Velichkovsky, and Jens R. Helmert (Dec. 2013). "The Perception of Egocentric Distances in Virtual Environments - A Review". In: *ACM Comput. Surv.* 46.2. DOI: [10.1145/2543581.2543590](https://doi.org/10.1145/2543581.2543590).

- Rhodes, Gillian (2006). “The Evolutionary Psychology of Facial Beauty”. In: *Annu. Rev. Psychol.* 57, pp. 199–226.
- Rijpkema, Hans and Michael Girard (July 1991). “Computer Animation of Knowledge-Based Human Grasping”. In: *SIGGRAPH Comput. Graph.* 25.4, pp. 339–348. DOI: [10.1145/127719.122754](https://doi.org/10.1145/127719.122754).
- Rio, Kevin W., Gregory C. Dachner, and William H. Warren (2018). “Local interactions underlying collective motion in human crowds”. In: *Proceedings of the Royal Society B: Biological Sciences* 285.1878. DOI: [10.1098/rspb.2018.0611](https://doi.org/10.1098/rspb.2018.0611).
- Rio, Kevin W. and William H. Warren (2014). “The Visual Coupling between Neighbors in Real and Virtual Crowds”. In: *Transportation Research Procedia* 2, pp. 132–140. DOI: [10.1016/j.trpro.2014.09.017](https://doi.org/10.1016/j.trpro.2014.09.017).
- Ríos, Alejandro and Nuria Pelechano (2020). “Follower behavior under stress in immersive VR”. In: *Virtual Reality* 24.4, pp. 683–694. DOI: [10.1007/s10055-020-00428-8](https://doi.org/10.1007/s10055-020-00428-8).
- Roth, Daniel and Marc Erich Latoschik (2020). “Construction of the Virtual Embodiment Questionnaire (VEQ)”. In: *IEEE Transactions on Visualization and Computer Graphics* 26.12, pp. 3546–3556. DOI: [10.1109/TVCG.2020.3023603](https://doi.org/10.1109/TVCG.2020.3023603).
- Rybarczyk, Yves, Tiago Coelho, Tiago Cardoso, and Rita F. de Oliveira (2014). “Effect of avatars and viewpoints on performance in virtual world: efficiency vs. telepresence”. In: *EAI Endorsed Transactions on Creative Technologies* 1.1. DOI: [10.4108/ct.1.1.e4](https://doi.org/10.4108/ct.1.1.e4).
- Sacau, Ana, Jari Laarni, Niklas Ravaja, and Tilo Hartmann (2005). “The impact of personality factors on the experience of spatial presence”. In: *Conference Proceedings of the 8th International Workshop on Presence*, pp. 143–151.
- Seyfried, Armin, Bernhard Steffen, Wolfram Klingsch, and Maik Boltes (Oct. 2005). “The fundamental diagram of pedestrian movement revisited”. In: 2005.10. DOI: [10.1088/1742-5468/2005/10/p10002](https://doi.org/10.1088/1742-5468/2005/10/p10002).
- Shields, Stephanie, Mary E. Mallory, and Angela Simon (1989). “The Body Awareness Questionnaire: Reliability and Validity”. In: *Journal of Personality Assessment* 53, pp. 802–815.
- Shum, Hubert P. H., Taku Komura, Masashi Shiraishi, and Shuntaro Yamazaki (2008). “Interaction Patches for Multi-Character Animation”. In: *ACM SIGGRAPH Asia 2008 Papers*. SIGGRAPH Asia '08. DOI: [10.1145/1457515.1409067](https://doi.org/10.1145/1457515.1409067).
- Skarbez, Richard, Solène Neyret, Frederick P. Brooks, Mel Slater, and Mary C. Whitton (2017). “A Psychophysical Experiment Regarding Components of the Plausibility Illusion”. In: *IEEE Transactions on Visualization and Computer Graphics* 23.4, pp. 1369–1378. DOI: [10.1109/TVCG.2017.2657158](https://doi.org/10.1109/TVCG.2017.2657158).
- Slater, Mel (2009). “Place illusion and plausibility can lead to realistic behaviour in immersive virtual environments”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 364.1535, pp. 3549–3557. DOI: [10.1098/rstb.2009.0138](https://doi.org/10.1098/rstb.2009.0138).
- Slater, Mel, Daniel Pérez Marcos, Henrik Ehrsson, and Maria Sanchez-Vives (2008). “Towards a digital body: the virtual arm illusion”. In: *Frontiers in Human Neuroscience* 2. DOI: [10.3389/neuro.09.006.2008](https://doi.org/10.3389/neuro.09.006.2008).
- Slater, Mel, Bernhard Spanlang, and David Corominas (2010a). “Simulating Virtual Environments within Virtual Environments as the Basis for a Psychophysics of Presence”. In: *ACM Trans. Graph.* 29.4. DOI: [10.1145/1778765.1778829](https://doi.org/10.1145/1778765.1778829).

References

- Slater, Mel, Bernhard Spanlang, Maria V. Sanchez-Vives, and Olaf Blanke (2010b). “First Person Experience of Body Transfer in Virtual Reality”. In: *PLOS ONE* 5.5. DOI: [10.1371/journal.pone.0010564](https://doi.org/10.1371/journal.pone.0010564).
- Slater, Mel and Anthony Steed (2000). “A Virtual Presence Counter”. In: *Presence: Teleoperators and Virtual Environments* 9.5, pp. 413–434. DOI: [10.1162/105474600566925](https://doi.org/10.1162/105474600566925).
- Starke, Sebastian, He Zhang, Taku Komura, and Jun Saito (2019). “Neural State Machine for Character-Scene Interactions”. In: *ACM Trans. Graph.* 38.6. DOI: [10.1145/3355089.3356505](https://doi.org/10.1145/3355089.3356505).
- Starke, Sebastian, Yiwei Zhao, Taku Komura, and Kazi Zaman (2020). “Local Motion Phases for Learning Multi-Contact Character Movements”. In: *ACM Trans. Graph.* 39.4. DOI: [10.1145/3386569.3392450](https://doi.org/10.1145/3386569.3392450).
- Starke, Sebastian, Yiwei Zhao, Fabio Zinno, and Taku Komura (2021). “Neural Animation Layering for Synthesizing Martial Arts Movements”. In: *ACM Trans. Graph.* 40.4. DOI: [10.1145/3450626.3459881](https://doi.org/10.1145/3450626.3459881).
- Tang, Jeff K. T., Howard Leung, Taku Komura, and Hubert P. H. Shum (2008). “Emulating human perception of motion similarity”. In: *Computer Animation and Virtual Worlds* 19.3-4, pp. 211–221. DOI: [10.1002/cav.260](https://doi.org/10.1002/cav.260).
- Thalmann, Daniel (2007). *Crowd Simulation*. DOI: [10.1002/9780470050118.ecse676](https://doi.org/10.1002/9780470050118.ecse676).
- Therien, James and Sylvain Bernard (2008). “Taming the Mob: Creating believable crowds in Assassin’s Creed”. In: *Presented at Game Developers Conference, San Francisco, CA, Feb 18-22*.
- Troje, Nikolaus F. (2002). “Decomposing biological motion: A framework for analysis and synthesis of human gait patterns”. In: *Journal of Vision* 2.5, pp. 371–87. DOI: [10.1167/2.5.2](https://doi.org/10.1167/2.5.2).
- Tsakiris, Manos, Gita Prabhu, and Patrick Haggard (2006). “Having a body versus moving your body: How agency structures body-ownership”. In: *Consciousness and Cognition* 15.2, pp. 423–432. DOI: <https://doi.org/10.1016/j.concog.2005.09.004>.
- Urtasun, Raquel, Pascal Ghardon, Ronan Boulic, Daniel Thalmann, and Pascal Fua (2004). “Style-Based Motion Synthesis”. In: *Computer Graphics Forum*. DOI: [10.1111/j.1467-8659.2004.00809.x](https://doi.org/10.1111/j.1467-8659.2004.00809.x).
- Usoh, Martin, Ernest Catena, Sima Arman, and Mel Slater (2000). “Using Presence Questionnaires in Reality”. In: *Presence: Teleoperators and Virtual Environments* 9.5, pp. 497–503. DOI: [10.1162/105474600566989](https://doi.org/10.1162/105474600566989).
- Vicovaro, Michele, Ludovic Hoyet, Luigi Burigana, and Carol O’Sullivan (2012). “Evaluating the Plausibility of Edited Throwing Animations”. In: *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. SCA ’12, pp. 175–182.
- Vicovaro, Michele, Ludovic Hoyet, Luigi Burigana, and Carol O’Sullivan (2014). “Perceptual Evaluation of Motion Editing for Realistic Throwing Animations”. In: *ACM Trans. Appl. Percept.* 11.2. DOI: [10.1145/2617916](https://doi.org/10.1145/2617916).
- Wallach, Helene S., Marilyn P. Safir, and Roy Samana (2010). “Personality variables and presence”. In: *Virtual Reality* 14.1, pp. 3–13. DOI: [10.1007/s10055-009-0124-3](https://doi.org/10.1007/s10055-009-0124-3).
- Waltemate, Thomas, Dominik Gall, Daniel Roth, Mario Botsch, and Marc Erich Latoschik (2018). “The Impact of Avatar Personalization and Immersion on Virtual Body Ownership, Presence, and Emotional Response”. In: *IEEE Transactions on Visualization and Computer Graphics* 24.4, pp. 1643–1652. DOI: [10.1109/TVCG.2018.2794629](https://doi.org/10.1109/TVCG.2018.2794629).

- Wang, Jing and Bobby Bodenheimer (2008). “Synthesis and Evaluation of Linear Motion Transitions”. In: *ACM Trans. Graph.* 27.1. DOI: [10.1145/1330511.1330512](https://doi.org/10.1145/1330511.1330512).
- Warren, William H. (2018). “Collective Motion in Human Crowds”. In: *Current Directions in Psychological Science* 27.4, pp. 232–240. DOI: [10.1177/0963721417746743](https://doi.org/10.1177/0963721417746743).
- Weibel, David, Bartholomäus Wissmath, and Fred W. Mast (2010). “Immersion in Mediated Environments: The Role of Personality Traits”. In: *Cyberpsychology, Behavior, and Social Networking* 13.3, pp. 251–256. DOI: [10.1089/cyber.2009.0171](https://doi.org/10.1089/cyber.2009.0171).
- Wolinski, David, Ming C. Lin, and Julien Pettré (2016). “WarpDriver: Context-Aware Probabilistic Motion Prediction for Crowd Simulation”. In: *ACM Trans. Graph.* 35.6. DOI: [10.1145/2980179.2982442](https://doi.org/10.1145/2980179.2982442).
- Xie, Zhaoming, Hung Yu Ling, Nam Hee Kim, and Michiel van de Panne (2020). “ALLSTEPS: Curriculum-driven Learning of Stepping Stone Skills”. In: *Computer Graphics Forum* 39.8, pp. 213–224. DOI: [10.1111/cgf.14115](https://doi.org/10.1111/cgf.14115).
- Yang, Dongseok, Doyeon Kim, and Sung-Hee Lee (2021). “LoBStr: Real-time Lower-body Pose Prediction from Sparse Upper-body Tracking Signals”. In: *Computer Graphics Forum* 40.2, pp. 265–275. DOI: <https://doi.org/10.1111/cgf.142631>.
- Yarbus, Alfred L (2013). *Eye movements and vision*. Springer.
- Yee, Nick and Jeremy Bailenson (2007). “The Proteus Effect: The Effect of Transformed Self-Representation on Behavior”. In: *Human Communication Research* 33.3, pp. 271–290. DOI: <https://doi.org/10.1111/j.1468-2958.2007.00299.x>.
- Yeh, Thomas Y., Glenn Reinman, Sanjay J. Patel, and Petros Faloutsos (2009). “Fool Me Twice: Exploring and Exploiting Error Tolerance in Physics-Based Animation”. In: *ACM Trans. Graph.* 29.1. DOI: [10.1145/1640443.1640448](https://doi.org/10.1145/1640443.1640448).
- Yin, Tairan, Ludovic Hoyet, Marc Christie, Marie-Paule Cani, and Julien Pettré (2022). “The One-Man-Crowd: Single User Generation of Crowd Motions Using Virtual Reality”. In: *To Appear in IEEE VR - Journal Track / IEEE Transactions on Visualization and Computer Graphics*.
- Yuan, Ye and Anthony Steed (2010). “Is the rubber hand illusion induced by immersive virtual reality?” In: *2010 IEEE Virtual Reality Conference (VR)*, pp. 95–102. DOI: [10.1109/VR.2010.5444807](https://doi.org/10.1109/VR.2010.5444807).
- Zell, Eduard, Carlos Aliaga, Adrian Jarabo, Katja Zibrek, Diego Gutierrez, Rachel McDonnell, and Mario Botsch (2015). “To Stylize or Not to Stylize? The Effect of Shape and Material Stylization on the Perception of Computer-Generated Faces”. In: *ACM Trans. Graph.* 34.6. DOI: [10.1145/2816795.2818126](https://doi.org/10.1145/2816795.2818126).
- Zibrek, Katja, Ludovic Hoyet, Kerstin Ruhland, and Rachel McDonnell (2013). “Evaluating the Effect of Emotion on Gender Recognition in Virtual Humans”. In: *Proceedings of the ACM Symposium on Applied Perception*. SAP ’13, pp. 45–49. DOI: [10.1145/2492494.2492510](https://doi.org/10.1145/2492494.2492510).
- Zibrek, Katja, Ludovic Hoyet, Kerstin Ruhland, and Rachel McDonnell (2015). “Exploring the Effect of Motion Type and Emotions on the Perception of Gender in Virtual Humans”. In: *ACM Trans. Appl. Percept.* 12.3. DOI: [10.1145/2767130](https://doi.org/10.1145/2767130).
- Zibrek, Katja, Benjamin Niay, Anne-Hélène Olivier, Ludovic Hoyet, Julien Pettre, and Rachel McDonnell (2020). “The Effect of Gender and Attractiveness of Motion on Proximity in Virtual Reality”. In: *ACM Trans. Appl. Percept.* 17.4. DOI: [10.1145/3419985](https://doi.org/10.1145/3419985).

References

- Zordan, Victor Brian, Anna Majkowska, Bill Chiu, and Matthew Fast (2005). “Dynamic Response for Motion Capture Animation”. In: *ACM Trans. Graph.* 24.3, pp. 697–701. DOI: [10.1145/1073204.1073249](https://doi.org/10.1145/1073204.1073249).