Animation of Quadrupedal Animals and Perceptual Evaluation of their Gaits

by

Ljiljana Skrba, B.A. (Mod)

Dissertation

Presented to the

University of Dublin, Trinity College

in fulfillment

of the requirements

for the Degree of

Doctor of Philosophy

University of Dublin, Trinity College

March 2010

Declaration

I, the undersigned, declare that this work has not previously been submitted as an exercise for a degree at this, or any other University, and that unless otherwise stated, is my own work.

Ljiljana Skrba

Ljiljana Skrba

March 3, 2010

Permission to Lend and/or Copy

I, the undersigned, agree that Trinity College Library may lend or copy this thesis upon request.

Acknowledgments

Firstly, I would like to thank Carol O'Sullivan for giving me the opportunity to work on this project. Secondly, I'd like to thank all the people who have helped me along the way, in particular Lionel Reveret, Franck Hetroy, Marie-Paule Cani from EVASION group at Inria for helping me with the State of the Art Report, and Alan Baird and Patsy Kearns from School of Agriculture, Food Science & Veterinary Medicine at University College Dublin which helped me greatly with the gaining access to and the filming of animals.

LJILJANA SKRBA

University of Dublin, Trinity College
March 2010

Animation of Quadrupedal Animals and Perceptual Evaluation of their Gaits

Publication No. ____________________

Ljiljana Skrba, Ph.D.
University of Dublin, Trinity College, 2010

Supervisor: Carol O'Sullivan

Abstract

Computer generated animals have become a common feature of today's digitised society, often found in animated films. In computer games highly realistic animals are simulated in real time.

High quality characters and their animations are usually generated by hand by skilled artists, but there has also been automatic generation of animation for multi-legged creatures in games like Spore^{TM}(Electronic Arts). In both cases the process is very challenging - in hand animation the timing and movement of the motion is difficult to represent correctly, and in physically based simulations it is equally difficult to set up the correct functions that deal with timing of the motion, and the physical rules associated with a character's motion.

In some cases it is possible to reuse a motion of one character and apply it to that of another. For example in Evan Almighty (Universal Pictures) different motions of one monkey were recorded and applied to virtual monkeys of different species in the film. This process makes the animation task easier as it allows for the reuse of motions that have been created or motion-captured.

The main focus of this work has been on realistic quadruped animation driven by knowledge of human perception of animal gaits. Towards this end we have built various tools and tested different methods for the creation, capture and replication of animal motion. We propose a study of how sensitive humans are to different gaits of quadrupeds, and whether it is possible to use the captured motion of one and apply it to a model of another. We also investigate which animations look best for animating flocks and herds of animals.

To study quadrupedal gaits we captured the walk and trot gaits of a horse, cow, sheep and pig and designed perceptual experiments that allowed us to test how sensitive humans

are to the different motions. We made point-light representations of all the motions captured, and we were also able to re-target these motions to 3D models to find out whether viewers could recognise the different gaits.

We found that humans can recognise the pig motions from point-light representations and when they are disguised on a 3D sheep model. Similarly, people are also able to recognise the horse trot gait from a point-light and when it is applied to different 3D models. This suggests that these motions are distinctive and we would not recommend their use for animation of other animals.

We also found that the horse walk and cow walk motions are not distinguishable (both from point-lights and 3D model animations) which would suggest either looks genuine when applied to a horse and cow model. This is a welcome result because it means less work is required from an animator if they can use the same motion to animate a few different models without it becoming apparent.

For herds, our perceptual experiments show that viewers were more sensitive to slower walks and overall they preferred herds and flocks that were animated using a faster motion. In addition, we found that animating a flock or herd using one gait that is out of step looks as effective as animating a flock or herd using two different gaits. Once again, this is a welcome result as it means less resources are used for simulating flocks and herds.

Animals are a popular topic in the animation field but they are difficult to represent. For this reason our work is an important contribution to the field as it can be used as a guideline for the re-use of motion between different 3D models. This applies to both singular movement and to the movement of flocks or herds.

Related Publications and Outputs:

1. "Animating Dolly: Real-time Herding and Rendering of Sheep", Ljiljana Skrba, Simon Dobbyn, Rachel McDonnell and Carol O'Sullivan. Eurographics Ireland Workshop, pages 7-12, 2006.

2. "Quadruped Animation", Ljiljana Skrba, Lionel Reveret, Franck Hetroy, Marie-Paule Cani and Carol O'Sullivan. In Eurographics *State-of-the-Art Report (EG-STAR)*, pages 1-17. Eurographics Association, Eurographics Association, 2008.

3. "Waltzing Sheep", Ken Perlin and Ljiljana Skrba, A short film at Eurographics Electronic Theatre 2007.

4. "Eye-tracking dynamic scenes with humans and animals", Ljiljana Skrba, Ian O'Connell, and Carol O'Sullivan. In APGV '08: Poster Proceedings of the 5th symposium on Applied perception in graphics and visualization, page 199, New York, NY, USA, 2008. ACM.

5. "Human perception of quadruped motion", Ljiljana Skrba and Carol O'Sullivan. In APGV '09: Poster Proceedings of the 6th Symposium on Applied Perception in Graphics and Visualization, page 130, New York, NY, USA, 2009. ACM.

6. "Join the Dots: Insights into motion of quadrupeds", Ljiljana Skrba and Carol O'Sullivan. Eurographics Ireland Workshop, pages 81-88, 2009.

7. Submitted a paper entitled "Perception and Simulation of Quadruped Motion" to CASA 2010.

Contents

List of Tables

List of Figures

"Never work with animals or children",

W.C. Fields

Chapter 1

Introduction

Since ancient times people have been depicting animals using different media. The Chauvet caves in France (Fig. 1.1) contain images of animals such as horses, rhinoceroses and deer that are up to 30,000 years old [VCG$^+$01]. With the advent of moving pictures, large New York and Hollywood studios started producing many cartoons depicting animals, often giving them very human characteristics, physical as well as behavioural e.g., Warner Bros.' Looney Tunes or Metro-Goldwyn-Mayer's Tom and Jerry. In the 1980s and 1990s Disney was a driving force in America and Europe in terms of animal animation, continually improving the quality of animal representation e.g. in the Lion King [Cra99].

Today CGI animation given the incentive for very realistic depictions animals, in cartoons as well as films and other media. We are now able to depict animals in more complex and refined ways than ever before. For example, in films like The Chronicles of Narnia (Disney), The Golden Compass (New Line Cinema), I Am Legend (Warner Bros. Pictures) and The Twilight Saga: New Moon (Summit Entertainment), some animals look and behave very

Figure 1.1: A painting of a rhinoceros found in Chauvet cave in France [VCG$^+$01], CNP Ministere de la culture.

Figure 1.2: A varied herd of cows used in perception experiments.

realistically - albeit some of the animals do talk. To achieve such realism, large teams of animators, artists and directors are involved, and even then it is a monumental task. Realistic representations of virtual animals are also used in documentaries and in veterinary sciences where the gaits of animals are analysed [Fro09, Cap05, Gil02].

The technical challenges associated with creating highly realistic, computer generated creatures have been receiving increasing attention recently. The entertainment, education and medical industries all contribute to the growing demand for simulation of realistic animals in the computer graphics area. In order to achieve this, several challenges need to be overcome: gathering and processing data that embodies the natural motion of an animal – which is made more difficult by the fact that most animals cannot be easily motion-captured; building accurate kinematic models for animals, with adapted animation skeletons in particular; and developing either kinematic or physically-based animation methods, either by embedding some a priori knowledge about the way that quadrupeds locomote and/or adapting examples of real motion.

Furthermore, the creation of realistic crowds, flocks or herds can be computationally demanding for a regular desktop PC and it takes render farms and large teams of artists and animators to create films such as Ratatouille (Pixar Animation Studios), Happy Feet (Kingdom Feature Productions) and Australia (20th Century-Fox Film Corporation) to name but a few. We ask the questions: is it possible to use the gait of one quadruped animal and realistically animate a 3D model of another? How sensitive are people to the different animal gaits? Re-targeting of motion for the animation of different 3D animals has already been used by Rhythm and Hues in films such as Evan Almighty (Universal Pictures), [Rij09].

In this thesis, we investigate different techniques for the animation of quadrupeds. Using re-targeting we have examined how sensitive humans are to the gaits of some animals; e.g., horse, cow, sheep and pig. We have also investigated whether the gait of one animal can be used to animate a 3D model of another. Furthermore we are aware that variation is key to illustrating realistic herds of animals, and in this work we have included a study that involves the animation of herds of different animals, Fig. 1.2. We explored whether different gaits within a herd produce better results than using a single gait. The results of our work can be used as a guidance for the creation of realistic herds in different applications where resources are limited and real-time simulation is required.

1.1 Motivation

To date there are almost 200 films that feature animals with human-like characteristics [Wik], for example the lion in Narnia (Disney), or daemons and bears in The Golden Compass (New Line Cinema). In addition, there are films in which the animals do not have anthropomorphic characteristics. In these types of films CG animals are used in scenes that are too dangerous for live animals or in places where huge herds are necessary, for example horses in Troy (Warner Bros. Pictures), sheep in Brokeback Mountain (Alberta Film Entertainment), deer, lions and dogs in I Am Legend (Warner Bros. Pictures) or cattle in Australia (20^{th} Century-Fox Film Corporation).

Reproducing realistic animal motion is very complex and only at the end of the 19^{th} century have actual animal movements been captured on film. Photographs do not convey the correct speed of a motion and so it is easy to get the timing of different gaits wrong. This can result in motions which may look as if they are either too rigid, too slow or too fast. In addition, photographs only contain 2D information which may not be enough for the recreation of certain poses such as in asymmetric gaits (the gallop), where the visual information about the legs is missing on the hidden side of the body. Similar drawbacks apply to video capture in a natural environment, where it is only possible to capture the movement from one side of the animal and occlusion of body limbs is very common, while hooves and paws are often hidden by grass or uneven terrain, Fig. 1.3. The quality of videos and photographs can often be poor where the limbs are blurred or difficult to differentiate from the background. In order to capture full 3D motion, motion capture systems can be used but these are very expensive and the equipment available is not always easy to adapt for use with animals. Large areas and trained animals are required, especially if one wishes

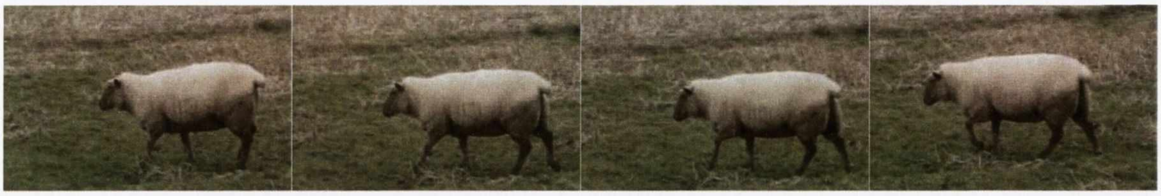

Figure 1.3: Four frames of a video show a sheep moving from right to left. The feet are hidden by the grass and the sheep is walking on uneven terrain. Due to the low quality of the video it is difficult to track any movement of the upper body.

to capture quick or turning movements.

There are many ways to animate characters using computers, but the process still retains some of the same problems as before when artists created animations exclusively by hand. Key-frame animation is still challenging and relies greatly on the skill of the animator considering that slight changes to a movement can take hours to perfect. Using procedural techniques can also cause problems, as an animator does not have complete control over a character. This means that only certain limbs can be moved and only in certain directions. Fully automatic techniques do exist but once again these take time to set up, for example a user may need to specify the step size, limb size, frictions, amongst other parameters, for each character.

To simplify the process of animation it would be useful to minimise the amount of animation required per character. Figuring out which areas of a body people look at most may allow animators to concentrate their efforts there. There is not much use in perfecting the motion of the legs if the upper body remains stiff. This can detract from the realism of an animation as it is immediately spotted. To makes matters more complicated, some animals are found in groups rather than on their own (naturally, animals stay close together to reduce chances of being killed by a predator), so realistic herds or flocks of animals are needed. People are very sensitive to the different ways in which humans walk, and it is very easy to spot when there are "clones" in a crowd, but are we as sensitive to the movement of animals? Can we tell them apart easily? Knowing the answers to these questions could inform us whether we can simplify and thereby speed up the process of animal animation. Re-using motions saves a great deal of preparation effort, while finding a way to modify a generic motion makes it easier to animate a range of animals. This also could allow for variation to gaits within a group of similar animals.

Capturing the motion of different animals is a difficult task, especially when those animals have not been trained. Would it be possible to capture the motion of one species of animal and apply it to another, so long as their over all physique is similar? For example, can we capture the motion of a horse (a relatively tame and easily trained animal) and apply that motion to a deer/sheep (an animal that is difficult to control and film)? In this work we address the issue of motion re-targeting from one animal to another and examine how well people recognize different quadrupeds under these conditions.

1.2 Contributions

The main focus of this work has been on realistic quadruped animation driven by knowledge of human perception of animal gaits. Towards this end we have built various tools and tested different methods for the creation, capture and replication of animal motion. Below is an outline of the contributions of our work,

- A published STAR at Eurographics 2008 - In collaboration with members of the EVASION group at Inria we published a comprehensive STAR, bringing together many different aspects of Quadruped Animation for the first time.
- A publication of a paper at Irish Chapter of Eurographics 2006 - "Animating Dolly" discusses the techniques we used to create sheep and dog impostors that were added to a real time system where shepherding and flocking behaviours were implemented.
- An IK system - To investigate whether it is possible to animate different quadrupeds using a single method we developed a simple automated system that animates quadrupeds using IK (Inverse Kinematics). A user is able to change the size of the legs or else pick from a set of pre-defined shapes representing a horse, cow, sheep and pig.
- Capturing the motion of animals - While many sophisticated 3D motion capture systems are available on the market today, using these on animals remains a difficult challenge in most practical situations. We have implemented a methodology for the motion capture of animals that is a relatively simple and cheap technique, where we paint the primary joints on the animals we wish to capture and then we film them using a video camera. A simple tool allows a user to track the motion of all the points marked on the animal. This is stored in a text file which can be easily used to visualise the motion in other software packages.
- Re-targeting captured motion to different 3D animal models - We were able to use the motion of the point-light walkers to animate the 3D models. As each point-light defines the position of a joint, it is possible to use this information to place the joints of the skeleton of the 3D model into the corresponding positions. Using scaling we can apply

the motion of a horse to a 3D model of a sheep or cow while still retaining the original movement.

- Perception Experiments (single animals) - To gain insights about whether we are sensitive to motions of different animals we designed experiments where we used genuine animal motion to animate point-light walkers and 3D models. From our novel experiments we have found that a horse's trot is distinct and can be spotted easily on point-light walkers and on different 3D models (horse, cow and sheep). Similarly the pig motion is easily recognised on both the point-light and 3D model. Horse and cow walk gaits are difficult to distinguish, suggesting that for this particular gait it is possible to re-target the motion of one animal to the model of another.

- Perception Experiments (herds of animals) - We created herds of 3D animals (horse, cow and sheep) using the captured motions of the horse and cow to see which animations people found more aesthetically pleasing. Once again the horse trot came up as distinct as it was the only motion that was preferred on a horse model. Overall the cow trot was the preferred gait for the cow and sheep models. We added variation to the herds by animating half the herd with a horse gait and the other half with a cow gait. We found out that this variation does not have any advantage over using a single gait - a favourable result as it is easier to implement.

- Real-time simulation - It is possible to use our animated models in real-time systems such as games or other interactive environments. The results from our perception experiments give insights about what animations to use for simulations of herds of animals.

Various games contain different animal characters, for example, Lucinda Green's Equestrian Challenge©(Red Mile Entertainment), Pet Vet 3D Animal Hospital©(Take Two Interactive), Paws and Claws Pet Vet (ValuSoft), Zoo Tycoon®(Microsoft Game Studios), Afrika™(Natsume) and Fable 2©(Microsoft Game Studios), to name but a few. Reducing the amount of animations and textures required allows interactive applications to run smoothly and quickly while still providing a realistic environment. Likewise, in the film industry it is useful to know which animations can be used to animate herds of animals and whether one needs to add much variation to the movement of animals. Our techniques can also be used to further develop visualisation tools that can be used in veterinary science to teach students about animal motion and possibly for better diagnosis of gait abnormalities - Structure and Motion Laboratory at The Royal Veterinary College have captured the motions of an elephant to learn how to spot injuries early [Fro09].

Our work involving the perceptual aspect of animal motion is novel within the computer graphics field and indeed, there is very little research in the perception field also. Therefore this work produces novel and interesting results for both fields and will hopefully stimulate many ideas for future work.

1.3 Scope and limitations

Our primary aim was to discover how sensitive humans are to the motions of different animals and if it is possible to use the captured motion of one to animate the model of another. As discussed, animals feature in many films in the industry but they are usually difficult to access and control.

We have applied the captured motion of farm animals to different types of models and have tested how well people recognised the different motions. We concentrated on animating models of farm animals (horse, cow, sheep and pig) as these animals were available for tracking through our collaboration with the School of Agriculture, Food Science & Veterinary Medicine at University College Dublin (UCD). Another reason for using these animals is that they have a similar physical build (they are all ungulates) and so we wanted to see whether the motion of these animals is similar even though they clearly look different externally. It is interesting to see whether the appearance has an effect on the perception of the motion, as if this is true then it suggests that the appearance of an animal is more important than its motion. We chose not to compare the animals against other species such as dogs and cats. Dogs and cats are digitigrade animals, which means that they step onto the ground with their whole paw as opposed to a hoof, a significant difference that has a likelihood of being easily noticed. Another reason for avoiding dogs is that they vary greatly in physique between different species, and so it would be difficult to select one type to capture for perception studies as the results could be specific to only one breed.

To test the human perception of animal motion we needed animations that are as lifelike as possible. We tried many different techniques using photographs and videos as guidelines, but hand animation never produced results that were of high enough quality, as it was not "real" motion and also we did not have access to expert animators during the process. Similarly, procedural animation resulted in very stiff unrealistic motions that could not be used in the experiments. For this reason we used a technique where we marked joints on animals in order to capture their motion from video, which we then applied to the 3D models. We captured two different gaits for each animal to see if the speed of the motion had any effect on our perception - can we recognise a horse when it trots but not when it walks? As we were comparing between gaits we did not make any attempt to blend between any of the motions nor to create any transitions from a walk to a trot motion.

To achieve our goal we built a simple tool which allowed us to track the motion of each animal by recording the position of each marker for every frame of the video. We were able to create simple point-light walkers from this, which we used in our initial perception experiments. We then built other tools that allowed us to apply the motion of the point-lights to the 3D models. We only captured 2D motion of the animals, which was adequate for our work as we mainly wished to animate animals walking or trotting from the side. The two gaits captured are symmetrical, so even though we were missing the motion for one side of the body we could mirror the motion of the visible side. To track the motion of animals we did not build any system that would be able to extract the motion automatically, as this is a computer vision problem and beyond the scope of this project. Our original plan was to use existing commercial software to extract the motions of animals but it was unable to extract the marked joints from the videos. In addition, as we are not investigating the behaviour of these animals, we did not implement multiple gaits or turning motions for each animal as they wander in different directions or transition between gaits. However, it would be interesting to investigate full 3D motion for a variety of movements and transitions in future work. Currently all our tools require some input from the user and are essentially there to aid the user and minimise any hand-animation. However, in the future it would be useful to build a system that is able to automatically apply or generate realistic quadrupedal motion for 3D models.

1.4 Methodology

We used 3ds Max® to rig and animate 3D models of animals using key-framed hand animation and scripting. It helped us create furry impostors for use in a real-time system, point-light walkers and herds and flocks of animated 3D models.

We used C++ to simulate flocking and shepherding behaviours in the existing real-time system. To investigate the aspects of automatic generation of quadruped motion we used OpenGl to build a simple IK system and we used glui libraries to provide an interface that helped a user select which quadruped animal they wanted animated.

To create the point-light walkers we had to extract the motion of the real animals from videos. We used OpenCV to create an interface which allows users to mark and track the animals in the videos.

For each experiment we got ethical approval from our university to ensure we could carry them out, and for some experiments volunteers were given a voucher for the national bookstore. To run the experiments we used Presentation® which also stored the results for each participant. In order to analyse the results of all the experiments we used Statistica 7© software.

1.5 Summary of Chapters

Background research is summarised in Chapter 2. This includes a variety of animation methods such as hand animation, procedural animation (including fully automatic simulations of quadruped animation) and mesh animation. We discuss works that try to simulate realistic behaviour of animals, such as packs of wolves and flocking algorithms, and we give an overview of previous perceptual research in this and related domains.

To better understand the motion of animals, in Chapter 3 we include a description of the anatomy of quadrupeds. We also compare the skeletons to those of humans, as this provides us with an intuitive parallel through which we can better understand the biomechanics of animal movements.

In Chapter 4 we present an overview of the many techniques used for character animation. In particular we focus on the methods we have used throughout this project and we discuss the advantages and disadvantages of the various approaches.

In Chapter 5 we describe our pilot and main perception experiments and analyse the results. These experiments made use of eye-tracking, point-light walkers, and 3D models onto which the motion of different animals had been re-targeted.

Chapter 6 considers how to create varied herds of animals. Here we also discuss the results from our perception experiment involving herds of animals with appearance and gait variation.

Finally we conclude with Chapter 7, discussing the results and their implications. We also suggest possible directions for future work.

Chapter 2

Background and Related Work

In our research we have addressed the problems of quadruped motion for simulating realistic flocking and shepherding animations. Burt presents the history of animals in films and the technical challenges overcome in the film industry throughout the twentieth century [Bur03]. In this chapter we will give an overview of the techniques that have been used to date for producing quadruped motion including video-based acquisition, physics based models, inverse kinematics, or some combination of the above. We also give a short summary of the related work that deals with flocking and shepherding methods. First, we provide a short introduction to quadruped motion where we give a brief historical overview of research on the motion of animals in different fields. This is followed by the different methods used to gather the data needed to create data driven quadruped animation. We then describe the animation of characters with skeletons using inverse kinematics, and without skeletons using mesh deformation. We also cover the work that has been done on physical simulations of animals, the creation of anatomically correct creatures, controlling the behavior of animals through user interfaces, and also simulations of group behaviours. Finally, we finish with a short case study of "The Chronicles of Narnia". Here we look at the techniques and tools used for the creation and animation of different quadruped characters that feature in the film.

2.1 Introduction to quadruped animation

Film animals, such as the lion Aslan in Disney's The Chronicles of Narnia: The Lion, the Witch and the Wardrobe, look and behave as convincingly as their real counterparts [HDK+06]. They have changed people's expectations of how realistic the animations of furry creatures should appear. In addition, there has been some recent commercial work on realis-

Figure 2.1: Aslan from "The Lion, the Witch and the Wardrobe". Image © Copyright Disney Enterprises Inc. and Walden Media Llc. All rights reserved.

tic simulation of large herds of animals as can be seen in games such as Afrika®(Natsume). The game features replications of much of Africa's flora and fauna, where large herds of wildebeest and zebras can be displayed in real-time. Over the years, much research has been concentrated on the simulation of realistic humans and there are many courses and surveys relating to this topic (e.g., [TT04, BK05, WBK+07]). However, as can be seen from the examples above, others have also been active in developing realistic simulation methods for a variety of non-human characters, particularly for quadrupeds.

With respect to traditional animation techniques, standard authoring tools are now available for animating quadrupeds. However, as with all animations that are created using key-frame animation, the realism of the final motion depends on the knowledge and skill of the animator. Procedural quadruped animation allows the automation of this task, but has the drawback that the animator has less direct control over the motions, which can result in a loss of realism [vdP96]. Difficulties arise when trying to animate complex articulated structures, with the compounding factor that humans are very familiar with such motions and can detect anomalies quite easily. Motion capture can provide large amounts of detailed data that can be used to replicate the motion of a character's performance with a high level of fidelity. Such data is hard to come by for animals due to a number of drawbacks. Sturman provides a short history of the use of motion capture for computer character animation in [Stu94]. While animal motion capture may not always be feasible, there have been some recent successes with using video-based methods to extract motion data. However, it should

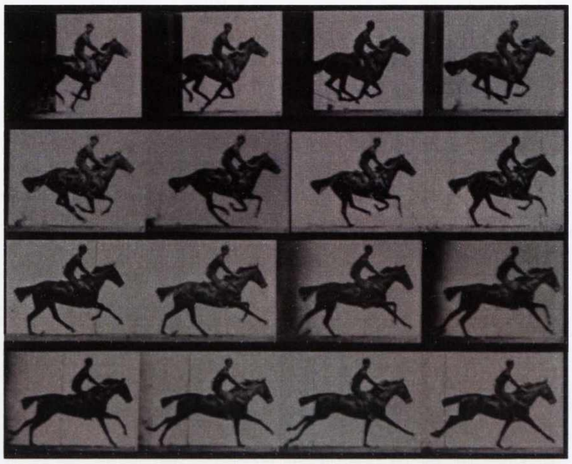

Figure 2.2: Sequence of pictures capturing the motion of horse and rider (Muybridge [MB57]).

be noted that motion capture can only model one animal's existing motion; the user will need to edit it, parameterise it and adapt it to their purposes; very often, they will simply want to learn from or be inspired by it, then apply what they have learned to other animals with varying morphologies, perhaps locomoting on another terrain, with a different trajectory and combination of successive gaits.

2.1.1 Optical recording of quadruped motion

The biomechanics and animal physiology literature provides several sources of data upon which the synthesis of realistic motion for quadrupeds [MB57, EKB95] is based. In the late 1800s, the works on motion photography of Eadweard Muybridge in the US and Etienne-Jules Marey in France greatly improved the understanding of animal gaits. Using a trigger, Eadweard Muybridge was able to set off a series of twenty-four cameras in order to capture galloping horses, deer, buffaloes, camels, dogs, cats and many more animals. Over 3,919 such photographs are shown in [MB57]. Etienne-Jules Marey used a different technique called chronophotography, where several phases of a single moment could be recorded on the same image. Images of animals in motion can show, for example, that there is a suspension phase during a horse gallop where all the hooves are above the ground but under the body, as can be seen in the top row of Fig. 2.2. Still images such as these are still valuable sources of motion even today, especially in situations where hand animation is used for animated feature films, video games and entertainment applications.

Rotoscoping is a technique which was developed in the early 1900s and is still widely used today in animation and for special effects. The basic process consists of a video projected

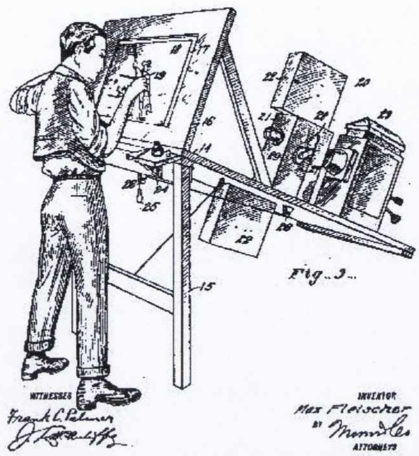

Figure 2.3: Rotoscoping by Max Fleischer (Image appears in [Cra82]).

one frame at a time onto a glass screen. A sheet of translucent paper is placed on top of the glass and the image of the video frame is traced onto the sheet, see Fig. 2.3. A different sheet is used for each frame. The images are then re-photographed onto the cinematic film using a rotoscope machine [Mil98]. This results in character motion that is very smooth and life-like, which is very difficult to achieve by hand animation alone. The results of this technique can be readily found in many Disney feature animations, such as Snow White, Fantasia, and Mary Poppins. The most relevant films to us are Lady and the Tramp, Bambi, The Jungle Book, and The Lion King where frequent trips to nearby farms and zoos took place in order to film the animals [TJ95]. Similarly, in the film 101 Dalmatians, rotoscoping was mainly used for animating human and animal motion.

Today the same principles apply, except that computers are used to digitally scan the film onto disk. Then, specialised software is used to apply stylisations to the video (e.g., Shake, FFI and Pinnacle Commotion) [Sil04]. Rotoscoping can even be thought of as an alternative to motion capture. The motion of the actors is captured by hand and, as already mentioned, is often used to capture and recreate the motions of four-legged creatures. However, this is a very tedious and time-consuming process, as the animator must work on one frame at a time. To overcome this problem, computer vision solutions have been found for tracking contours and extracting animal motion from video footage in an efficient way [AHSS04, FRDC04]. This is discussed in detail in Section 2.2.2.

2.1.2 Biomechanics of animal motion

Quadruped motion has long been an active area of research in biomechanics and zoology, which can provide valuable insights into generating automatic gaits for various quadruped characters. Alexander has published numerous reference works on animal motion [Ale68, AJ83, Ale84, Ale96, Ale03]. In particular, he developed a hypothesis about the relationship between size, speed, mass and external forces. The resulting *dynamic similarity hypothesis* states that the movements of two bodies can be described as dynamically similar if the motion of one can be made identical to that of the other by multiplying a) all linear dimensions by some constant factor, b) all time intervals by another constant factor, and c) all forces by a third factor. Under the assumption that the weight is the dominant external force involved in locomotion, the dynamic similarity hypothesis shows that motions dynamically equivalent have a constant quantity v^2/gh, where v is speed, h is the height of the hip from the ground in normal standing position and g is the gravitational acceleration. This quantity is known as a *Froude number*.

The consequence of this hypothesis is that any animal, whatever his size and weight, tends to change gaits at equal Froude numbers. This rule predicts that species change from a symmetric gait to an asymmetric gait when their Froude number is between 2 and 3. This has been verified experimentally by Alexander and Jayes on wide range of animals, from rodents to rhinoceros [AJ83]. It can be used as a guide for animating models of different shapes and sizes by providing guidelines on how fast an animal should be moving depending on its size and what gait should be used. Keeping the relative gait speeds consistent in an environment of varying animals should serve to increase the realism of the simulation.

2.1.3 Animal robotics

Physical reconstructions can also provide great insights into the motion of quadrupeds for such things as balancing, managing difficult terrain, changing gaits and steering. Examples of such robots are "BigDog" and "Little Dog" built by Boston Dynamics [Bos]. BigDog's legs are shaped like that of an animal which can absorb shocks and recycle energy as it moves from one step to the next. Similarly LittleDog is used for the study of locomotion. A more familiar example is that of the Sony AIBO robot. Shaped like a dog, the robot can be programmed and used for vision studies (so that the robot can "see") and gait studies (so the robot can move around while keeping balance) [GLH04]. However, all of these machines still move like robots. BigDog's legs are synchronised differently to the more common quadruped gaits (the diagonal legs move at the same time), whereas LittleDog on the other hand has the typical quadruped motion (for example: front left, back right, front right, back left) but the motion is still very jerky and slow as it tackles rough terrain.

2.2 Data driven approaches

One approach to generating quadruped motion is to capture and apply the actual motions of real animals. This may be as simple as visually observing their motion, or as complex as using multiple cameras for full-body motion capture. However, reliable and realistic results can be difficult to achieve.

2.2.1 Standard motion capture

With respect to the motion capture of animal motion, the first challenge is in actually attaching equipment or markers to animals in order to track their movement. Afterwards, the performance of an action must be restricted to the space covered by the cameras or magneto fields. In some cases treadmills can be used, for example with horses and dogs. However, this method can produce uncharacteristic motion - walking on grass is different to walking on a treadmill, and it is unsuitable when working with wild animals. Furthermore, any data captured will be very specific to a particular animal and its movement, making it difficult to change or blend with other motions.

In animal biomechanical studies, force plates are often used to track the force exerted by particular limbs in animals, thus determining which are the weight bearing legs. For example, in the case of dogs, the front legs act as the main support for body weight (ribs, head and other parts all rest on the front legs) while the back legs are used to apply the pushing force [CDNDMN85]. However, this method does not provide information about the movement of the spine, neck and head. Due to an increase in demand for computer generated animals in the film industry, motion capture technology has been used to capture and analyse the gait of larger animals such as horses. However, this is expensive and done by specialised companies such as [Law] and [Wid]. Motion capture is also used in veterinary sciences for measuring the kinematics of horses' joints in order to better understand their movement and the origin and causes of dysfunctions which occur among horses used in competitive sports [BPB06].

De Aguiar et al. present a markerless approach to capturing human performances in [dATSS07] and [dAST+08]. Currently applying motion captured data directly to a 3D model can sometimes look unrealistic. This is because it is very difficult to get a direct correspondence between a 3D model and the captured human actor. Generally, the motion captured is that of a person in a tight fitting outfit and this might not always be suited for a 3D character that is wearing loose clothing. Using a 3D laser scanner, de Aguiar et al. scan an actor while they are fully dressed. Afterwards, the motion is recorded using a number of cameras to capture the motion from different angles. With their technique, they can deform the 3D model according to the data captured on video. The silhouette of a character in the video is compared to the silhouette of the model and this is refined until there is a close match. This technique could also possibly be adapted for use with animals, although the authors do mention that capturing the motion of hair or fur is not yet possible.

Equine locomotion is a very active area of research, with dedicated conferences and journals [Cla97, Cla98, Cla01, ICE], so it is no surprise that motion capture methods have been applied to horses more than any other animals. It also helps that they are tame and easily-trained beasts. However, this captured motion is very specific to horses and thus not applicable to other animals. The difficulties with reuse and retargetting of such data, together with the cost and effort involved in its capture, means that motion capture for animals currently remains largely within the remit of specialised companies.

2.2.2 Capture from video

There are obvious difficulties that arise when working with wild animals such as elephants, cheetahs, tigers and lions. Traditional motion capture is clearly unsuitable, so video processing is the most practical solution. While there are numerous wildlife documentary videos featuring such animals, the main obstacle to exploiting this footage is the single viewpoint, making it impossible for a standard 3D measurement of motion. To overcome these problems, two types of approaches have been taken: standard tracking (e.g., [AHSS04]) and statistical analysis (e.g., [FRDC04]).

Wilhelms and Van Gelder [WG03] use a technique to extract the motion of a horse from video, which is then applied to their three dimensional models of horses. The video image is processed using active contours. The model is scaled to match the animal in the video and it is aligned with the video. The active contours of the video are then anchored to the model of the horse. Playing the video changes the shape of the contour lines which in turn change the positions of the horse's limbs, as can be seen in Fig. 2.4. This method is very sensitive to noise so the user has to reinitialise active contours every few frames.

Following the same approach but with a stronger focus on the user interface, Agarwala et al. [AHSS04] introduce "roto-curves" (similar to active contours) to outline areas of interest, which are specified by a user for the first and last frames of a video. For all in-between frames the curves can be calculated automatically and even be corrected by the user at any point in time, such as in cases of complex motion. For example, a person walking and waving their hand can be outlined with roto-curves. The results can be used for both special effects and animation. Different effects, like stylistic filters, can be applied to areas outlined by the roto-curves. To create cartoon style animation, user-drawn strokes can be attached to the

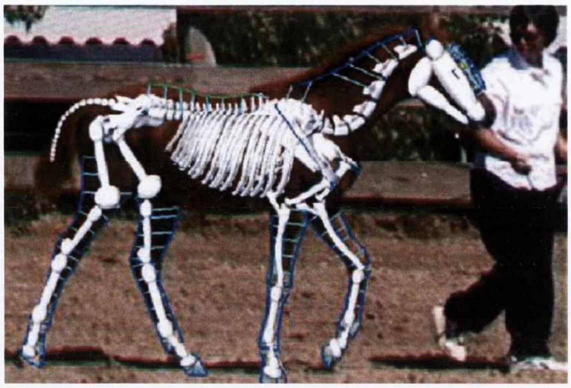

Figure 2.4: Horse model is the same size as the one in the video. The active contours are anchored to the model so it follows the horse in the video (Wilhelms and VanGelder [WG03]).

roto-curves and are then pulled into place according to the shape of those curves, see Fig. 2.5.

Gibson et al. [GODC05] present a system that captures the motions of very small scale creatures, such as spiders and ants. Three cameras are used to film the movements and then tracking techniques capture specific points on the animals. While capturing the animals, their size, the speed of their movement and the lighting in the room (diffuse lighting reduces harsh shadows) must all be taken into account. Spiders are slow moving, so cameras set at 25fps are used, whereas ants are very fast and require a frame rate increase to 150fps. For the points found in the primary view, corresponding points are found in the other two camera views, thus creating a 3D point. In order for this to work even when there is occlusion, 2D motion histograms are used to determine whether 3D points from corresponding pairs predict the expected 2D frame to frame motion. Due to the size of the creatures, very little video footage is used as input data for characteristic motion. However, the use of motion synthesis algorithms allows for the generation of large amounts of new realistic motions. This work has been used for a Natural World production by BBC's Natural History Unit.

Statistical analysis is used by Favreau et al. [FRDC04] to analyse video data and make it applicable to the generation of 3D motion. A live video sequence of a wildlife documentary is segmented into binary images in order to isolate the foreground subject from the background. This can be done automatically by image segmentation or from a rough user sketch of the parts of the animal when the automatic segmentation fails. The input visual data do not have to include detailed joints, just the main features. Principal Component Analysis (PCA) is applied directly on all the segmented images of the video sequence, each image being considered as a single high dimensional observation vector. Using PCA directly on images was first proposed by Turk and Pentland [TP91] for facial images recognition. In [FRDC04]

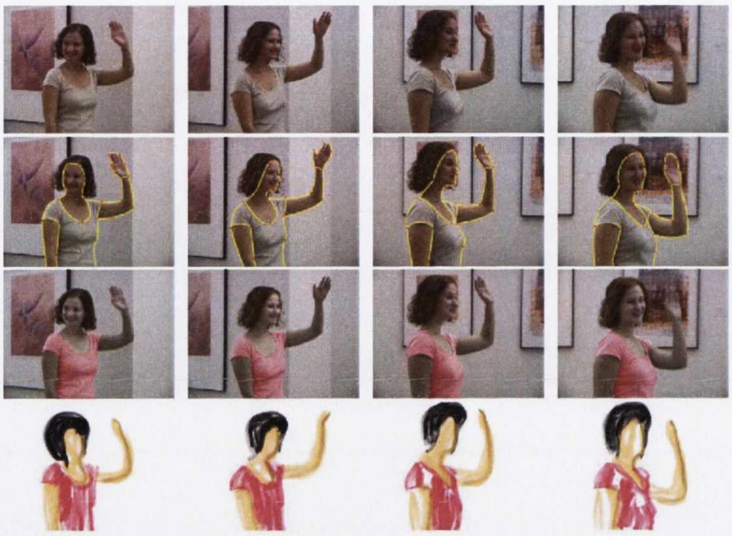

Figure 2.5: Roto-curves are used to outline areas of interest (Agarwala et al. [AHSS04]).

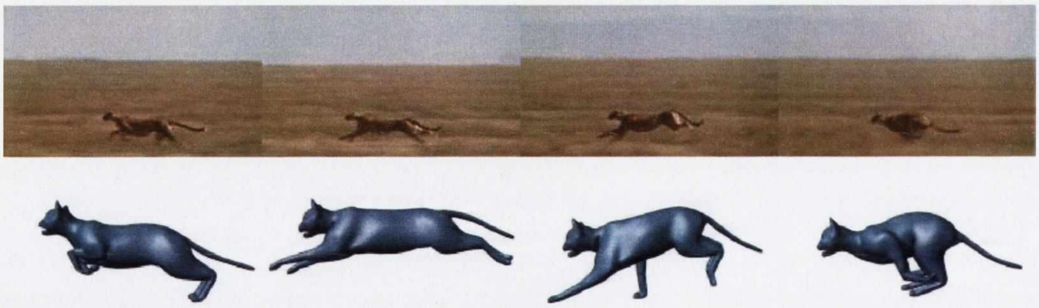

Figure 2.6: Images extracted from a video are used as a learning set for prediction of continuous 3D motion (Favreau et al. [FRDC04]).

it is used to find regular motion patterns in the images and serves as an input parameter to continuously and robustly predict 3D motion of animals using Radial Basis Functions (RBF). In cases where the profile silhouettes of an animal are ambiguous, as can be the case in symmetric walks, a user is prompted to validate the silhouette as a different pose. The final collection of images is used to provide information for controlling 3D animation, as seen in Fig. 2.6.

Gibson et al. [GCT03] have also applied PCA on visual cues from video footage of animals in motion. Gibson et al. work with dynamic outdoor scenes of walking animals. Videos of wildlife film footage can be large and contain a vast range of content (anything from large sky regions to complex foliage). They focus on recognising animal movement within such videos. The background image tends to move in a consistent manner and is extracted, leaving the foreground as the point of interest. This method works even when the camera is moving, assuming it moves in a relatively consistent manner. An eigengait model enables spatio-temporal localisation of walking animals. The eigengait space is generated using a video of a horse walking on a treadmill. Each frame is hand labelled with 12 points which are tracked over the video sequence. Using PCA analysis on all points found in the video sequence, an eigengait space is generated. Gait frequency in the outdoor scene videos is detected by tracking points found in the foreground and using their vertical differences. If the gait frequency is in the range specified by the eigengait space, then the presence of an animal in the video is assumed.

Hannuna et al. [HCG05] also propose a method to detect animal movement in wildlife videos. Foreground points are identified and a bounding box is fitted over them, which specifies the dense flow area. Once again PCA is applied to a set of dense flows which describe movements. Projection coefficient variation reflects changes in the velocity and the relative alignment of the components of the foreground object. To understand the movement, motion flow images are constructed, where different colors represent the direction of motion. The dense flows can describe the internal motion of the object's torso and not just the legs. As the motion detection is very sensitive in order to pick up subtle torso movements, other movements are also detected such as imperfect bounding box placements, head and tail movements. Furthermore, the animals tend to move along uneven terrain with varying velocity, which detracts from the correctness of the motion flow with respect to what is being tracked. In order to improve the results, motion frequencies outside a certain range are discarded as drift or noise. Results show that quadruped gait patterns are detected even in low quality videos. Using this system, movements previously missed when scanned by eye were detected from videos.

Ramanan and Forsyth [RF03] build 2D models of animals' appearance from a video sequence. The animals are presented as a kinematic chain of rectangular segments. They derive an appearance model from video footage by identifying the main pool of pixels corresponding to the animal in each frame. Once the appearance of a specific animal can be recognised by the system, the result can be used to find such animals in a collection of images. Calic et al. [CCC+05] address the problem of animal recognition and classification from videos. However, as with the models of [GCT03] and [HCG05], they only deal with the processing of videos, finding patterns in the individual frames and classifying them, rather than actually tracking the motion of the animals.

In summary, computer vision techniques have proven effective in the detection and extraction of animal data from wildlife videos. However, the data is not always suitable for use in the simulation of 3D animal characters. The approaches of Favreau et al. and Wilhelms and Van Gelder are the only cases reviewed, where the authors have successfully managed to apply some of the motions captured by video to 3D models.

2.3 Animation methods

Character animation methods fall into many different categories including manual, procedural, physically based or inverse-kinematics approaches. In practice, a combination of techniques is used and the amount of user interaction can vary. We now provide an overview of animation methods used to date for quadruped animation.

2.3.1 Creation of the animation skeleton

Animation skeletons (sometimes called I.K (Inverse Kinematics) skeletons) are the standard control structure for animating 3D character models. They consist of a simplified version of the true anatomical skeleton, with only the main joints represented: an animation skeleton is defined as a hierarchy of local reference frames, each frame corresponding to a joint. Pairs (parent, child) of frames are called the "bones" of the skeleton. The appeal of using skeletons comes from the fact that they provide a very natural and intuitive interface for animating characters. For the final rendering, it is possible to deform a corresponding 3D mesh (corresponding to the skin of the character) in real-time according to its pose using standard linear blend skinning algorithms. Animation skeletons are usually created with respect to such a target 3D mesh. However, one still needs a good understanding of anatomy in order to produce adequate skeletal structures, especially for quadrupeds.

In the literature, a lot of methods have been proposed that try to approximate an animation skeleton from the geometry of the skin alone, represented as a mesh [KT03, LKA06, WML+06] or a set of voxels [WP02]. However, since the anatomy of a character is not always directly linked to its shape (see Fig. 2.8: the horse's spine should be along the top of the back), these skeletons cannot be used as such for animation: a user's input is usually required to adjust them.

A few recent methods have been proposed that insert a priori anatomical knowledge in order to create animation skeletons that are more related to the real skeletal structure of the animals. This knowledge can either be inferred from a set of examples [RFDC05, SY07], or provided as a template [MDMT+04, AHLD07, BP07].

In order to quantify the morphological variations of anatomically plausible quadruped animation, Revéret et al. [RFDC05] developed an intuitive method of morphing quadruped skeletons using only a few parameters, thus allowing the creation of skeletons for animals not in the initial database. Statistical analysis is applied to the collection of animation skeletons including a horse, goat, bear, lion, rat, elephant, cow, dog and pig. All skeletons share the same topology in terms of number of articulations and joint hierarchy. Each skeleton contains information about the position and orientation of each articulation. All the skeletons are normalised so that each has the pelvis in the same location and the spine column is the same length for every animal. Therefore, the variability in skeletal structure between animals can be explored independently of the size of the animal. Using Principal Component Analysis (PCA), the authors show that, using global frame coordinates and quaternion rotations, the translation and rotation data of the different models can be efficiently described by a linear model controlled by three parameters. These three parameters correspond to m_1: animal height (height from the hip to the ground); m_2: bending of the spine; and m_3: hoofed vs plantigrade; which are the three parameters used to define the shape of a geometrical model, as can be seen in Fig. 2.9. Geometrically controlled parameters are given more preference as they are easier to change and intuitive to use. The results show that the morphable skeleton can be fitted to any quadruped model using the three measurements mentioned above. Additionally, convincing animations can be achieved when combining the skeleton with standard smooth skinning for geometry attachment.

Rather than a database of input models and skeletons, Schaefer and Yuksel propose to use a set of input poses in order to compute an animation skeleton for a given model [SY07]. Their idea is to cluster the faces of the skin mesh into regions corresponding to rigid bones. Of course, the set of input poses should contain a rotation for each desired joint of the skeleton, unless this joint will not be recovered. The bone weights for each vertex of the mesh are then estimated using the technique described in [JT05]. Finally, these weights are used to define both the connectivity of the skeleton and the joint position, hence allowing direct creation of new poses. This method is illustrated in Fig. 2.10. In a similar approach, de Aguiar et al. propose to compute a skeleton from an existing mesh animation (see section 2.3.3)

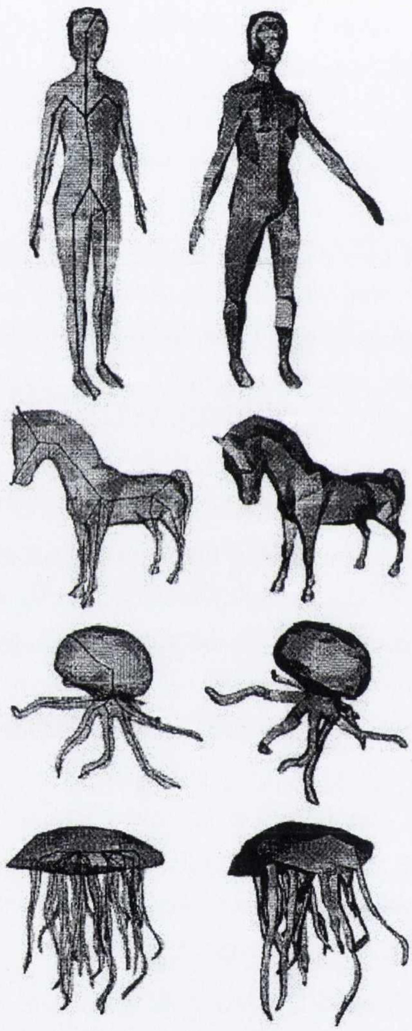

Figure 2.7: Skeletons can be automatically generated for different shapes of varying complexity (Wade and Parent [WP00]).

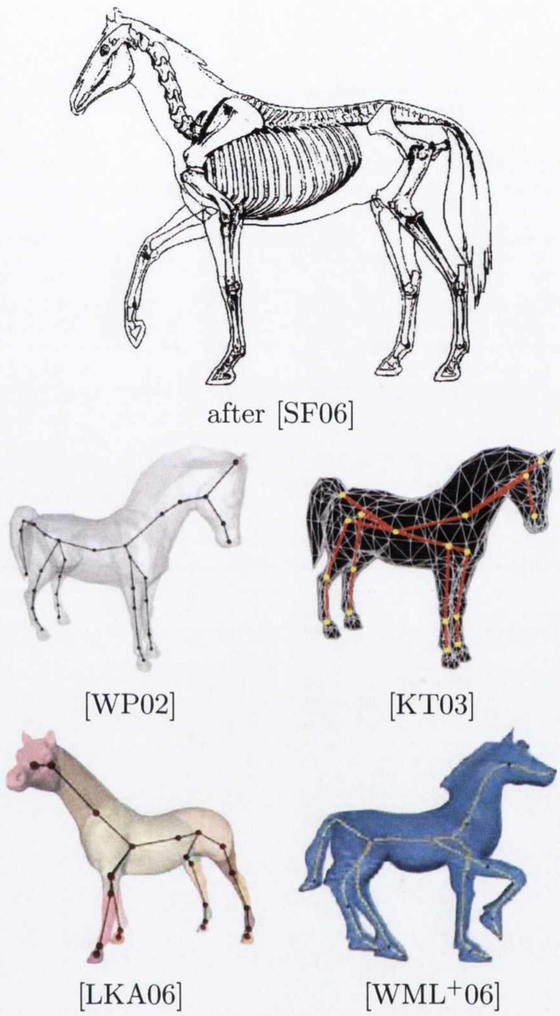

Figure 2.8: Top: skeletal structure of a horse. Middle and bottom: animation skeletons inferred from the geometry of the shape with different methods.

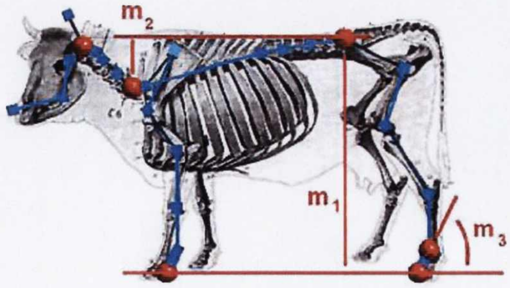

Figure 2.9: Three parameters controlling the morphable model (Revéret et al. [RFDC05]).

Figure 2.10: A set of input poses, mesh clusters representing rigid bones, computed skeleton (the yellow dot corresponds to the root of the hierarchy), and new poses created with this skeleton (Schaefer and Yuksel [SY07]).

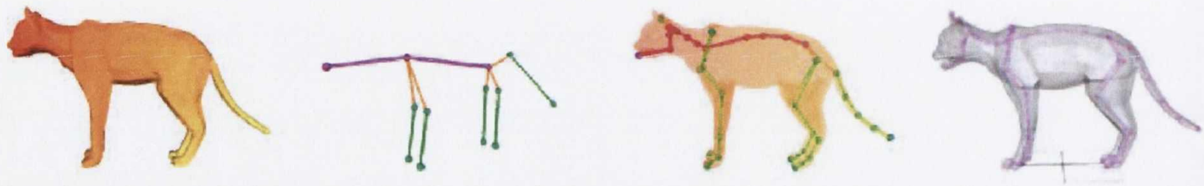

Figure 2.11: A cat model, its harmonic graph and the computed harmonic skeleton, along with a handmade animation skeleton in the last image (Aujay et al. [AHLD07]).

[dATTS08]. In this case, the set of input poses is replaced by a greater set of meshes with additional information, corresponding to the motion of the model.

In the case of bipeds, Moccozet et al. [MDMT$^+$04] propose to fit a template model (described in [SMT03]), which contains both skin and animation control information (including an animation skeleton), to any scanned model. The correspondence between the two models is set manually, but the authors propose some tools to assist the user, for instance for the location of characteristic points on the input data.

Aujay et al. describe a system which is able to create a skeleton for any given model [AHLD07], with anatomically plausible locations for joints in the case of bipeds and quadrupeds. A single point is selected by the user on the model as the starting point for the generation of the skeleton. This should be on the head of the character, in order that the system is able to recover the anatomy of the model. A harmonic function and a Reeb graph based on this function are then computed. The harmonic function is used to filter and recover the symmetry of the character's structure, resulting in a graph with unnecessary joints. This harmonic graph is refined starting from the user specified point and is embedded into 3D space, with anatomical information attached; this is based on a priori knowledge coming from an infographist's work. For quadrupeds the generated skeletons have been compared

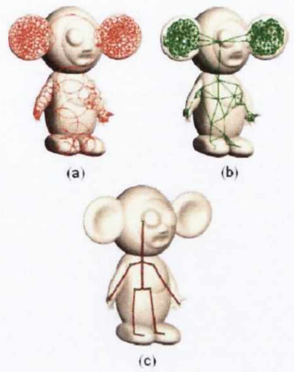

Figure 2.12: (a) Spheres packed into the mesh (b) Graph constructed based on sphere intersections (c) Embedded skeleton (Baran and Popovic [BP07]).

with those created by hand, see Fig. 2.11. The results show that, not only visually but also when comparing the parameters from [RFDC05], the automatic method produces valid skeletons ready for animation. Moreover, the method works even when the character's poses are different. A user interface allows for more bones to be created manually, for example if more bones were needed for better control of an elephant's trunk.

Another recent paper describes a similar system, called Pinocchio, which embeds a skeleton into a model [BP07]. It also attaches the skin of the model to the skeleton automatically. The skeleton construction works as follows: Spheres are fitted into the character's body and a graph is constructed based on their centres and by adding edges between spheres that intersect. This graph is then used to help embed a template skeleton. The different stages can be seen in Fig. 2.12, in the case of a biped. To attach the skin, a heat equilibrium approach is used, where the character is treated as a heat-conducting object. To find the weights of the vertices, the temperature of a bone is kept at 1 while the others are at 0. Equilibrium temperature is calculated at each vertex on the surface which then specifies the weight of the bone. This results in smooth skinning, due to the second order nature of the heat equation. Note that the straight forward solution of using the weights on the proximity of a bone to a vertex does not work: as the skeleton is two-dimensional, for fat characters the vertices of the belly can end up attached to the arm bones. Results show that the Pinocchio system can embed a template biped or quadruped skeleton into different models, which are then ready for skeletal animation.

Skeletons that better reflect the nature of an animal should lead to better animations. On one hand, generating the skeleton automatically saves a great amount of time and does not require specific skills or knowledge of anatomy from the user. On the other hand, pa-

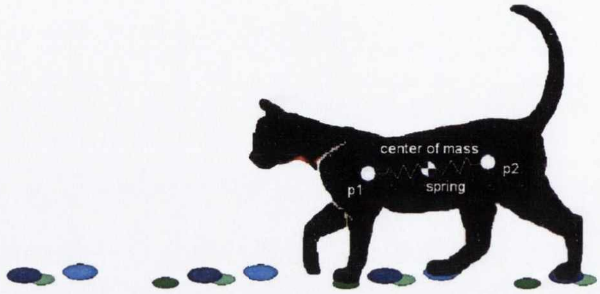

Figure 2.13: Two point masses are used in the simplified physics model to make the overall body movement look realistic. The footprints guide the position of the feet, the small circles represent the front legs, and the big circles represent the back legs (after [TvdP98]).

rameterising the skeletons provides a way of simulating new animals very quickly. It can also give better control over the levels of detail that may be needed for specific parts of a model.

2.3.2 Direct and inverse kinematics

One of the earliest examples of computer generated quadruped animation is the PODA computer animation system [Gir87]. The user can control the movement of the animal through a combination of inverse kinematics and dynamics simulation of the limbs and body motion. Positioning of the limbs can be achieved either through inverse kinematics, by moving the end-effectors (hands and feet) to a desired position, or by rotating the joints using forward kinematics. Additionally, the user can adjust the angles of the joints while keeping the end-effector in place. The different "postures" are assembled to produce the final motion path. The system interpolates between the different limb positions so that the movement is smooth. The limb trajectories are recalculated during body movement to accommodate variations in foot placement and body speed. Co-ordination is allowed, where each leg can have a state indicating whether it is in support phase (on the ground) or transfer phase (in the air). Depending on gait, the time spent in the different leg phases varies. An animal's body trajectory is solved for by taking into account its vertical motion (using Newton's equations of motion for upward velocity), horizontal motion (defined by the animator using a cublic spline) and angular motion (found by calculating the direction of movement from changes in the horizontal position and solving Newton's equations of motion). While individual limb positions are controlled by inverse kinematics, the overall movement, the timing of foot placement and body dynamics are controlled using Newton's equations in order to keep the body balanced. This can allow for better timed and more natural motion.

Torkos uses a similar approach of a hybrid physically-based/kinematic system to animate a 3-D cat [Tor97, TvdP98]. Spline trajectories represent the state of the quadruped over time. Footprint locations and their timings are additional constraints used to generate various motions such as walking, galloping, jumping, push-ups and skating for four-legged animals. Physics is used to make the overall movement look realistic while inverse kinematics is used for details, such as positioning the leg joints appropriately. Two point masses are connected by a spring that models the internal forces of the back. The point masses are located at the hips and shoulders of the animal, see Fig. 2.13. The motion of the complete skeleton is reconstructed from the point mass trajectories. A trajectory is calculated from one footprint to the next. Optimisation ensures that the foot is at a reasonable distance from the body and collision avoidance is also handled. Inverse kinematics determine the arch of the spine and the position of the leg limbs, as well as the movement of the head and tail. A user can control where to place the footprints on a plane, specifying the x and z co-ordinates, while y co-ordinates are calculated according to the terrain. A user interface provides suggestions about the order and timing of the feet and their positions. Additional features include automatic generation of footprints where the user specifies the path the quadruped is to follow (Fig. 2.13). For fully autonomous behaviour, the system can generate footprints in real time, e.g, for a random wandering behaviour. To ensure that the overall movement of the 3-D cat is visually appealing, Torkos adds a "comfort" term as a constraint to make sure that the quadruped does not deviate too far from allowable positions.

Kokkevis et al. [EKB95] describe a way of creating autonomous animal motion using kinematics, dynamics and control theory. The quadruped is divided into body and legs subsystems. The body subsystem consists of position and velocity of centre of gravity, orientation and angular velocity of the body. The legs subsystem consists of the position of each leg, the state of each leg (up or down) and the force it exerts of the ground. Depending on the target position and velocity of the animal, forces are combined for the body using the dynamic controller and are distributed to the legs. The leg forces are exerted on the ground, thereby generating a reaction force on the legs, passing up to the body, thus making the animal move. Dynamics is used to control the legs when they are on the ground and kinematics is used to position them when they are in the air, all calculated by the gait controller, which also computes the velocity of the body and the step-size based on the gait pattern used. They successfully simulate a dog performing walking and trotting over different terrains, with minimal or no user interaction. Providing the means to generate character motions automatically allows the user to concentrate on other aspects of the animation, such as creating realistic behaviour.

While inverse kinematics is an intuitive approach to creating character animations, Torkos points out that physics can also be a useful tool for making the overall movement believable,

Figure 2.14: (Top) the triangle mesh, (Middle) Estimated bones, (Bottom) Vertex weighting to the virtual bones (James and Twigg [JT05]).

allowing for a level of detail approach.

2.3.3 Mesh animation

It is also possible to create realistic animations by manipulating the skin of the mesh directly. Vertices of an articulated mesh tend to move relative to their neighbours rather than independently. By modelling these deformations correctly, mesh manipulation is possible without the creation of a skeleton. While James and Twigg do not use skeletal bones in their work, they do calculate "proxy bones" that deform the skin [JT05]. The input models do not have any predefined bone transformations, so proxy bone transformations are estimated before calculating vertex weights. To calculate these proxy bones, triangles with similar transformations are clustered. Once the number of bones is known the vertices are weighted. The results are shown in Fig. 2.14. During animation, a character may end up in atypical positions. Furthermore, large numbers of mesh triangles are needed for deformations, so the algorithm that assigns vertices to their respective bones has to be very robust. James and Twigg use a mean shift clustering algorithm to calculate the number of bones and the triangles associated with each bone. Each vertex can be deformed by one or more bones. To animate the model, a skinning transformation is found that maps an undeformed mesh to a sequence of deformed models. Depending on the pose of a deformed mesh, new proxy bones might be created. These are referred to as flexi-bones and they ensure smoother skin deformation. Using this method, James and Twigg [JT05] were able to simulate a large herd of skinned animals (horses, camels and elephants, see Fig. 2.15) at interactive frame rates.

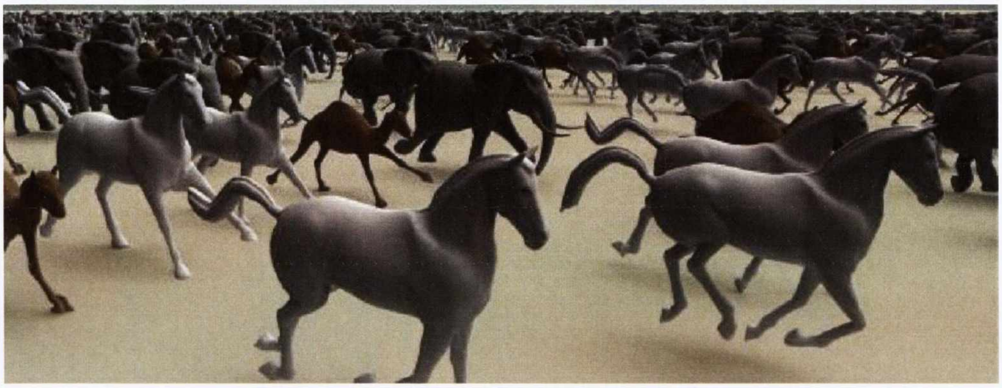

Figure 2.15: Skinned mesh animations (James and Twigg [JT05]).

Der et al. [DSP06] reduce the deformability of a model, thereby making it easier for a user to animate various creatures. They demonstrate their algorithm on a horse, elephant, mouse, gorilla, human and a dragon. By moving vertex handles created for the hooves, a horse model can be placed into a galloping pose. The deformability of the model is reduced by taking into account the fact that neighbouring vertices move together. In this way, a compact set of control parameters can be used to change the pose of the model and allow for direct manipulation of the model. Groups of vertices that are affected by the same control parameter are replaced by a proxy vertex. This in turn speeds up the calculation of transformations for those vertices. Vertices that are affected by more than one control parameter are wighted depending on which control parameter has the greater influence. However, the user is still allowed to manipulate each vertex of the model. Additionally, inverse kinematics (similar to the more conventional IK algorithm for jointed rigid skeletons) is applied to the reduced model to re-position the vertices as the controllers are being moved around.

Shi et al. address the problem of mesh manipulation at interactive framerates [SZT+07]. New poses of 3D models such as an armadillo, camel, horse, cat or dinosaur, are created by specifying constraints on leg length preservation, fixation of the foot on the ground, balance preservation and self-collision handling. These limitations help to maintain the volume and shape of the object. This technique combines the more traditional skeleton based rigging and mesh deformation. The input consists of a triangle mesh and an associated skeleton, where the vertices of the mesh have already been matched to the bones of the skeleton. Each vertex is defined as a linear combination of the locations where it would be if it were following the movement of the associated bone. For each bone, virtual vertices (tetravertices) are added to create a tetrahedron (tetrabone). As the tetrabone moves, so too do the tetravertices, thus capturing the motion in the process. Tetrabones are used to keep the length of the bones

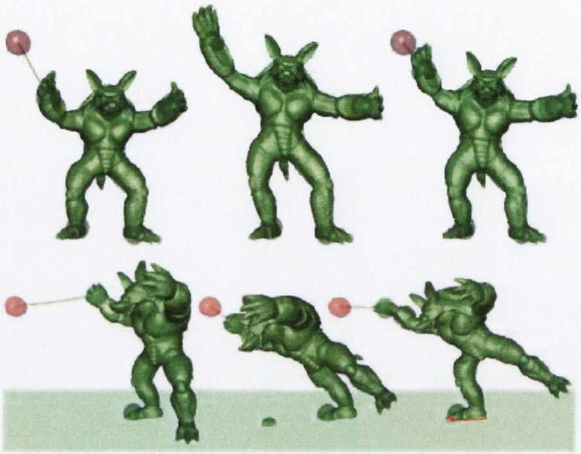

Figure 2.16: (Top) When the length constraint is used, the whole body moves (right) as opposed to just dragging the arm (middle) column. (Bottom) Using the centre of gravity constraint, the small red sphere shows the target position of the barycenter (Shi et al. [SZT+07]).

constant and also to deform the skin around bones in such a way that character volume is preserved. Centre of gravity is calculated by computing the barycenter of each bone when the mesh is in the original pose, and then recalculating it once the transformations are applied. By ensuring that the projection of the barycenter is on a given position on the floor, the method ensures that the position of the mesh stays balanced, see Fig. 2.16. In order to prevent sharp bends in joints, restrictions are placed on the ranges of the bones involved.

Wilhelms and Van Gelder [WG97, SWG02] modelled and animated animals using accurate models of bone, muscle and tissue. In [WG97] these components are represented using triangle meshes and ellipsoids. As the bones move the muscles change shape, which in turn deforms the skin during the animation. Modelling a character involves building a body hierarchy, designing individual muscles, voxelising the components and mapping the skin vertices to the nearest underlying component. The bones and the tissue are parameterised so that the same components can be used for different animals. Muscles are represented as deformable cylinders or ellipsoids that are divided into discrete parts. The origin and ends of a muscle are attached to specific bones. A user can change the shape, orientation and location of the origins and ends of a muscle. Once a joint is opened or closed, the muscle shape is recalculated. The cross section of a muscle increases as it contracts and decreases if the muscle relaxes. The skin's vertex is attached to the closest underlying component and it can move relative to the muscle or tissue. This is because the skin is attached using virtual anchors which give the skin a more elastic appearance.

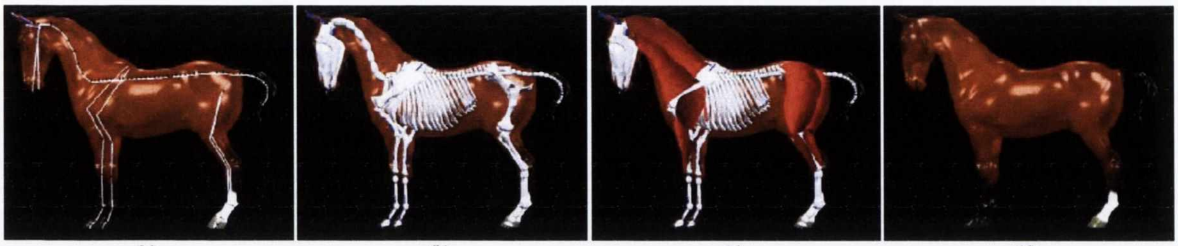

Figure 2.17: (a) segment hierarchy - input into the system (b) skeleton (c) muscles (d) skin (Wilhelms and Van Gelder [SWG02]).

Simmons et al. [SWG02] similarly deal with the creation of 3D models of creatures (see Fig. 2.17), where they deform a canonical representation of a character to create new creatures. The transformation is applied to the segments hierarchy (this is input into the system) and this is then propagated to the bones, muscles and skin. The input information can be gathered from either taking measurements of a real horse, or using a 3D horse model. The skin is represented by a triangle mesh. A subset of the skin vertices, designated as "feature" vertices, is used to deform the attached mesh skin. The feature vertices are created in addition to the vertices that make up the mesh of the model. Unlike in [WG97], only the feature vertices are anchored to the underlying components, which move as they move and the rest of the vertices are interpolated. This generic model can be morphed to another target animal. In order to achieve this, eighty-nine marker locations are placed on the generic model, thus denoting the points that are necessary to estimate joint locations and the structure of other components such as muscles. This information is then used to transform the geometric hierarchy of the canonical model, along with its bones and muscles and skin, to match that of a target animal (which can be in a different pose). Using this technique, one is able to create very realistic animals of different sizes which can then be used in animation.

While Walter et al. [WFM01] describe a technique for growing patterns on different mammalian models, their method can also be used as an animation technique. A user segments the model into cylinders so that it is represented as a group of geometries. In order to better control the cylinder shape, a user can change its parameters or "features". The cylinders make up a hierarchy where translation and rotation is inherited from the parent cylinder but, due to the pattern generation, scaling is kept local to a specific cylinder. By changing the orientation of the limbs, one can create different poses for the models which are then used for animation.

There is a wide application of mesh manipulation techniques. It can be used for real-time animation, as seen in the work of James and Twigg, where skeletons are not used explicitly but the technique is still based on the idea of vertices being controlled by bones. Wilhelms and Van Gelder use mesh manipulation to create anatomically realistic animal models (bones, muscles and skin tissue are modelled) ready for animation. For both techniques the emphasis is put on the resulting *shape* of a pose. This avoids unsightly kinks and bends in the skin of a character and also ensures that volume is preserved.

2.3.4 Automation

The quadruped motion we are all familiar with these days is usually achieved through hand animated key frames, which are then interpolated to provide smooth motions. Commercial or freely available software contains many features to aid hand animation. However, to produce quality motions, the user needs to be very skilled and even then it is still a very painstaking process. Hence, there is a great demand to automate the animation of some of the more common motions characteristic to quadrupeds. Feature animations such as Madagascar and Shrek 2 combine hand animations with some procedural techniques to aid the animators [GMC+05, GAD+07]. This idea has also been adopted by other researchers.

Kuhnel [Kuh03] used footage of horse training videos to recreate realistic animations of virtual horses and riders in Maya. Automating the movement of the horse and the rider leaves more time to animate the finer details of the scene such as arm movements and facial expressions. By simulating a horse performing different gaits, the motion of the rider can be automatically generated (balancing, bouncing, leaning forward or shifting their centre of gravity). To determine whether the simulated motions of the horse and the rider are correct, they are displayed alongside reference footage and a mathematical model of the movement is iteratively modified to better reflect the original video. These mathematical comparisons thereby create a base set of rules, supporting a predictive model of what will happen to a rider in any new situation. In order to achieve this, the rider and horse are rigged in Maya. The horse rig includes control points for the hooves, rear rotational pivot and a mid-back pivot (which is located in the centre of the spine). However, the authors focus on the generation of the rider's movements where separate control points are used to calculate the movements of the rider's hips, legs, spine and arms.

Francik and Trybicka-Francik [FTF03] created a system called KINE+ that imports model data (e.g., skeleton information for bipeds or quadrupeds) from either 3D Studio Max or Maya software. The workflow pipeline includes modelling the object in 3D Studio Max or Maya (this includes adding a skeleton to the object), exporting the skeleton of the model into the KINE+ system where it is animated using C++. The animation data is stored in a file which is then imported by 3D Studio Max or Maya to be skinned and rendered. The KINE+ system is based on the Discreet Character Studio package and uses its bone hierarchy. To help with the animation of a character, KINE+ supports the following features: forward kinematics, inverse kinematics, collision detection and avoidance, key frame management and exporting of the animation.

Above we have discussed the different approaches that can be used for animating quadrupeds. While skeleton manipulation is the most popular method for creating such motion, some researchers have found that equally good results can be achieved using mesh manipulation. By parameterising the motion and the skeletons it is possible to create a range of gaits suitable for different animal sizes.

2.4 Control, Interaction and Behaviour

2.4.1 Physically-based modelling of animals

Physically-based animation involves simulating the effects of a combination of forces such as torque and gravity, thus resulting in body movement. Some researchers have created simulations of physics-based animal motions with varying degrees of user interaction.

There are many advantages to physics based modelling. By considering that any movement is a result of forces that move a body, gravity, balance and speed can be used as controlling parameters. Characteristic motions that vary depending on a character's situation are easily achieved once physics controls are set up. Some of the latest work on procedural physics simulation can be seen in the new game Backbreaker, where American football tackles are all generated as the game is played, so they look different every time. While humans have been the primary focus in state of the art physically based techniques, some researchers have been active in developing methods for animal motion also.

Marsland and Lapeer [ML05] minimise the need for user input when animating a physics-based trotting horse, by using video footage to position the horse's legs correctly. The horse is modelled as a collection of connected bodies, which are moved from the current state to a desired state through the application of torques. Each bone in the skeleton has mass, centre of gravity, inertia, width, depth and height. The bones are connected using hinge joints, the limits for which are set based on the results of studies of horse motion. Information about the size of the angles at the joints is taken from each frame and used to calculate the torques needed to move the physics based model into the desired position. This is achieved by the animation controller. A heading controller keeps the physics-based horse walking towards a target position. An angle is calculated between the current heading direction and the desired direction, which is then added to the joints of the horse's limbs.

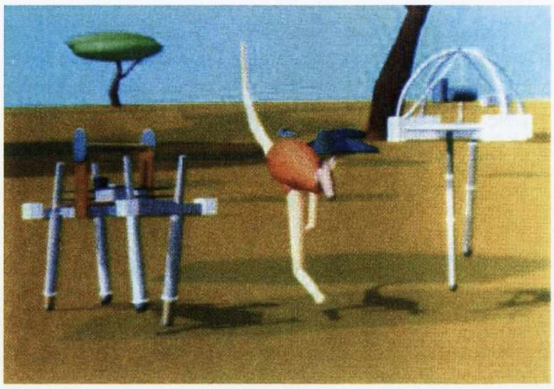

Figure 2.18: Different models used in actuated systems (Rainbert and Hodgins [RH91a]).

2.4.2 Designing locomotion controllers

One of the aims of computer graphics is to create systems where the user input is limited to directing a character around a scene. The movement of the character is controlled by a system that integrates equations of motion derived from physical models, as in the work of Raibert and Hodgins [RH91a]. Users can specify the input to algorithms that control the behaviour, but they do not manipulate the model directly. Their technique was used to create a computer animated cartoon that depicts automated computer characters with one, two and four legs [RH91b] (Fig. 2.18). The challenge here is to simplify a high-level request so that it is understood by the underlying processes, while at the same time creating realistic motions. All characters in the system have control algorithms that maintain balance, hopping, speed, posture and elastic energy stored in the legs. The control inputs (speed, gait, path and the initial position of the legs) guide the actuator forces, which in turn make the animal move with a desired speed and direction. The movement of the legs is modelled on a spring system. The torques calculated during the support phase are used to keep the body upright. A finite state machine is used to synchronise the legs by invoking the correct algorithm. The same technique is used for all the animals: a bipedal run is treated as a one-legged hop where the functions for a one-legged hop are applied to each leg in turn. A quadruped trotting is treated as a biped, where diagonal legs form pairs. One of the advantages of this system is that, because of the higher level of control, the motion of individual legs does not need to be specified.

Similarly, in the work of Laszlo et al. [LvdPF00], the details of the physically-based animations are hidden from the user. In this, way the user is free to experiment with the behavior of a character in the scene rather than having to control the movement of their bodies or legs explicitly. At interactive frame rates, a user can position objects around the scene and as a result can quickly produce satisfactory animations. One of the more important features of this approach is that a user can control the speed of motion through the user interface. This improves upon previously mentioned methods where the parameters have to be specified before the program is run. An animator is allowed to use his intuition about motions to create animations for planar characters such as lamps, cats and humans. The success of such a user interface relies on the design of effective motion controllers. For simple cases, mouse movement can be used to control the joint movement of a character. This can produce jumps, shuffles and flips. For complex characters, such as a cat, keystrokes are assigned a set of desired angles assumed by the legs. By selecting different positions of the legs, a user can create different animations. If there is a need to correct the position assumed by the cat at any stage during the animation, checkpoints are provided that allow for more detailed positioning of the cat's limbs. The dynamics of the character are important as it means that the user will retain the necessary rhythm and timing of the action. Other complex motions like bipeds traversing monkey bars or climbing a ladder are aided by the use of IK primitives. A position that a hand is to grasp is marked and from there inverse kinematics is used to establish the desired joint angles. In addition, a finite state machine chooses the next arm that is to grasp a target. From the above we can see that, if the user interface is intuitive enough to use, and the underlying physics is abstracted appropriately from the user, many different effects can be achieved using different characters.

Ideally a motion controller should be able to control the overall movement of a character by enabling it to perform a single motion in different ways, e.g., a walking simulation that can display varying speeds and turning as well as realistic movements over uneven terrain. Different applications require different styles of animations: feature films require tightly-scripted motion control while computer games require intelligent autonomous motion. Simulated creatures capable of autonomous behaviours potentially require less effort to animate and are reusable. This involves the varied motions listed above and it also requires planning for a given task. For example, given a task such as "walk to the door", a character needs information about the environment in order to plan its route. To achieve such high level character control, one needs a rich family of well-studied control techniques that are also easy to use.

Van de Panne provides an overview of methods used for controlling simulated human and animal motions [vdP98, vdP97]. An important point illustrated is that, in order to build a successful system, there needs to be a way of evaluating the choices used to represent control functions. Analysis of real animals such as that provided by Alexander [Ale84], provides many insights into the movements of animals and the parameters that control them. However, this does not necessarily provide enough information to build a working controller.

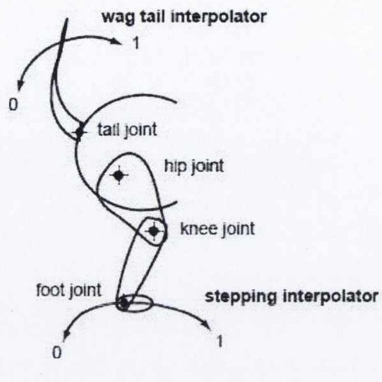

Figure 2.19: Inverse kinematics is used to move the foot. A value between a 0 and 1 sets the leg along the stepping cycle, similarly for the wagging of the tail (Blumberg and Galyean [BG95]).

In [vdP96], van de Panne uses parameters to physically simulate different motions. By studying the motions produced, for example by looking at the speed with which a creature moves with different parameters, the optimal values can be found. Similarly, in order to animate a character, each has a starting pose and a desired pose. The torques needed to move the limbs from one pose to the next are calculated. Different trials are run to see which values for limb position and orientation give the quickest and most accurate result. Further enhancements can be used to speed up the process, such as using different levels of detail, as in the work of Torkos [Tor97], where the cat is simulated using physics for the overall body movement and IK for the legs. Whichever technique is used to help control the movement of a character, the end result needs to be a smooth realistic motion in real time. There is usually a compromise between the required level of control over the motion and the effort required to specify the motion.

2.4.3 Interaction and behaviour

Realistic behaviour is important for compelling quadruped simulation. Even if the generated gait is perfect, if the animal does not behave characteristically, then its motion may be perceived to be anomalous.

Blumberg and Galyean [BG95] explored these issues by creating an interactive environment with virtual characters. They describe a system architecture that is able to propagate high level instructions to the geometry of a character, resulting in an animal that responds to human commands. The Behaviour system deals with instructions such as "find food" or "sit down and shake" as well as more simple commands like "move to" and "avoid obstacle".

These commands are translated into motor commands such as "forward", "left", "sit". The Motor system then uses these instructions to produce coordinated motion such as walking, which in turn deforms the geometry of the character. This architecture is applicable to all creatures, as commands such as "move to" can be used for animals, humans or even dynamic objects such as cars. A controller ensures that "move to" is interpreted correctly by each creature. The movement of the limbs is controlled by inverse kinematics. As can be seen in Fig. 2.19, a value between 0 and 1 controls the position of the foot along the curved line, and the same technique is used to wag the tail. The behaviour of the actor depends on its relationship with the environment, whether a task is achievable or not, and also the priorities of actions.

Using artificial intelligence including action selection, machine learning, motor control and multi-agent coordination, Tomlinson and Blumberg [TB03] were able to bring their characters to life, displaying the same characteristics found in real wolves. They combine hard-coded behaviours (walking) with learned behaviours (being able to react to the environment), resulting in wolves that "grow" into different personalities as time goes on. The wolves are able to store previous interactions in their memory, which can lead to new emergent behaviours. The wolves will learn from past experiences how to interact with other wolves in the pack, e.g, whether they are the dominant figure or not, and interpret other social cues. Users can also interact and control the virtual wolves by howling, barking, growling and whining into a microphone. By studying the social interactions between animals, a lot can be learned about behaviour within a pack. These actions can be broken down and even applied to human characters. This is especially useful in environments where autonomous behaviour is required.

The desire to bring characters to life dates back as far as the first Disney cartoons, only today the demands are much greater. People are no longer satisfied with sitting back and watching a feature animation of hand animated characters telling a story. Today, interaction with characters in different and new ways is growing in popularity. The challenges to be met include creating flawless motion for any movement imagined and providing a way for the characters to think. Using animal characters as a testing ground for autonomous behaviours is appropriate since in real life they are independent beings.

Hecker et al. describe a system that is capable of animating highly-varied characters [HRE$^+$08]. Through user interaction and procedural animation, unique characters can be created and animated within minutes. An animator can control his character using familiar key-frame animation techniques. They can set key-frames and edit motion curves. During this animation phase the motion data is recorded (positions and orientations of body parts). This motion can then be retargeted according to different character morphologies, resulting in characteristic movement for each unique model.

In SPORE, a character is made up of different body parts and each part has its own built-in function. Locomotion of legs, the opening and closing of mouths, eyes and hands are all controlled procedurally. Movement modes specify whether the movement is relative to the ground, the size of the character or the length of its limbs. The secondary relative movement mode can specify tasks such as "shake hands" or "clap hands". The leg movement is completely procedural, generating a cyclical motion. Legs of different size on the same character can display different gaits at the same time (two short legs trotting while long legs are walking). A Particle Inverse Kinematic Solver is used to generate the movement for the characters. The advantages of using this system is that it is high performing, produces realistic results and gives path independent solutions (the solution at time step t does not depend on previous solutions).

A bind phase resolves how an animation will be played on a character. During this phase, the shape of a character is evaluated depending on whether the spine is prone or upright and how many graspers and/or feet a character has. Animations are played for different tasks and a user can pick whichever variation of the animation suits their character the most. Depending on the task and the number of feet/graspers, the animation may vary.

More recent work by Wampler and Popovic [WP09] describes a fully automatic way of generating gaits and morphologies for a monoped, biped, velociraptor, horse and a pentaped. Given the shape of an animal, they are able to re-proportion the skeleton so it can move more efficiently and also solve for the timing and the pattern with which the feet should move using constraints dictated by the laws of physics. This includes conditions such as, the skeleton should remain above the ground and the feet should have the correct velocity when in contact with the ground. Each animal is defined as a kinematic tree of connected limbs. When a foot is in contact with the ground there are 6 variables which calculate the force and torque exerted on the foot. There are an additional 3 variables that define the passive actuation characteristics representing the stiffness, rest state and dampening. These variables make the motion smoother and more believable. An objective function is used for optimisation where spacetime optimisation minimises the muscular exertion of the animal, high velocity joint motion is penalised to avoid low-torque wiggling motions, perpendicular motion of the head to the rest of the body is also penalised to keep it stable and an extra function is used to keep the animals head facing in the direction of the motion. To optimise the skeleton two variables are added that control the radius and the length of the limbs. The mass of the bones is also scaled so that the bone proportions maintain a constant density. To optimise the process a hybrid technique is used combining a derivative-aware spacetime constraints with a derivative-free approach.

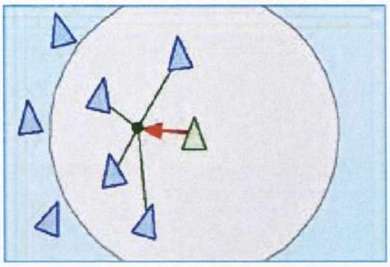

Figure 2.20: Cohesion: A boid moves towards the average position of the nearest neighbours (C.Reynolds [Rey87, Rey99]).

Using [HRE+08] and [WP09] as most recent examples we can see that within the research area the simulation of more complex characters is becoming more popular. This is further encouraged as animals often feature in the industry, more recent examples are Dug and Alpha dogs in Disney's Up [ABHM09] and the cattle simulation in "Australia" (20th Century-Fox Film Corporation) [KBF+09].

2.4.4 Flocking and shepherding

There are many commercial companies that specialise in simulating large crowds for popular feature films such as Lord of the Rings (New Line Cinema), Narnia (Disney) and WALL-E (Pixar Animation Studios). In general, such simulations follow the same basic principles set out by Reynolds, which are then combined with fuzzy logic that controls the "brain" of the characters (this includes the actions a character makes and other things such as foot placements). Traditional and laborious techniques for animating flocks and herds using scripting were replaced by automatic methods thanks to the pioneering work of researchers such as Reynolds [Rey87, Rey99], Tu and Terzopolous [TT94], Hodgins, Brogan and Metoyer [HB94, BMH98] and many others. Reynolds showed how a realistic flock could be simulated by isolating a few elementary rules that each agent obeys, in particular motivated by the opposing desires to both stay close together and to avoid collisions. Individual behaviours such as cohesion, separation, pursuit, evasion and obstacle avoidance can thus be simulated to produce quite complex patterns of behaviour, see Figs. 2.20 and 2.21. One of the first examples of using scripted herds in computer animation can be seen in Disney's The Lion King wildebeest stampede [Wil94].

Lien et al. [LBS+04, LRMA04] addressed the problem of shepherding flocks. A key rule is that a sheep reacts to a dog by moving away from it. The shepherd's movements are broken down into approaching and steering: To move the flock, the shepherd moves to steering points

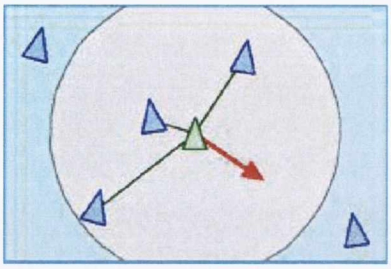

Figure 2.21: Separation: A boid moves away from the nearest neighbours in order to avoid any collisions (C.Reynolds [Rey87, Rey99]).

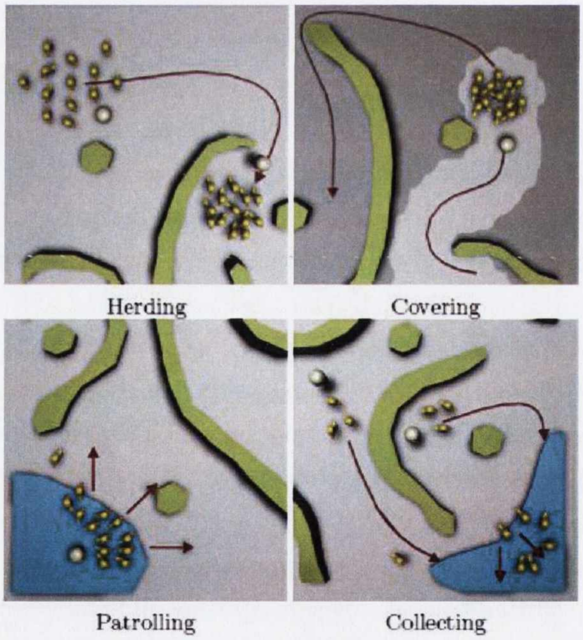

Figure 2.22: Shepherding behaviours: herding, covering, patrolling and collecting (Lien et al. [LBS+04, LRMA04]).

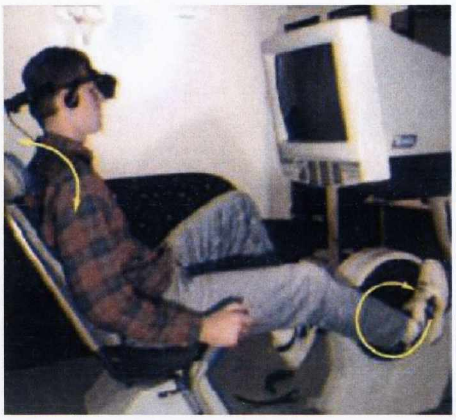

Figure 2.23: User acting as the shepherd (Brogan et al. [BMH98]).

Figure 2.24: One legged robot sheep outside their corral (Brogan et al. [BMH98]).

which are located near the flock - the approaching stage. Then it needs to be able to steer the flock towards a target in the most efficient way possible. A variety of movements are combined and tested to form different shepherding behaviours; herding (steering flock from start position to goal position), covering (flock guided to visit all positions in an environment), patrolling (preventing the flock from entering a forbidden area) and collecting (gathering a scattered flock). This work concentrates on implementing optimised shepherding behaviours, while the "sheep" observe the fundamental flocking behaviours mentioned above, see Fig. 2.22.

Flocking simulations can be found in many different research areas such as robotics and also in biology, where these simulations are used to study the behaviours of many different animals. Vaughan et al. have built a robot system that exploits and controls animals' behaviour to achieve herding of flocks of living ducks [VSH^{+}98]. The algorithms used to simulate shepherding and flocking behaviour in this virtual system are similar to those presented by Reynolds. Brogan et al. [BMH98] also simulate a Border collie environment populated by a small group of hopping one-legged robots, in which the user acts as the shepherd and herds the robots into a corral.

Figure 2.25: A battle scene from "The Lion, the Witch and the Wardrobe". Image © Copyright Disney Enterprises Inc. and Walden Media Llc. All rights reserved.

More recently, Kolve et al. from Rising Sun Pictures discuss the different components used to produce realistic cattle herds of 5000 bovines in Australia (20^{th} Century-Fox Film Corporation) [KBF$^+$09]. A crowd container holds all the data necessary to generate a crowd, including the initial position and orientation for each agent in the crowd. The animals are animated using an animation library that contains interconnectable loops and transitions of cow animations that are applied to the characters using functions that find the next best fit (similar to the way motion graphs work). Collision avoidance is resolved using a relaxation technique that checks whether the roots of animal models are intersecting. This is quicker than checking for collision detection at high levels of detail such as when you use the whole mesh. If the roots of the animal characters do not intersect, the rest of the character animation is computed. Smaller groups of herds are first computed and later these are put together in different ways to produce much larger herds that can take different shapes.

2.5 Case Study: The Chronicles of Narnia

The characters created for the Chronicles of Narnia represent some of the best examples of quadruped animation seen to date. In this section we will review the work of companies such as Rhythm and Hues, Sony Pictures Imageworks and describe some of the techniques they have used for character creation and animation in "The Lion, the Witch and the Wardrobe" [WB05].

Rhythm and Hues had the task of creating Aslan (the lion) for the first film, which involved building, rigging, lighting and controlling, muscle/skin dynamics and fur collision detection. Other creatures that were created included centaurs (with the head of a human and the body of a horse combined), fauns, minotaurs, minobars, cyclopses and goblins.

Maquettes were created for all the characters, which were then scanned to create their virtual counterparts and rigged for animation. To gather information about the look and behaviour of certain animals, animators from the studios spent time in a cage with real lions, leopards, cheetahs, bears, hawks and other animals needed for the production. Other sources included books and nature videos. A lot of high definition detail was recorded in order to examine specific movements as well as the deformations of the muscle and skin.

Muscles and soft body creatures such as rhinoceros have skin that tends to wiggle, so Harmonics were used to create such skin dynamics. Oscillations move the skin in response to body movement. This technique was applied to Aslan to achieve realistic muscle vibration. For speed gains, volume preservation was ignored and any skin artefacts were corrected by artists. Fur dynamics were used to model the clumping of fur on Aslan's mane and associated collision detection. Two controllers were built for simulating wind blowing through the mane. Dynamic wind controls handled the motion of the wind against the mane, while Pelt wind controlled the movement of individual hairs, such as the tips that can move independently, to create a frizzled effect.

The facial animation of Aslan was created using a shape-based blend system that uses pre-built facial poses, selected by an animator, which can be combined and animated. On top of this system is an additional layer of muscle and traditional deformers. Another system was built to allow modellers to create facial poses by hand. This system used the defined poses to determine which muscles to deform on the model. Other features allowed for larger wrinkles to be added and harmonics were used to add vibration to the skin of the face. An important consideration when animating fantastical animal characters is to retain the intrinsic animality of the character, while simultaneously endowing it with antromorphic properties. In order to achieve this, plausible facial expressions that were anatomically possible by a real lion were created to express to emotions (happy, sad, angry).

To create convincing centaurs, motion capture data was collected from both a human and a horse. For a more natural transition between the two species, a human torso was placed where a horse's neck would have been. To gather ideas about the movements of centaurs in the battle scenes, video clips of horse actions were superimposed onto the clips of a human performer. The actions between the "actors" were matched to ensure that, on video, it looked like the head controlled the actions of the body. Later, a motion-tree was used to select the movements a character could perform based on what state they were in. During the battle scenes the number of motions needed caused the motion tree to grow to include several hundred actions.

For the faun creature, an actor was required to walk on tip-toe. His motions were captured and retargetted onto the faun rig. The software compensated for the difference in leg proportions between humans and animals. This means that the heel of the human was positioned to match the heel of the goat, which is at a much steeper angle and higher off the ground. The studio used motion capture and hand animation, or a mixture of both called "supanim", for the animation of the characters. The faun character is completely motion captured, while the gryphon (half eagle-half lion) is animated by hand. Werewolf motions were created using a combination of both methods, as the motion is that of a chimpanzee applied to a quadrupedal rig.

Sony Pictures Imageworks created some characters such as the wolves, fox and beavers. Beavers, who are by nature quadrupedal animals, were created to walk as bipeds for the film. Once again, the challenge was to ensure that the animals were realistic but with human qualities. In order to animate the beavers, they were given two sets of muscles. One set was used when they were in biped mode, while the second set was used for the quadrupedal mode. These muscles could be turned on or off depending on the motion being performed. Furthermore, their muscular structure allowed for the simulation of fat, jiggly bodies.

Finally, real dogs were used for the shots that included the wolves. During the filming, the dogs were wagging their tails and their tongues were hanging out. In order to make them look more menacing, the wagging tails were replaced in many of the shots.

2.6 Discussion

In this chapter we have discussed the rich and varied history of quadruped synthesis ranging from 2D cartoons to 3D anatomical representations. In all methods reviewed, the main challenge has been the representation of animal motion in such a way as to be easily understood and controlled. Even if humans are not overly familiar with the details of quadruped motion, anomalies in computer generated simulations can still be easily identified.

The most accessible sources of information today come from video footage. As we discussed, the extraction of motion from video is a very active topic in computer vision and has also proved very useful in computer animation. Animating characters through the use of skeletons is the most intuitive approach, but it is difficult to create realistic surface deformations using such methods. Even for very natural skeletal motion, the skin of a character can look incorrect if the angles of joints are too big or small, thus detracting from the overall realism of the motion. To counteract this, mesh animation methods have been developed, which in effect use the idea of a skeleton in order to position the vertices of the skin in the right place. However, it is harder to animate most animals in this way, as joints and bones are not explicitly represented. Hybrid methods involve the anatomically based simulation of animals. However, this is difficult to implement in real time, especially if the aim is to create a large herd of animals. The control of the character is ultimately what enables them to be moved. With improved interfaces and effective representations of animal motion, it should be possible to create realistic movements very quickly.

Physically-based models are particularly promising because they enable the generation of not only a single, but a full family of plausible motions based on some a priori knowledge of the animal's morphology and on the way it uses its body to locomote. Requiring adaptive control mechanisms for controlling the gait while maintaining balance, these techniques are still very challenging to implement. While the generation of a variety of motions on different terrains has already been achieved, there is still room for research in the automatic generation of locomotion controllers for different animals, for instance based on the passive dynamics of their body [KRFC07].

By facilitating the task of creating compelling characters, animators can spend more time directing scenes rather than tediously creating the movements for each individual character along with any resulting secondary motions. Other challenges involve creating different motions for similar characters, so that crowds or herds look like they are made up of individuals rather than clones. Parameterising motion is a good way of achieving this goal, whereby different stylisations can be added to gaits. In Australia(20^{th} Century-Fox Film Corporation) and The Chronicles of Narnia (Disney), motion graphs are used to animate herds and crowds of characters, motion clips are put together depending what state a character is in and which clip is the next best fit. This helps to create variation and autonomous behaviour of characters improving the overall simulation.

Chapter 3

Anatomy and Zootomy

Knowledge of anatomy and zootomy leads to a better understanding of how animals move. Most movement comes from the contraction and extension of muscles, which in turn move the bones that support the body. In computer animation the muscles are not generally present as it is sufficient to attach the mesh of the 3D model to a skeleton, which is then used to position the 3D model into various poses. Available software can interpolate between these positions to fill in the gaps in the animation. The quality of the animation, realism and its fluidity depends on the knowledge and skill of the animator, director and artist. Horvath et al. show that even today many museums, taxidermists, anatomists and painters illustrate the postures of walking quadrupeds incorrectly [HCN+09]. Of the 307 walking illustrations that were collected, almost half were wrong. This suggests that we do not understand the motion of animals well enough to be able to tell when they are positioned incorrectly and therefore it is difficult to say whether we can picture the *movement* of animals correctly. This problem is easier to solve for walking bipeds: left leg followed by right leg, and the movement is repeated. Similarly for the fast gaits such as running, left, right, left, right. Given a quadruped, it becomes more difficult to visualise the positions, order and timings of all the limbs. This is made more complicated by the fact that quadruped gaits change from slow symmetrical gaits to fast asymmetrical gaits. In this chapter we provide a concise overview of different animal skeletons with the aim of achieving a better understanding of their motions.

3.1 Anatomy of Humans and Animals

Surprisingly, the anatomy of humans and animals is almost the same, even down to the finger details. The biggest differences, we find, are in the bone proportions. To understand

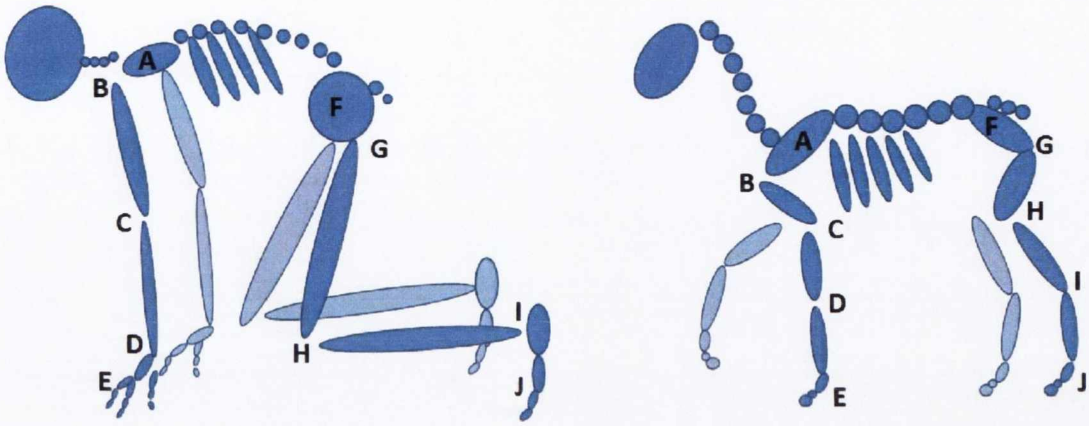

Figure 3.1: A direct comparison of the human (left) and horse (right) skeleton. **A:** shoulder blade, **B:** shoulder, **C:** elbow, **D:** wrist, **E:** finger joints, **F:** hip bone, **G:** hip, **H:** knee, **I:** ankle, **J:** toe joints. Images based on the work of Szunyoghy and Feher [SF06]

the skeleton of a quadruped easily, it is conceptually simplest to consider it as a human on all fours, as shown in Fig. 3.1. We can see that the front legs of the quadruped are the same as the arms on the human - there are shoulders, elbows, wrists and finger knuckles, and the back legs are the same as the human legs with hips, knees, ankles and toes. In this section we will outline some of the main differences between quadrupeds, with our main focus on a particular class of quadruped called *ungulates* (see Sec.3.1.1) as we used them in our perceptual experiments described in Chapter 5 and 6.

3.1.1 Ungulates

Animals such as horses, cattle, zebras, sheep, antelope and pigs fall into the group of animals known as ungulates. An ungulate skeleton walks on the tips of its fingers. As the size of an animal increases, the number of bones in the legs is reduced. For example pigs have numerous bones for the metacarpals - as do humans for the palms, whereas in larger animals, such as the horse, these bones have fused into one. Bigger bones are stronger and are required to support large animals. The joints of the ungulates are also restricted to hinge joints, so they cannot move their wrists as humans do. Restricting movement to one axis prevents damage to the bones. Forces which cause a shearing effect on a bone are the most common cause of broken bones, so by preventing twisting action shearing is avoided.

While larger animals have skeletons similar to their smaller relatives their bodies have developed differently in order to suit the various lifestyles they lead. Horses and pigs are both quadrupeds but horses have long legs, possibly because over time they have evolved to

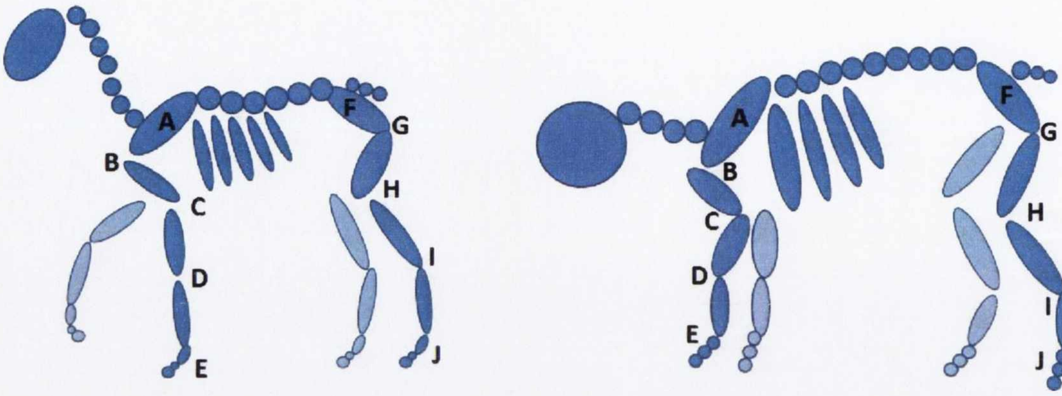

Figure 3.2: On the left is a skeleton of the horse, on the right of a pig. We can see that both animals walk on the tips of their "fingers" and "toes". Images based on the work of Szunyoghy and Feher [SF06]

escape predators [SH05], while pigs have short legs possibly as they have been reared as a food source for humans and hence rarely have to gallop. In addition, the natural habitat for pigs or wild boar is woodland, thus with their short legs pigs can hide in the dense bush to avoid being eaten by predators [SH05, kHD03]. It is not only the legs that are different, Fig. 3.2 shows that the shape of the spine also changes. Horses have a very distinctive dip in the spine, while the pig has a very curved spine which starts off low at the head and curves upwards towards the end of the body. Cows on the other hand have a very flat spine and a sheep's spine has a slight curve. It is easy to spot these differences when looking at the point-light walkers of these animals. The speed at which an animal moves is defined by the distance it covers in a certain amount of time. This distance is determined by the stride length and, in order for this value to be high, the length of the stride has to be maximised. Two different factors affect the stride length: 1) the physical make-up of the animal and 2) the way the animal uses its body during the motion. We will first look at how the physique helps to increase the stride length. In order to move the limbs quickly, the muscles are joined to the bone close to the joint, thus allowing the muscle to move the joints at a large angle quickly. As ungulates stand on the tips of their limbs, the length of their legs is thereby extended, thus increasing the stride length. The shoulder blade is found on the side of the body (as opposed to humans that have the shoulder blades across the body) and this not only adds to the length of the leg but, as it can be used for movement as well, it increases the angle through which the whole limb can move and results in a longer stride.

Hinge joints ensure that all movement is in a forward direction, so the velocity of the hoof can be calculated by adding all the velocities of the connected bones (hip, thigh, calf, carpal

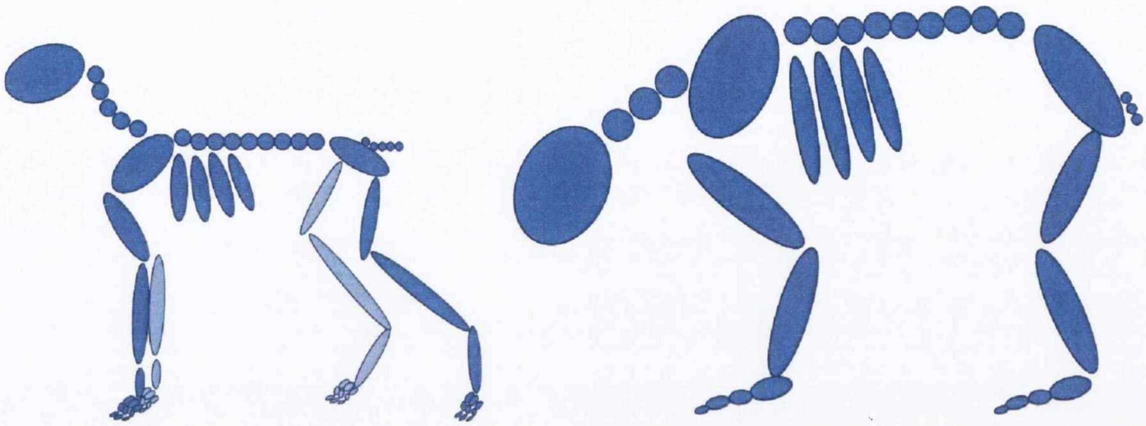

Figure 3.3: On the left is a simple representation of a dog skeleton; the paws are the same as human fingers and toes. On the right is a simple representation of a bear skeleton; they step down on the whole palm and foot.

and metacarpals). As the momentum increases, suspension phases develop during the gallop, which further increases the overall length of the stride.

3.1.2 Carnivores

Carnivores, walk on what could be considered their fingers and toes (digitigrade animals). Bigger animals such as bears step onto the whole palm or foot as humans would when on all fours (plantigrade animals), see Fig. 3.3. By retaining their digits, carnivores are able to better manipulate their food. A reduced clavicle allows them to use their shoulder blade as part of the limb, lengthening their stride when they need to run fast. To increase the stride further carnivores are able to flex their spine. Extending the spine when the hind legs are on the ground, increases the length of the body in the forward direction (the planted legs prevent any backward movement). Likewise, deceleration is prevented when the front legs are planted and the spine is flexed. All this helps to maximise the distance covered during one stride. Large animals keep their spine stiff in order to minimise oscillations that would otherwise keep them unstable and render their legs incapable of supporting such weight. The structure, motion and evolution of animals has been studied in great detail by Hildebrand and Goslow [HG01].

3.2 Animal Motion

Animal gaits can be classified as either symmetric (walk, trot and pace) or asymmetric (transverse and rotary gallop). During a slow walk, three legs are on the ground at any one

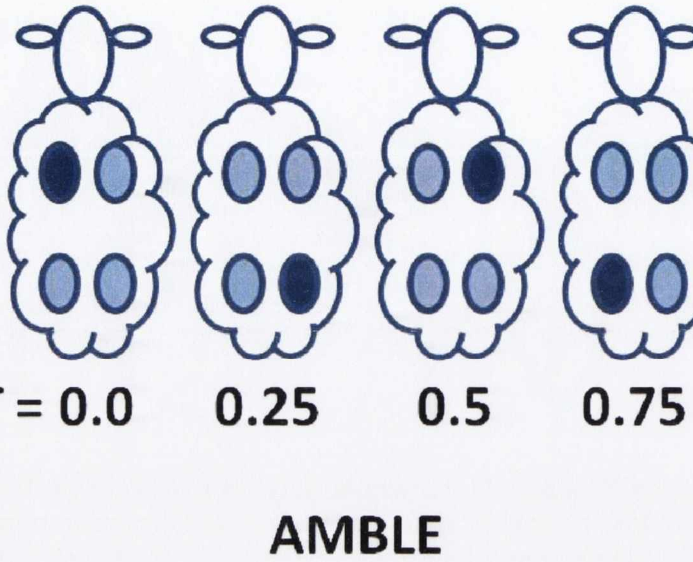

Figure 3.4: A top view of a sheep showing the different phases of a slow quadruped walk. Front left, back right, front right, back left. Image based on work of Alexander [Ale84]

time and the timing of the cycle can be divided into four phases. For example, the order of motion for a slow walk (amble) can be: left front, right back, right front, left back, as illustrated in Fig. 3.4. For the quicker gaits the time the legs spend in contact with the ground is reduced. During the pace the legs on the same side of the body move at the same time, for example the legs on the left are moved together and then the legs on the right. Animals with relatively long legs are natural pacers, such as camels or big dogs whereas horses are usually taught how to pace. For the trot, the legs on the diagonal are moved together as is more natural for most quadrupeds, see Fig. 3.5.

During a gallop gait there is a suspension phase where all the legs are off the ground. In horses this occurs when all the legs are underneath the body. Dogs, such as greyhounds, have two suspension phases, one when they push off with the front legs, and one when they push off with the back legs. If the front legs are looked at as a pair, then the leg that is second to touch the ground is the leading leg as it is used to push off for the next suspension phase. Similarly for the back legs, the second leg to land is the leading leg. If the leading leg of the front legs is the same as the hind one (e.g., both are left) it is called a transverse gallop and is generally performed by horses. If the legs are different (e.g., one is left and the other right) it is called rotary gallop and is performed by dogs and cheetahs, Fig. 3.6.

Hildebrand tries to classify the various gaits that horses are able to perform [Hil65]. He includes different breeds of horses in his study to check for variation and concentrates on the

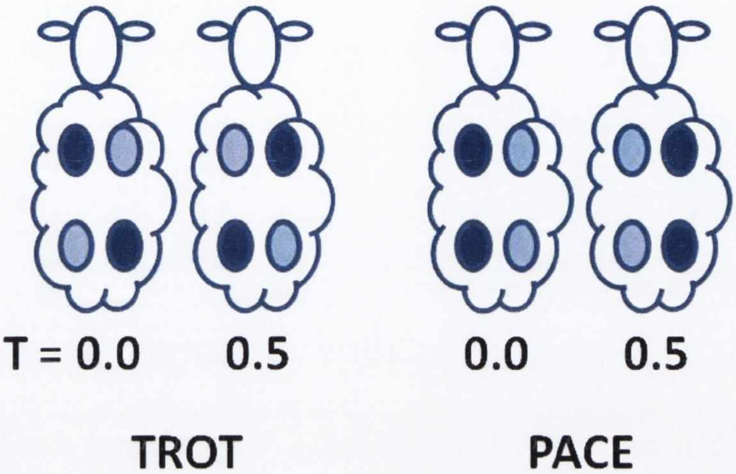

Figure 3.5: The phases of a trot and a pace gait. The trot involves moving the diagonal legs at the same time, and the pace involves moving the legs on the same side of the body at the same time. Image based on work of Alexander [Ale84]

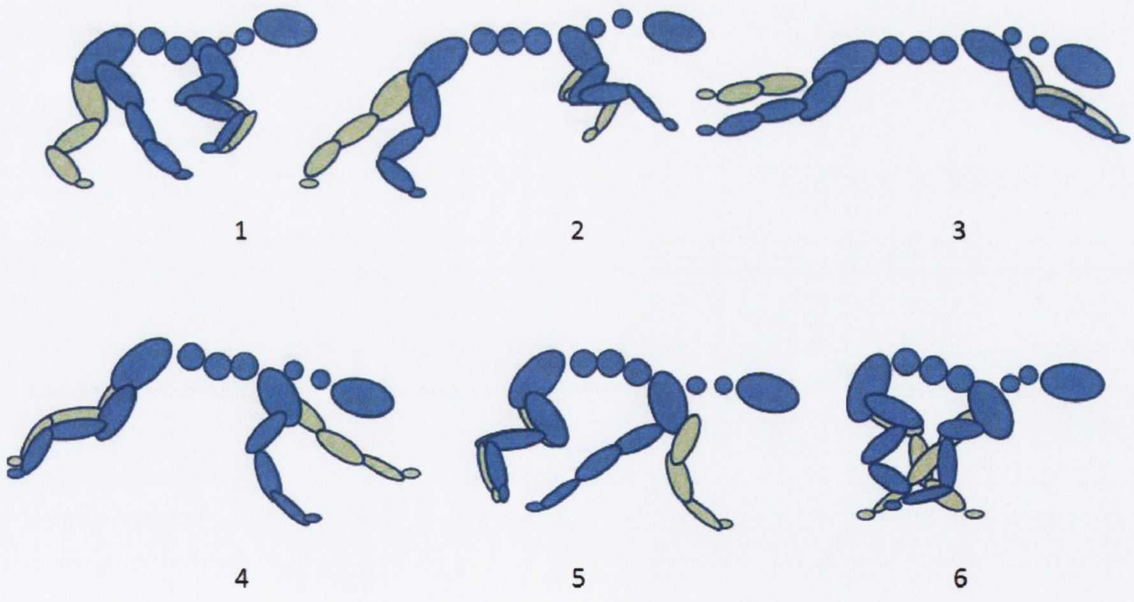

Figure 3.6: A an example of a rotary gallop. 1 - 3 pushing off with the back right for a suspension phase, 4 - 6 pushing off on the left front for the second suspension phase. Image based on photographs from [MB57].

pace and trot motions. The variation between the gaits, for example the timing of the footfalls, mostly occurs between the different species of a horse rather than between individual horses. This study included 68 horses of many breeds, such as Belgian, Clydesdale, Arabian, Peruvian Paso, Thoroughbred, Quater Horse, Standardbred, American Saddle Horse, Tennessee Walking Horse, Roadster, Hackney Pony and Harness Pony. The study also found that there is no considerable difference in motion between the colts and the adult horses. The movement of animals has also been studied in great detail by Alexander [Ale68] (as mentioned in the previous chapter) where it is described by a Froude number: u^2/gh, where u is the speed the animal is travelling at, g is the gravitational pull, and h is the height distance from the hip to the ground. If the Froude number is the same for different animals, then they are travelling with the same gait. For example for a walk gait, a cat travelling at 0.5m/s with a hip height of 0.22m has a Froude number of 0.1. A camel which is a much bigger animal, travelling at 1.3m/s with a hip height of 1.7m also gives a Froude number of 0.1. It makes sense that a camel with longer legs covers a much greater distance when it walks compared to the cat. We can also deduce from this that, for a cat to keep up with the camel, it would need to increase its speed and by doing so the Froude number goes up indicating that it is using a faster trotting gait. This classification of gaits comes in useful when you need to relate the motion of one animal to another, as in the case of a scene where different sized animals are seen to move together.

3.3 Discussion

From this we can see that quadrupeds move in complex ways. With high speed cameras and videos we can replay, slow down and stop the motion entirely in order to study the order and timings of the footfalls, as well as the secondary motion that includes the head and tail. However, even with a thorough understanding of the motion of animals, it is difficult to model and recreate the movements using robotics and mathematical models. It is possible to use a physics engine to model the locomotion system but there are many factors to take into account e.g., forces and torques at every joint, spring-like behaviour due to shock absorption, friction with the ground, ensuring that the feet stay planted on the surface, the motion of the upper body and the motion of the head and neck during the different gaits. The resulting motions can be stiff and slow to calculate thus not look very realistic at all.

Aesthetically pleasing animations are difficult to produce and require many man hours to perfect, as they are usually tweaked by hand. In the animation industry today virtual animals are usually animated with fast complex motions, some of which may not even be commonly used by real animals. In addition, consumer expectations are continuously rising e.g., in The

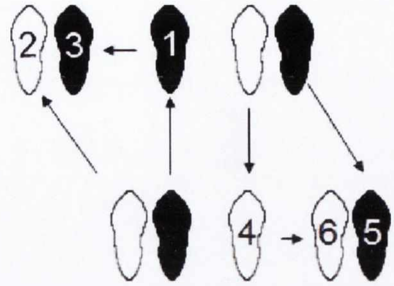

Figure 3.7: The box step for the Waltz. Left foot - white, right foot - black.

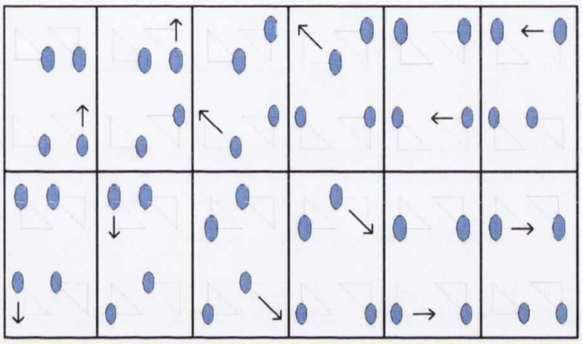

Figure 3.8: The twelve-phase cycle of the sheep box step for the Waltz, starting in the top left corner moving towards the right.

Golden Compass (New Line Cinema) the animals are made look very real but with their anthropomorphic qualities they find themselves in fast paced action scenes - chasing and fighting, so their actions are exaggerated. However, no matter how complicated the positions, a model will always be guided by the limits of its corresponding physique. Therefore, understanding these structures is vital in order to produce good quality animation.

To see if it is possible to apply human motions to quadruped animals while staying true to their natural movements, we adapted the footsteps associated with a Waltz dance to suit a 3D sheep model while working in collaboration with Ken Perlin (Media Research Lab, New York University). A quadruped's natural walk cycle is adapted to create a twelve phase step that imitates the Waltz box step, see Figs. 3.7 and 3.8. Using Inverse Kinematic chains the legs of the sheep are animated and the body is positioned so that it moves towards whichever hoof is reaching out farthest to keep the overall movement looking "natural". A method produces the choreography for the dancing sheep by specifying the orientation and position of each dancer. We were able to create a new "gait" for the sheep characters while staying true to the physiology of real sheep and the timing of their more natural walking gait. The final result can be seen in Fig. 3.9.

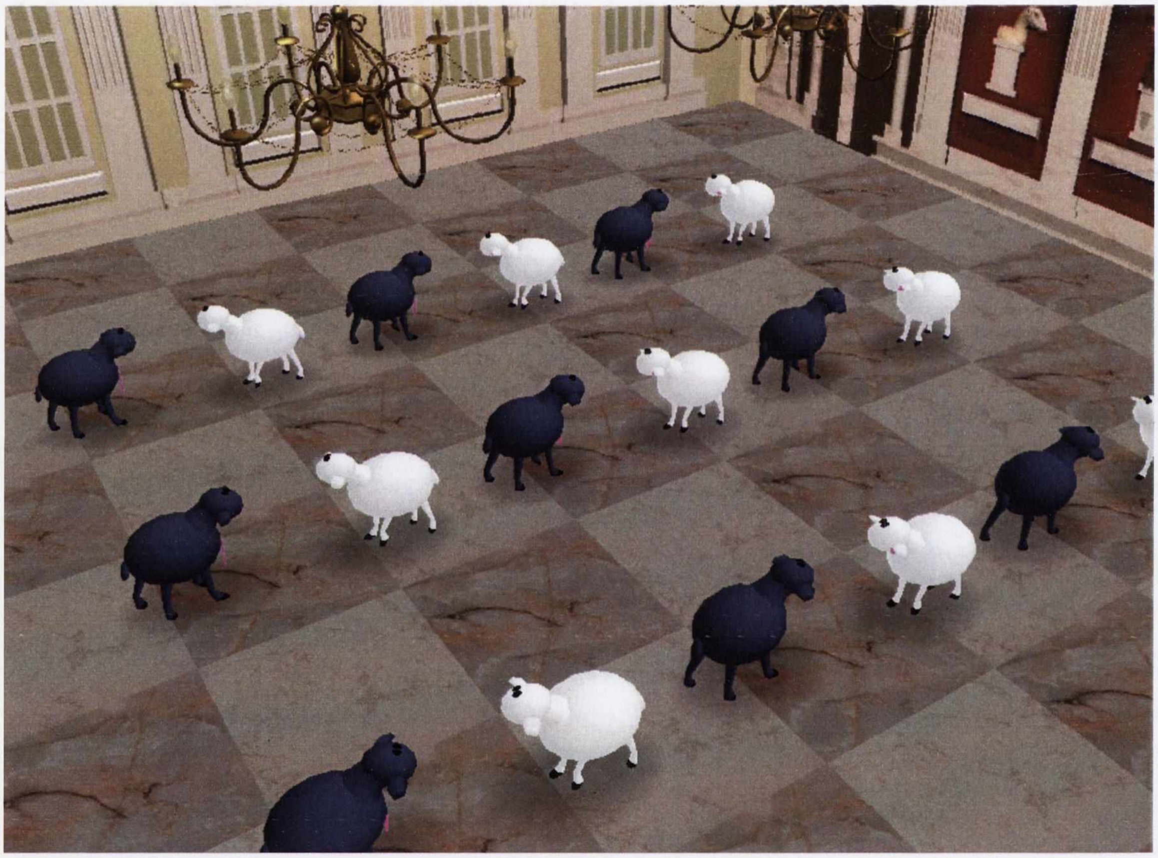

Figure 3.9: A screen shot of the virtual sheep dancing the Waltz in a virtual ballroom.

Chapter 4

Quadrupedal Animation

How do we achieve good quality animations? Johnson and Thomas provide a comprehensive overview of the principles of traditional animation used during the 20^{th} century [TJ95]. There are eleven basic principles: squash and stretch, timing, anticipation, staging, follow through and staging, straight ahead action and pose-to-pose action, slow in and slow out, arcs, exaggeration, secondary action and appeal. Although in recent times the trend has been towards increasingly realistic animations rather than cartoon-like and stylised motions, these principles can nevertheless still be relevant, particularly when preforming facial or skin animation and these principles can still be useful for artists, as explained by Lasseter [Las87]. The animals in Chronicles of Narnia (Disney) and The Golden Compass (New Line Cinema) are designed and animated to such a high standard that they almost seem real. This is very difficult to achieve and many skilled artists are involved, who spend hours studying the different motions of animals and how they look and behave. The companies that work on such expensive projects build their own in-house tools that help them with the process (e.g., in the film Australia, Rising Sun Pictures used their *Venom* infrastructure to create the cattle herds). However, this option is not always practical and many have therefore tried to develop procedural techniques to animate quadrupedal animals, as discussed in Chapter 2.

In this chapter we will discuss the different techniques we used for the animation of our chosen animals and the advantages and disadvantages of each. We will first look at where relevant information about these animals can be found and then how it can be best exploited for the purposes of realistic quadruped animation. Then, we will describe the different approaches we took to achieve realistic animations that we used in experiments to study the motion of real animals.

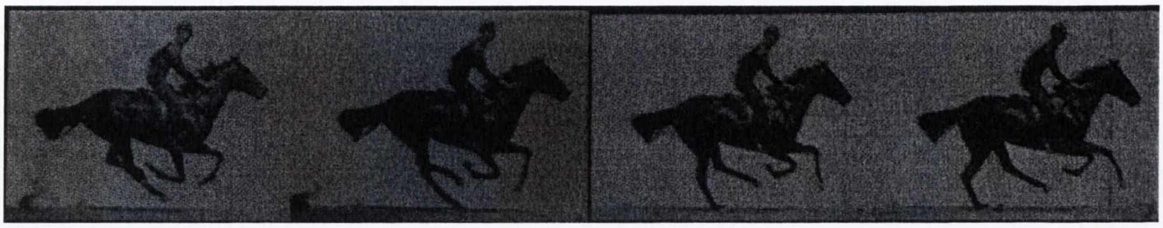

Figure 4.1: These images represent a few frames of a gallop gait. For a fast paced gait, it is particularly difficult to tell the exact movement of the head from the frames but this type of motion is vital for realistic portrayal of animals [MB57].

4.1 Data Collection

As mentioned before, in order to animate any subject there has to be an understanding of how that subject moves. This information can be obtained from different sources, such as live action, a cartoon strip, a selection of photographs or a video of someone or something preforming the motion. Notwithstanding this wide choice it is difficult to acquire appropriate and good quality data. In particular, it is difficult to gather animal motion that can be directly used for animation. When animating human characters it is relatively easy to ask someone to act out the poses or motions while they are being tracked or recorded. There can be numerous takes after which the actor can be instructed about where to stand and how to move and where and when to go and all this can be achieved in a relatively short amount of time. This does not necessarily apply to working with animals. Many hours of video footage usually provides a few seconds of data that can be used. Burt presents the history of animals in films and the technical challenges overcome in the film industry throughout the twentieth century [Bur03]. Below we discuss the more popular techniques used for gathering data on animal motion and mention some of the drawbacks found in each one.

4.1.1 Photographs

Muybridge's photographs are often used as a source of information when the motion of animals needs to be studied. In Fig. 4.1 we can see a few frames of a horse's gallop gait that Muybridge was able to capture with his high-speed cameras. The different leg positions are relatively easy to see but the head motion looks almost unchanged from picture to picture. This shows how difficult it is to replicate the real motion exactly from still images without the upper body looking relatively stiff. It is also difficult to see any details on the legs due to the low resolution, and it is challenging to decipher which leg is the left and which is the right, in particular for the back legs, mainly due to the lack of colour and poor lighting.

Today however it is possible to acquire very detailed photographs of animals as the technology has developed and high definition has become standard. Some problems do remain however, including the fact that it is difficult to deduce how fast the animal is moving from a series of static poses. It would be possible to set up a location where animals are walked in front of a grid, as Muybridge had done, however we have not found a modern equivalent to Muybridge's work - it has not been possible to locate any collections of good quality coloured photographs that could be used for the study of animal motion.

4.1.2 Videos

Videos have an advantage over photographs in that they capture the speed at which the animal's limbs are moving. If the camera is set up correctly it is possible to determine what distance the animal covered in one cycle of motion or for the whole duration of the motion captured. Video cameras can also capture the motion over a longer period of time, so it is possible to leave the camera recording the animal as it performs many different motions and later to edit the video and keep only what is necessary. Furthermore, an animator can play a video alongside the animation they are creating, thereby making it easier to compare between the two to ensure that the motions have been reproduced correctly.

The disadvantages of videos are some of the same ones as found in photographs, in that they only contain 2D data, the quality of the image can be poor or blurry if the frame rate is low, the animals can be difficult to see if the lighting is constantly changing - for example, due to changes in sunlight and shade when filming outdoors. Occlusion is also a big problem, particularly when one wishes to film animals in their natural environment. Therefore, it is of utmost importance that all recording sessions are carefully planned.

In our work we painted circles on the joints of animals we wished to film, on one side of the body only (we completed ethical forms for this, see Appendix A.1). We let the animals walk around an enclosure until we found a clip in which the animal walked parallel to the camera plane and all the markers were visible. We planned to use commercial software to extract the marked points automatically but this proved infeasible for the video corpus we had gathered. The commercial software was unable to extract the markers in the video as the video was either too dark or there was too much contrast. Furthermore, as the animals were moving, the legs would cross each other so it was not possible to track the markers. We therefore built our own software that was able to keep track of all the points. Our software allows the user to pick out each marker to be tracked for every frame of the captured motion. This results in a simple text file that contains all the x and y positions for all the markers on the animals' bodies, see Fig. 4.3. We were then able to use this text file to animate a set of discs in 3ds Max® with the same number of frames as the original motion.

Figure 4.2: A sequence of images taken from a recording of a sheep's walk cycle. The sheep vacillated between standing and galloping/leaping without performing the normal walking gait we wished to capture.

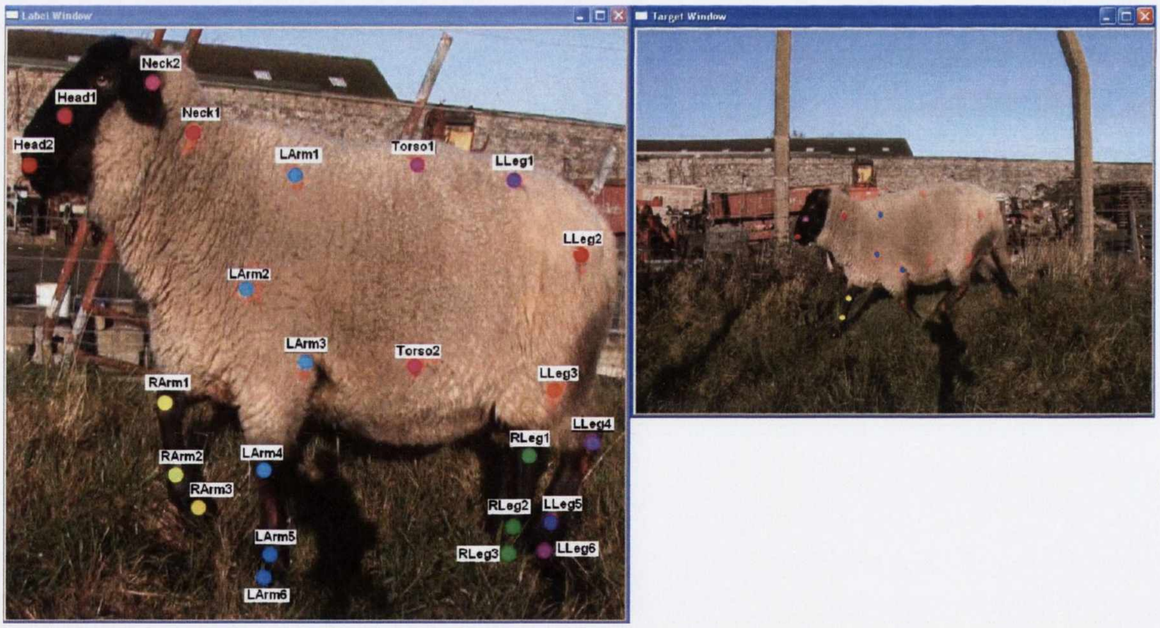

Figure 4.3: The image on the left shows the labels of all the markers to be tracked. The image on the right shows the current frame a user is working on. A user first clicks on the marker they wish to track in the image on the left, and then they click on the corresponding marker on the image on the right. In this way, the program knows which marker corresponds to which co-ordinates for every frame of the animation.

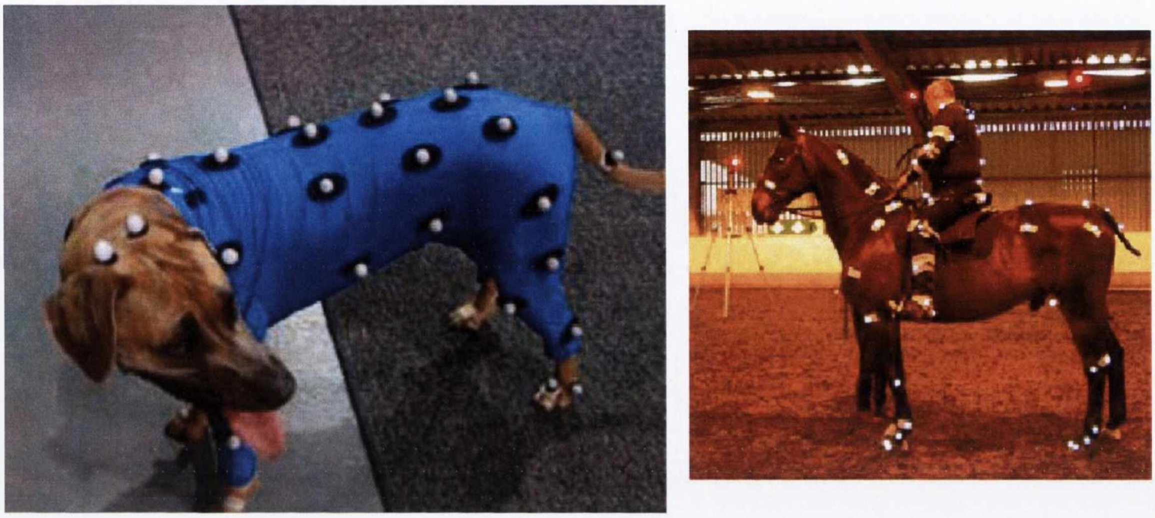

Figure 4.4: Customised equipment used for the motion capture of dogs and horses. Images from Peak Motus [Gil02] and Equine Mechanics [Law].

4.1.3 Motion Capture

Motion capture systems can be adapted for use with animals and there are specialised companies that use customised equipment, which they attach to animals such as horses and dogs in order to capture their movements (e.g., Equine Mechanics [Law], Kinetic Impulse, Horse Locomotion [Wid], Peak Motus [Gil02]) see Fig. 4.4. Thus full 3D motion is captured and applied to 3D characters, who then behave and move like the real animal. Horses and dogs are commonly used for motion capture as they are relatively tame and trained so it is easier to work with them. Interestingly, cats remove markers that are attached to them, thus proving difficult to capture and so they are still mostly hand-animated [1]. Placing animals in environments with expensive equipment may be less than an ideal situation as they can behave unpredictably, while the motion capture systems usually have a relatively small range which may not be suitable to capture such detailed movement. Dogs are small animals in comparison to humans, so markers would be small and difficult to see. However, the area required to capture a dog galloping or performing different sorts of actions has to be quite large, so the resulting compromises made severely restrict the range of motions that can be captured.

Finally, with respect to the capture of motion from animals, no matter what equipment is used, animals will always be difficult to work with. Recording sessions last for a long time,

[1] We learnt about this from an interview with Hans Rijpkema, senior production software engineer: Rhythm & Hues Studios

even when all that is required to film is a simple walk.

4.2 Preparing a 3D model for animation

In our research we used 3ds Max™ for the animation of the 3D models, as this software provides way of creating hierarchical bone structures that define the skeleton. We also use the built-in IK system for the animation of skeletons. The 3D models we used throughout this project are commercially available [2] [3].

4.2.1 Skeleton

The skeleton controls the movement of the character and as it moves the mesh deforms accordingly. The skeleton can be as detailed or as simple as the user wishes. However, a compromise usually needs to be reached whereby the skeleton is simple enough for ease of control, but detailed enough to capture all the necessary motions of the character. As we saw from Chapter 3, an animal leg consists of: the shoulder blade, upper arm, lower arm, palm and 3 finger joints, which is a total of 7 bones. However, in animals such as horses and cows, the movement of the last 3 bones can be described by just one bone, thereby reducing the number of leg bones to 5. Similarly, depending on the character being animated, the spine can have either many bones e.g., if a lot of flexible spine motions are required, or (as in our case) the animal models only have 2 bones, the neck is reduced to one bone and there is another one for the head, see Fig. 4.5. To transform the bones we created an IK chain between adjoining bones on the skeleton. Each IK-chain controlled the position of a joint and the bones below it (further down the hierarchy).

4.2.2 Skinning

During skinning an animator specifies which vertices on the mesh are controlled by which bone. This can be a particularly difficult job to get right as it is difficult to keep the mesh smooth around the joints. One also needs to trick the viewer into thinking there are muscles under the skin which shape the mesh. All this requires careful positioning of skeleton bones and weighting of the mesh vertices, as unsightly deformations can take away from the quality of the animation.

---

[2] http://www.deespona.com/3denciclopedia/index.html

[3] http://www.turbosquid.com/FullPreview/Index.cfm/ID/305334

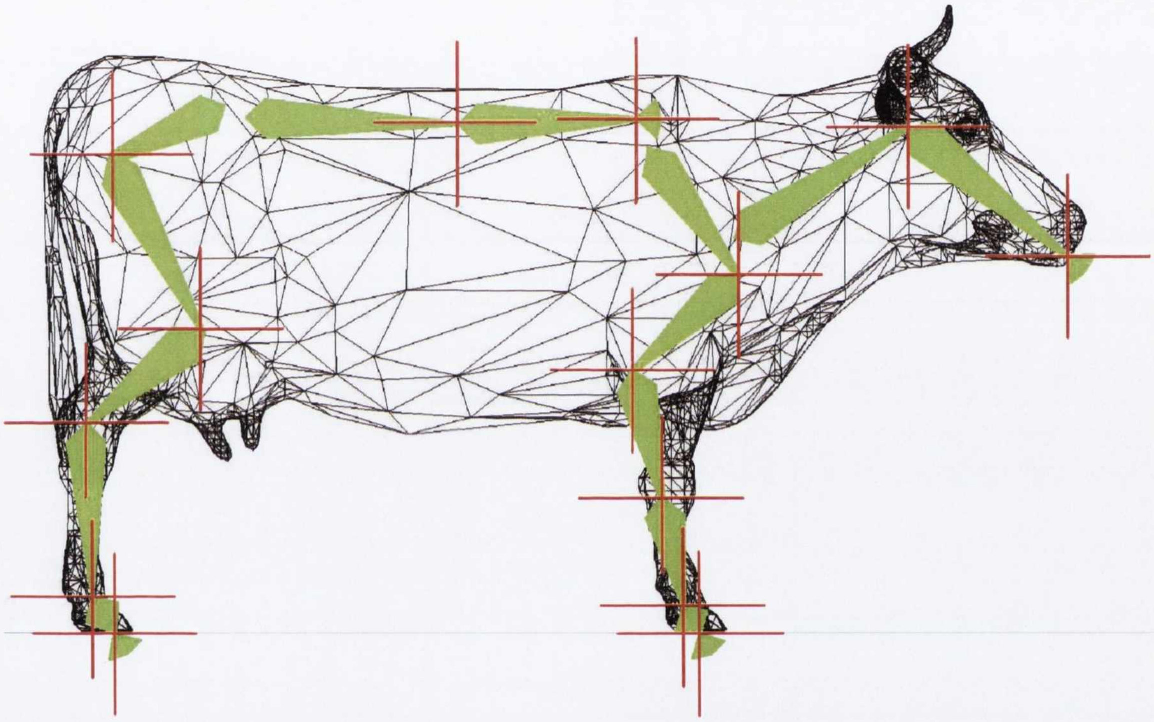

Figure 4.5: The 3D cow model we used and its simplified skeleton. The small green bones found at the end of each chain of bones (some are outside the body) make it easier to calculate the positions for foot, head and upper shoulder placements. The red crosses represent the IK-chains that move the skeleton.

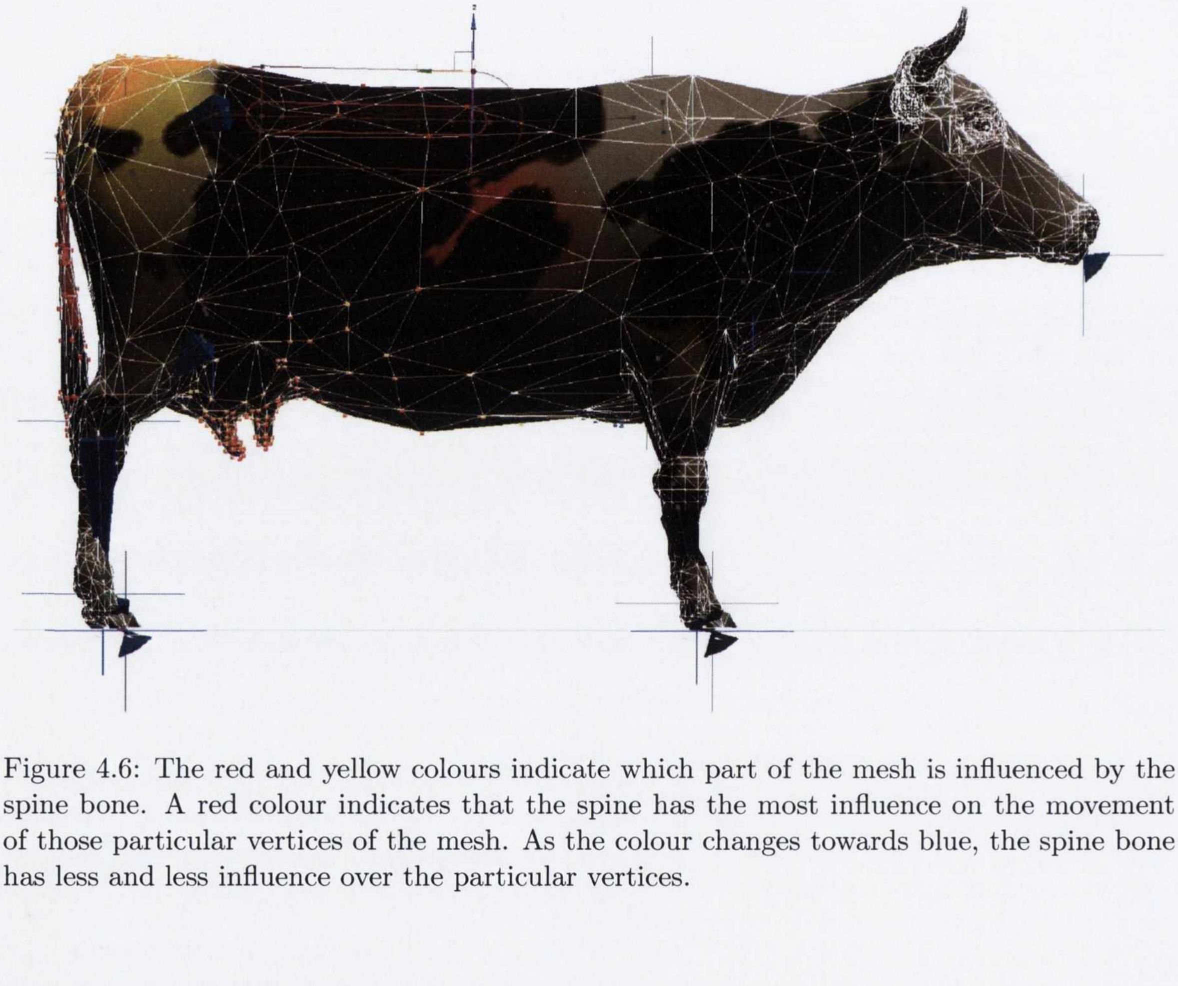

Figure 4.6: The red and yellow colours indicate which part of the mesh is influenced by the spine bone. A red colour indicates that the spine has the most influence on the movement of those particular vertices of the mesh. As the colour changes towards blue, the spine bone has less and less influence over the particular vertices.

4.3 Animation

There are different techniques used for animating characters depending on which effect needs to be achieved. For detailed and very precise realistic animations, hand animation is often used along with some motion captured data. This approach is common in the film industry, however it is very time consuming and expensive. In cases where high detail is not so important procedural animation is an option as the motion of the character is restricted due a smaller set of control points. We already described in Chapter 2 the existing techniques used to date for the animation of quadrupedal animals. Here we will discuss the different approaches we took to achieve our goals.

4.3.1 Key-framing

Key-framed animation involves an animator specifying different poses of a character. Each key-frame stores the position and orientation of the corresponding bones or controllers used. The software used (3ds Max® in our case) interpolates between the key frames, resulting in the continuous motion of a character. The animation can be played back and corrected many times over.

Hand

Hand animation is usually the easiest method to start with as there is built-in support in most 3D packages and it gives the animator a great amount of control over all aspects of the animation. However, it is very difficult for a non-expert to create good animations and even more difficult to create animations that are believable or convincing. Our first aim was to have a galloping dog shepherd a flock of sheep, so we first animated a dog and a sheep model using key-framed hand animation. To animate the dog model we used diagrams depicting the gallop found in [CDNDMN85], and for the sheep we used video footage that we could stop at intervals in order to re-position the bones correctly for each key-frame. These models were exported as impostors and used in a in-house crowd and urban simulation system (Virtual Dublin) [HO03] that demonstrated flocking and shepherding algorithms (as described in Chapter 6).

In order to add variation to the behaviour of the flock as a whole, the sheep model was also exported with a grazing animation. In the Virtual Dublin system, the sheep were able to pick randomly whether to walk or eat, which helped to create autonomous sheep behaviour. As we were using impostors for our simulation of shepherding and flocking we were able to add fur to the skin of the animals using the fur modifiers in 3ds Max™.

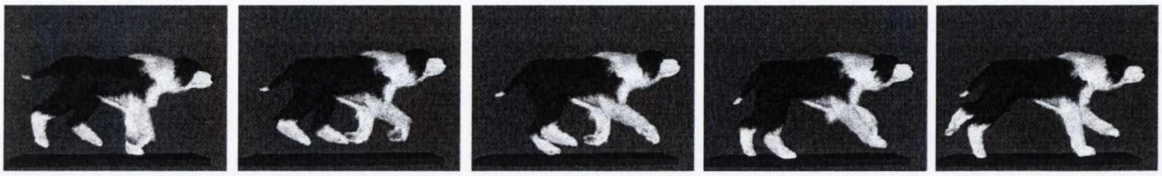

Figure 4.7: A few key-frames of galloping animation based on images found in [CDNDMN85]. As we were using a dog impostor in the real-time system we were able to add a fur modifier to the mesh in 3ds Max to make it more realistic before exporting it.

Scripting

A script can be used for many different purposes such as animation, behavioral simulation, exporting of animation from 3D software packages, to name but a few. In character animation a script can help keep movement smooth by creating key frames at regular intervals and by ensuring the motion is along a smooth trajectory. If the character is already animated, or has a few animations ready, a script can then choose to play one of the animations depending on how the character should act within a scene. For example, if a character needs to run, the script can load that animation and play it in a loop. However, scripting does not always provide a solution that blends between different animations and so there may be no transition between a standing and a running motion. It is also very difficult to allow the character to interact with the environment they are in when their animations are pre-generated. There is a large amount of research done in this area for human characters [ZS09, HG07, Wel09].

2D motion capture and re-targeting

We extracted the motion of the animals from the videos of the horse, cow, sheep and pig we recorded as described in Section 4.1.2. The 2D data was then represented by a set of discs in 3ds Max, shown by the green circles in Fig. 4.8. We scaled the distances between the green circles so that the motion would fit the size of the bones of the skeleton for each 3D model, and the new positions are shown as the red squares. In order to capture the real distance covered and also to avoid any foot skating later on in the animation, we kept the camera still during filming. As we were limited by the position of the camera and it was difficult to capture the motion for one full cycle, for example in the case of the horse, the head entered the frame first, followed by the rest of the body, then the head left the frame first. Unfortunately, due the tight space in which we filmed the full body was not visible in the frame for a complete cycle of each motion, see Fig. 4.9. However, we were able to use the data we had to calculate the position of the body for each frame of the cycle and use this to

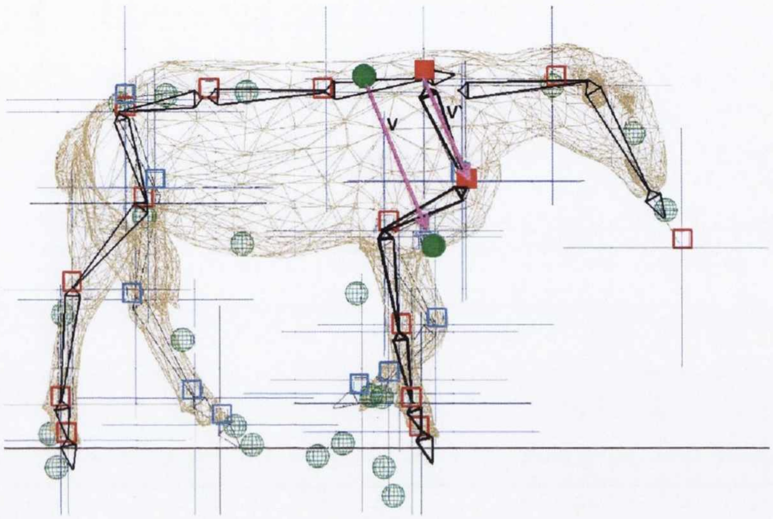

Figure 4.8: The green circles represent the original motion of the point-light walkers, while the red squares represent the same motion that has been scaled to fit the model. The blue crosses are positioned according to the motion of the red squares. The bones deform according to the blue crosses (IK-chains).

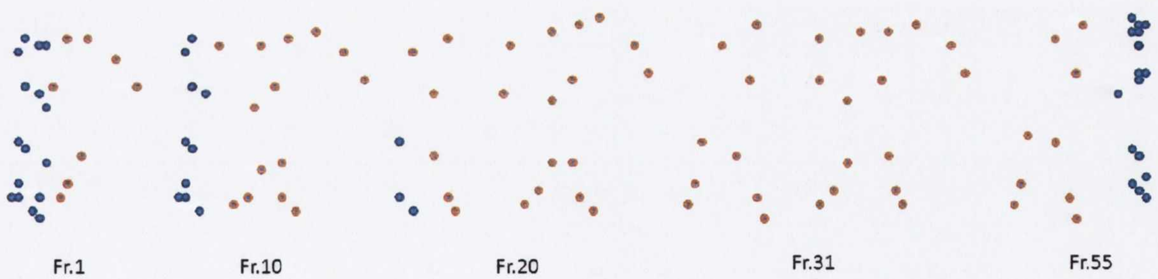

Figure 4.9: The horse enters the field of view of the camera from the left and then leaves again. The orange circles show the points that are moving and the blue that are not moving. The whole horse is only visible in the whole field of view for only a few frames, but not for the whole cycle.

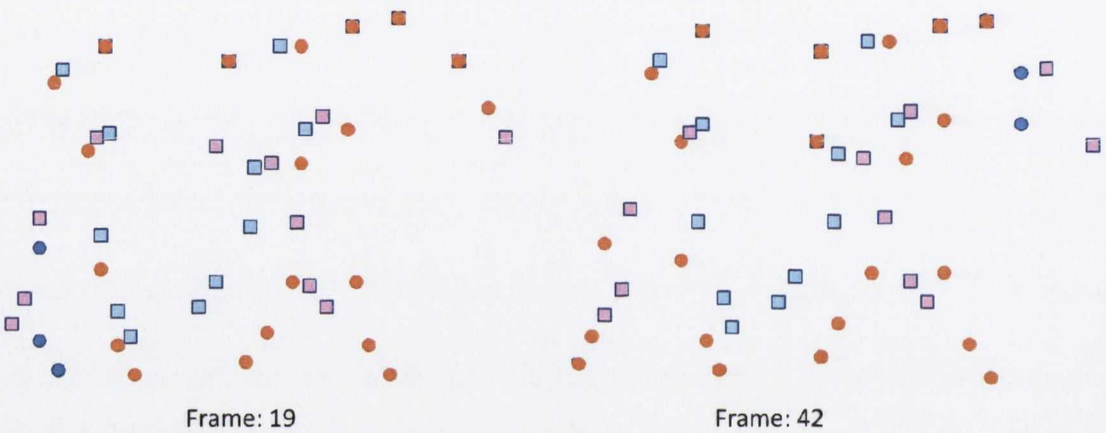

Figure 4.10: From the motion we captured we could see that Frame 19 and Frame 42 are the same during the horse trot cycle. The blue circles represent the original tracked motion where the points are not animated and the orange circles represent the original motion that is animated. In Frame 19 the key-frames are missing for lower right leg and in Frame 42 key-frames are missing for the head. The squares represent the scaled motion based on the size of the bones for that model. To calculate the missing key frames in Frame 19 we use the data from Frame 42 to calculate the positions for the lower right leg - this can be seen from the squares that represent the lower right leg, while similarly we use the information from Frame 19 to calculate the missing key-frames for the head in Frame 42 - this can also be seen from the squares.

create any extra frames we needed in order to make the animation longer, see Fig. 4.10.

To scale the motion, we calculated a vector between a pair of point-light positions, then we normalised and scaled this to match the size of the bones. The starting point for the calculation of the vectors was the head for the first few frames of the animation as it was the first to enter the field of view. For the last few frames of the animation we had to use the base of the spine as the starting point, along with data from previous frames to calculate the missing positions for the head.

Both of the gaits we captured are symmetrical, therefore we were able to use the motion of the right side of the body to animate the left by calculating the full length of a cycle and then offsetting the motion for the left side by half the cycle length. As the animals were walking in a straight line, the 2D data was acceptable for the recreation of the gait on the 3D models, and we used the width of the model to calculate where the scaled points should go. These points were then used to animate the IK-chains of the model's skeleton which in turn animated the 3D model. It is possible that we lost information about possible movement of the spine around the y-axis. However, we believe that this information is only minor in the overall movement that was captured. As described in Chapter 3, we know that large ungulates keep their upper bodies relatively inert. Also point-light walkers moving across the screen only represent motion in 2D, therefore any movement around the y-axis would be lost were it captured.

4.3.2 Exporting

During the course of our work we used two different real-time systems; first we exported the hand-animated dog and sheep models as impostors and used them in the Virtual Dublin system for real-time simulations of shepherding and flocking behaviours (described in more detail in Chapter 6) [HO03]. Then we exported the horse, cow and sheep models (animated using re-targeting) as geometric models along with their animations and textures and used them in a another real-time system based on the Ogre engine[4].

Exporting the Impostors

For the flocking demonstration, the models of the dog and the sheep and their animations were exported as impostors using a customised plug-in and rendered in a similar way to that described in [DHOO05]. The animal impostors were rendered from the 8 camera levels, at 32 different viewpoints, see Fig. 4.11, with hardware assisted colour variation for the sheep. The alpha channel values were used to display the sheep correctly, by using the silhouette of

[4] http://www.ogre3d.org/

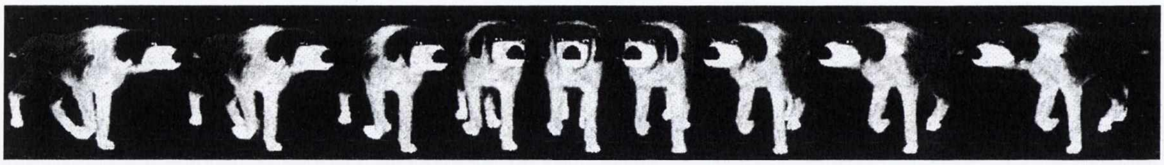

Figure 4.11: A few images of the dog impostor for a single frame of the animation. Since the gallop is not a symmetric gait, we had to generate the images around the whole model for entire duration of the cycle. For other symmetrical gaits, such as the walk, it would be possible to generate images only from one side of the body and mirror them.

the geometric model. While this reduces the "woolliness" of the sheep (we applied the fur modifier to this sheep model), it prevents a black halo (caused by the fuzzy wool combining with a black background) being displayed around the sheep, see Fig. 4.12. The number and size of pictures was large therefore, due to memory restrictions, the number of animations per animal was limited. For this reason, we had two animations for the sheep and one for the dog. In the bottom left corner of Fig. 4.13 we can see the sheep with the grazing animation and in the background with the walking animation.

4.3.3 Procedural Animation

As an alternative to using characters that have already been animated in other software and then replaying their motion in a real-time system, as in the case of various computer games, it is instead possible to animate characters in realtime using constraints, with other information, to move a character around a scene. Many such systems exist where foot placements can either be specified by a user or generated automatically and then either through inverse kinematic or physics based engines or a combination of both it is possible to simulate moving bodies around a scene (e.g., moving the cat in [Tor97], or Havok animation tools including features such as foot-IK modifier, AI, transitions between animation clips, ragdoll controller). As mentioned earlier, there are drawbacks to re-playing an animation in a real-time system and in particular in interactive environments, as there is a high possibility that the exported animation will not be able to react smoothly within its new environment - in particular if the character is an impostor. By creating animation on-the-fly, a unique motion is created each time, thus resulting in a more natural simulation than a simple play-back and the motion will be better equipped to react to objects in the current environment.

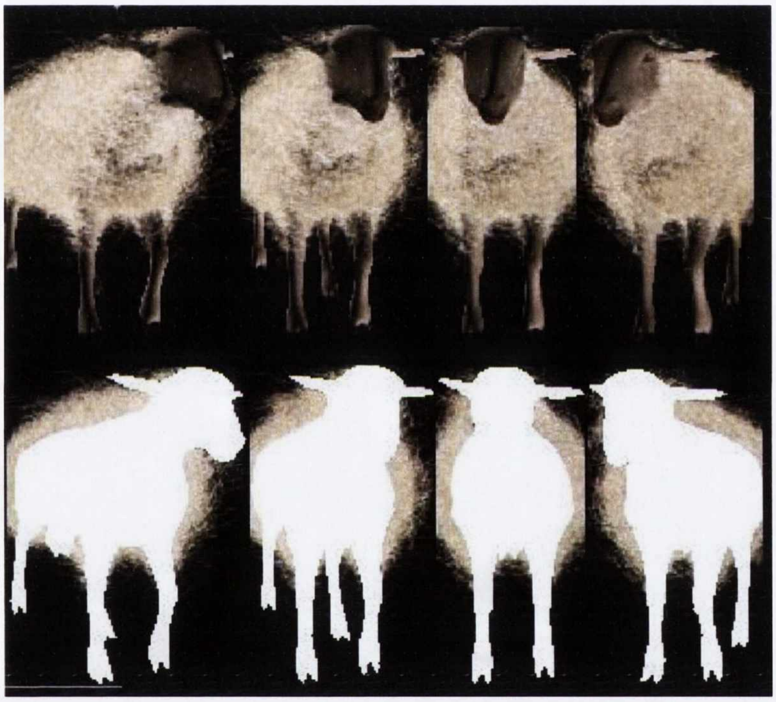

Figure 4.12: A few images of the sheep impostor, bottom row shows the cut-off wool which is how the impostors were displayed in the final simulation.

Figure 4.13: Top left shows the impostor dog chasing the sheep, and the bottom left is a close up of how the sheep look, on the right side we can see a flock of sheep made up of several hundred impostors.

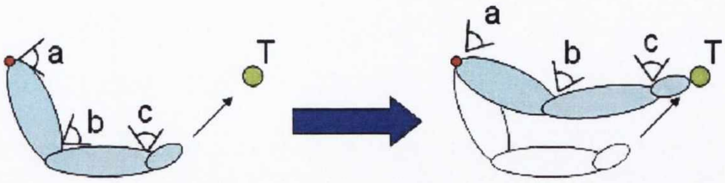

Figure 4.14: Inverse kinematics. Move the end-effector to the target position and then calculate the joint angles for the rest of the bones.

OpenGl

After animating the dog and the sheep by hand, we investigated how to make the system more reactive and scalable, perhaps by building a system that would be able to generate animations for different types of quadrupeds. Using parameters, it would be possible to specify the size of the legs and how long their stride should be, and how quickly the animal should be moving, thereby generating a realistic looking gait for that particular animal.

We built a simple system using OpenGl where the positions of the hoof determined the position of all bones for the rest of the leg using Inverse Kinematics, see Fig. 4.14. Each hoof has a trajectory that specifies its path, where the foot should be planted, how high it should lift off the ground and the distance it needs to move, see Figs. 4.15 and 4.16. This resulted in a somewhat stamping gait. The main issue with this approach is that we do not know how exactly the hooves and legs move and which path they follow (for real animals). In our simulation the hooves followed a relatively high arc, whereas in real life the hooves are barely lifted off the ground during the walking gait as this requires the minimum amount of energy. Similarly, the path the hoof follows is not one smooth curve, it dips at certain points during the motion (we only noticed this when looking at the actual motion data) and even though this movement is very slight and happens for one frame only, it is obvious when it is missing from an animation. Finally, the upper body of our model was not animated and looked very stiff. Therefore, we decided to capture the motion of real animals in order to be able to set up a system that would then be able to replicate it. For this reason we decided to study real animal movement and to investigate people's perception of it first. This could then be used to provide guidelines for procedural animation in the future.

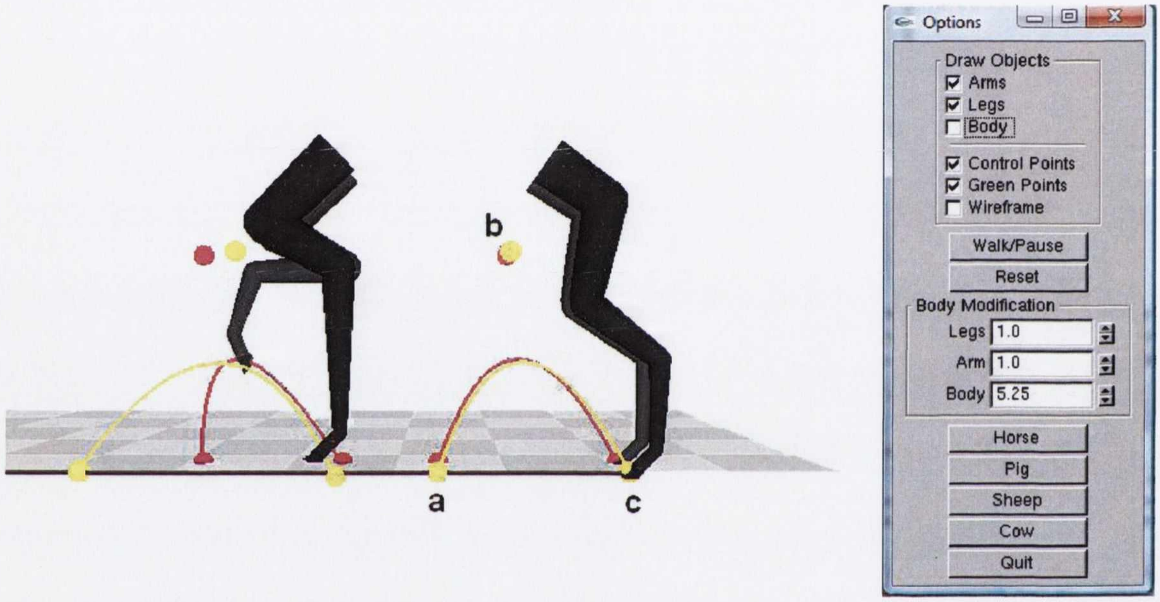

Figure 4.15: Points a, b and c are used to form a curve (a - next position, b - controls the height of the curve the hoof follows, c - current position). The curve represents the path each hoof is to follow. When the hoof is moving, the rest of the leg bones are calculated using IK. When the hoof is planted and the hip is moving, the rest of the leg bones are positioned by solving the IK in a reverse direction (from the hip down). The size of the legs can be changed using the interface on the right and pre-defined animals can also be displayed.

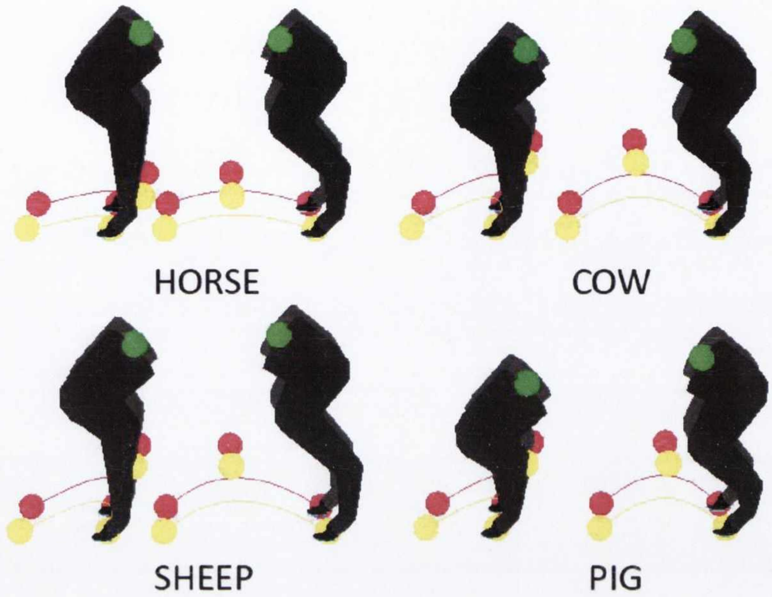

Figure 4.16: To implement a generic method to animate all different sizes of quadrupeds we test it on pre-defined legs for a horse, cow, sheep and pig. Their proportions are based on drawings from [SF06].

4.4 Discussion

In this chapter we presented the techniques we tried in order to achieve realistic animation of quadrupeds. We encountered many challenges along the way and decided it would be useful to know whether people can notice differences in motions between different types of animals, as this might simplify the task of animating a variety of species. Gathering required information about the motion that is to be replicated is of vital importance. However, this task can be very difficult if one wishes to animate animals that are difficult to approach, film, photograph and capture. Therefore, knowing whether a similar motion of an easily accessible animal can be used instead would be very useful, e.g., can you use the motion of a cat to animate that of a tiger? In addition, if one can animate various models of animals using the same motion, it would save a lot of time and effort, both for the information gathering task and also in the actual animation task itself. Similarly, it is possible that this information can be used when designing real-time simulations of animals.

Finally, we wish to note that, in our early work we used hand animated impostors to create flocking and shepherding behaviours. However, the animation of the characters used was not of a high enough standard (stiff animations, evident foot-skate) and it took away from the realism of the overall flock simulation. For this reason it would be useful to know which animations look good for particular herds or flocks. For example, a sheep model can be animated with a cow motion using re-targeting, it is then possible to generate a flock of sheep using this animation, if this simulation looks convincing (as a result of perceptual tests), it means that it is worthwhile creating their impostor representation for use in a real-time system. However, it is also possible that this particular animation does not look well for a flock of sheep and therefore having insights into what motions are perceptually appealing is useful.

Chapter 5

Perception of Animal Gaits

Why is it necessary to perform studies on human perception of animal motion? As demonstrated in previous chapters the animation of animals is quite complex and particularly difficult for non-expert animators. If we could find ways of simplifying the process using knowledge of human perception this would be very useful for the animation, games and other communities. When animating animals by hand, we found that we spent such a long time getting the motions of the legs perfect that we neglected the motion of the upper body, this was creating an unrealistic animation, in our opinion. We therefore thought it would be interesting to test whether humans looked at the motion of the animals' legs or if they looked at the upper body more. If they did look at the upper body more it may mean that more emphasis should be placed on animating it well. In addition, it is possible that we cannot tell the difference between the motions of animals, and this would mean that one animation can be used on several models, thereby speeding up the overall process.

There is much work published that includes many perception studies involving human motion and appearance. Johansson was one of the first people to attach lights to the joints of humans in order to isolate the motion itself, so that the perception of the motion was not influenced by appearance cues [Joh73]. His technique revealed that we are indeed able to identify human motion from just a small set of moving points. Later, Cutting and Kozlowski showed that we can tell the gender and identity of those actors represented by the point-lights [CK77]. Other studies involving the motion of human characters include the work of Hodgins et al. [HOT98]. In their experiment they use a stick figure and a polygonal model to see if people are more sensitive to the motion of one representation over the other. They found that humans are more aware of changes in motion when represented with the polygonal model, and suggest that these studies can be used as a guideline for optimization of animation sequences.

Studies by McDonnell et al. [MJH+08] conclude that both shape and motion influence sex perception of virtual human characters. They found that a male walk on a female model, and a female walk on a male model, were almost always perceived to be ambiguous or anomalous. Similarly we wanted to find out if it is possible to animate different models of animals realistically by using the motion of other animals. Would this also always be anomalous, or are we as humans less sensitive to such variations in animal motions? Their studies also revealed that a male model with a neutral walk is perceived as masculine, and a female character with the same neutral motion is perceived as feminine. Therefore, in scenes where there are male and female characters, for example in a crowd, using neutral motions could help optimize memory resources that would otherwise be needed to store male and female motions. Once again, it would be useful to see if there is a neutral animal motion that can look realistic on different animal motions. The work of McDonnell et al. [MLH+09] suggests that motion variety is not as important as appearance variety. If this is true for humans, then it may be even more so for animals with whom we are less familiar. As previously shown in Chapter 3, Hildebrand studied the motions of different species of horse to check for variation in the pace and trot motions. He found that the timing of the footfalls differs between the different species of a horse rather than between individual horses [Hil65]. It is outside the scope of this work to capture multiple versions of each gait for different animals, as it would be too labour intensive, but it is an interesting question for future work.

There are few relevant perception studies involving animal motions, indicating the novelty of our studies both in the animation and perception communities. Mather and West [MW93] have shown that people are capable of recognizing animals from point-light displays. They created the point-light videos using still images created by Muybridge [MB57]. They found that the participants viewing the videos could identify the animals represented but when they were presented with a still image of dots they were unable to identify the animals represented. This echoes the earlier work of Johansson where a static image of points was not judged as representing a human. Vuong et al. [QCVT06] performed experiments to test the extent to which observers use dynamic information to detect targets in natural scenes. Results showed that observers performed more accurately with dynamic than static target scenes. Target patterns used in the experiments consisted of a walking human, either static or moving, superimposed onto a distractor scene containing machines or animals. Users were asked whether the human was present or absent in the scene. Overall dynamic target scenes resulted in better detection performance than static target scenes. In experiments where the human was inverted the results showed the same pattern, so humans were able to quickly learn to detect unfamiliar inverted human motion.

In our work we used point-light videos that represent the real motion of animals and ask participants to identify which animal is represented in each video. We followed up with three more studies using realistic 3D models in a variety of settings. To our knowledge, these are the first experiments that investigate the perception of animal motion in such depth.

We would like to note that the number of participants in each experiment varies as participation is completely voluntary so the number of people that take part changes. In addition in some experiments, such as the eye-tracking experiment, participants become tired quickly (they have to sit very still in order for their eye movements to be tracked) therefore the number of volunteers is reduced further. However, we find that eye movements are consistent between people indicating that only a small number of participants are needed for these experiments. Another way to determine whether there are enough participants for an experiment is to analyse the results collected. If the results show a clear pattern of participants' behaviours/reactions and statistical significance can be computed then the number of participants are sufficient. Statistical significance indicates that the results have not occurred by chance therefore increasing the number of participants does not necessarily add to the experiment outcome. For example, as can be seen Section 5.3.3 we only have five participants in our trial but the results of are clearly divided so adding another twenty participants would not change the outcome of the experiment as their results would follow the same pattern of the first five.

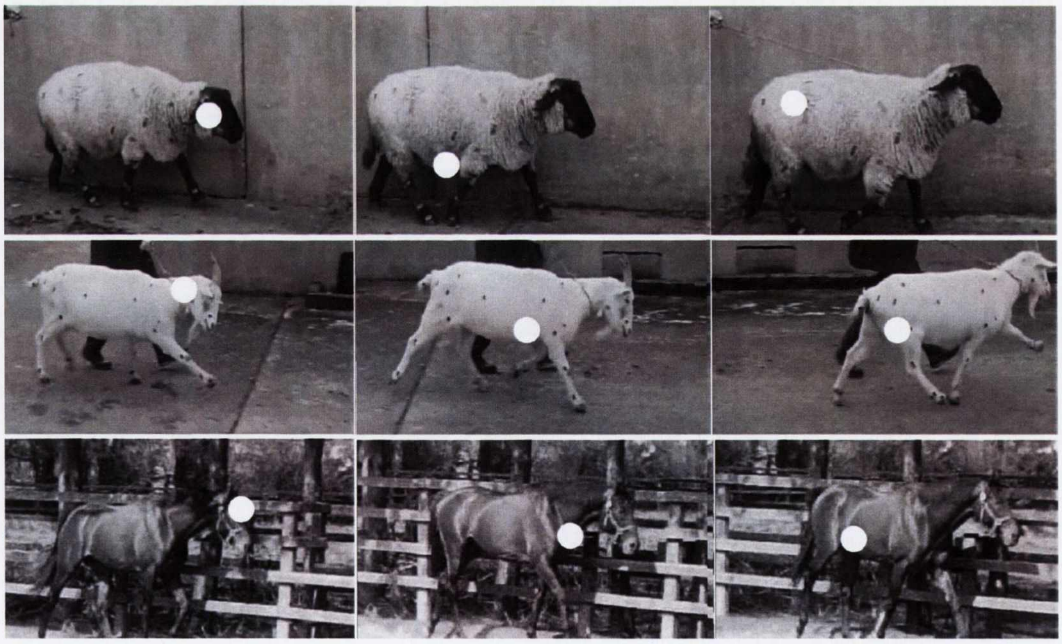

Figure 5.1: Frames from videos showing fixations. It was found that the head is looked at first, followed by the torso (front then back) and then the front legs. The front legs of the goat were looked at more so than for any of the other animals because the goat we filmed had a limp in the front leg and so it drew more attention.

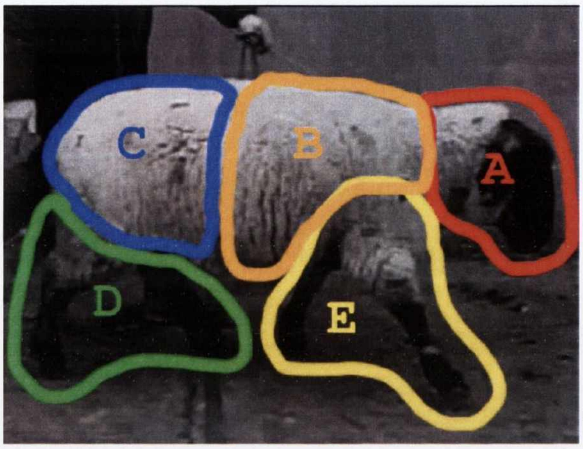

Figure 5.2: Regions of the body: A: Head, B: Torso, C: Hips, D: Back legs, E: Front legs

5.1 Eye-tracking

In our first experiment, our aim was to find out which area of an animal's body people looked at most. We therefore collected a set of videos depicting animals walking and trotting. Previous eye-tracking results have shown that faces are particularly salient for static images of animals and humans [HBaANS04, HHO05, MLH+09]. To explore whether similar eye-movement patterns are found for dynamic scenes depicting animals, we displayed multiple 4-second (56 frame) grey-scale video clips of farm animals (goat, horse, sheep) walking and trotting. Using an EyelinkII eye-tracker, we recorded the eye-movements of seven participants who were instructed to view the experiments with a view to subsequently answering questions about the movements. As it has been shown that human and animal motions activate different areas of the brain in children [MC03], we also showed the participants the same number of videos depicting humans walking and running.

Fig. 5.1 shows three frames from three video clips, with the eye-fixations of one participant overlaid. This depicts a typical eye-movement pattern. Most participants first looked at the head of the animal, then looked along the torso, finishing at the hips.

To examine the results more closely, we considered an animal's body as consisting of 5 main regions: head, torso, hips, back and front legs, see Fig. 5.2. We counted the number of fixations in each region for each participant across all video clips (note: for the comparison with humans, we averaged over the animals' front and back legs). We performed two-factor within-subjects ANOVAs for both sets of comparisons, and the results are shown in Figs. 5.3 and 5.4. For the animals, we found no main effect of species, but a main effect of body part ($F(4, 24) = 13.45, p = .00001$). Newman-Keuls post-hoc analysis confirmed that participants

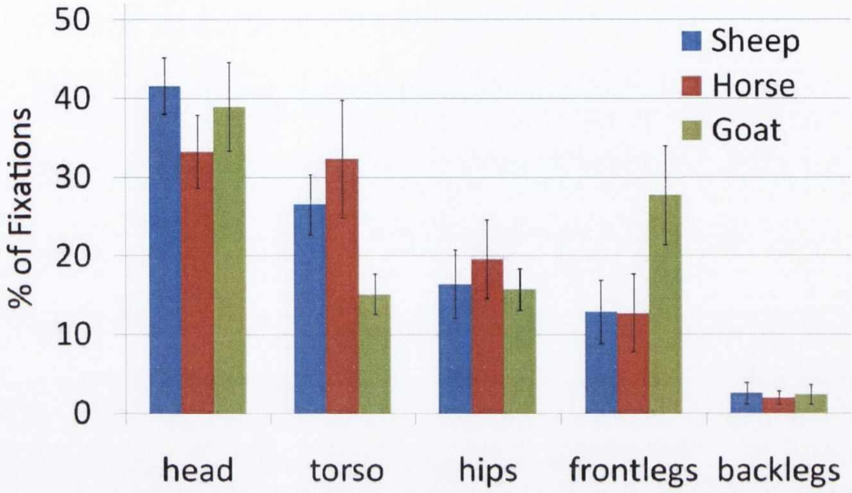

Figure 5.3: Results from tracking of different participants' eye movements. Overall the upper body is looked at more than any other part of the moving animal.

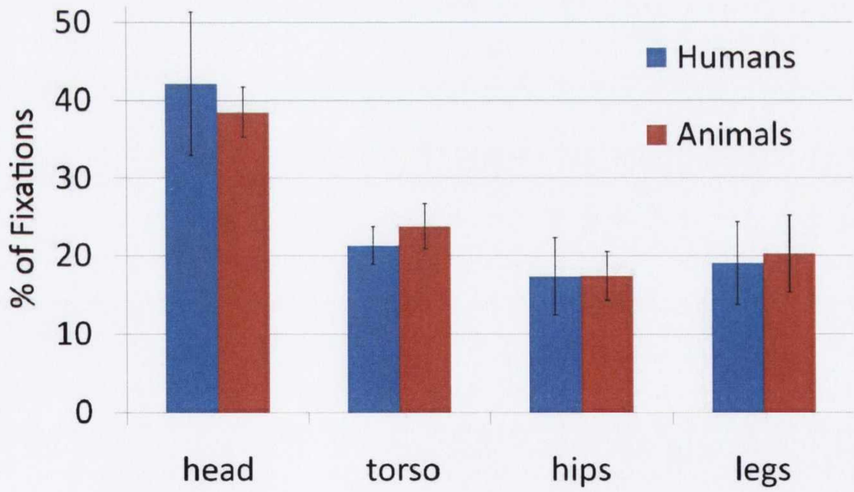

Figure 5.4: Results from tracking of participants' eye movements viewing movement of animals vs. humans. The same sort of pattern can be seen for both humans and animals where the head is looked at the most.

fixated on the head significantly more often than any other region, and on the back legs significantly less often. An interaction effect was also found, because the goat had a slight limp and hence the front legs attracted significantly more attention. For the animals vs. humans comparison, we again found a main effect of body part ($F(3, 18) = 5.19, p < .01$), with the head region again receiving significantly more fixations. However, there was no significant difference in results for humans and animals, indicating the equal importance of faces in both cases. Similarly, in the work of McDonnell et al. [MLH+09], eye-tracking experiments showed that for a 3D human model the face and the upper torso were looked upon first and that those areas received the most fixations.

It is interesting to see from our results that there is almost no difference between the way we look at humans and animals. The results also suggest that we are correct in saying that the animation of the upper body of animals is important, as this is the first place people look at. The legs on the other hand are only studied for a longer period of time if there is an anomaly as in the case of the goat with the limp. Therefore we can conclude that although viewers do not necessarily focus their attention on certain parts of moving animals, if there are anomalies they will detect this. Therefore it is important to avoid anomalies, such as foot skate, even if the final animation is not accurate.

5.2 Point-light walkers

In order to find out whether we can tell the difference between animal gaits we captured the motion data of real animals that are similar in physique. We continued to work with farm animals, as their size makes them easier to mark, they are relatively tame, and we also had access to these animals at the University College Dublin (UCD) farm used for veterinary studies. We had full ethical approval from the UCD ethics committee that oversees research with animals (Appendix A.1). Challenges included working in a relatively uncontrolled environment, where the animals were outside, sometimes on uneven ground, and could behave unpredictably. Using water-based paint we marked the animals by painting dots on their joints as shown in Fig. 5.5. They were placed in an enclosure so that they would not run away and we recorded the motions of each animal over a period of time. In order to keep the point-light videos of all the animals the same, we only used video clips where the animal was found walking or trotting parallel to the camera.

To track the points marked, we used the technique described in Chapter 4.12 and we animated a set of discs which replicated the real motion. Figure 5.7 shows the point-light walker for each animal. These animations were used to create the final point-light walker representations of the real animals used in our first experiment. Each video contained the

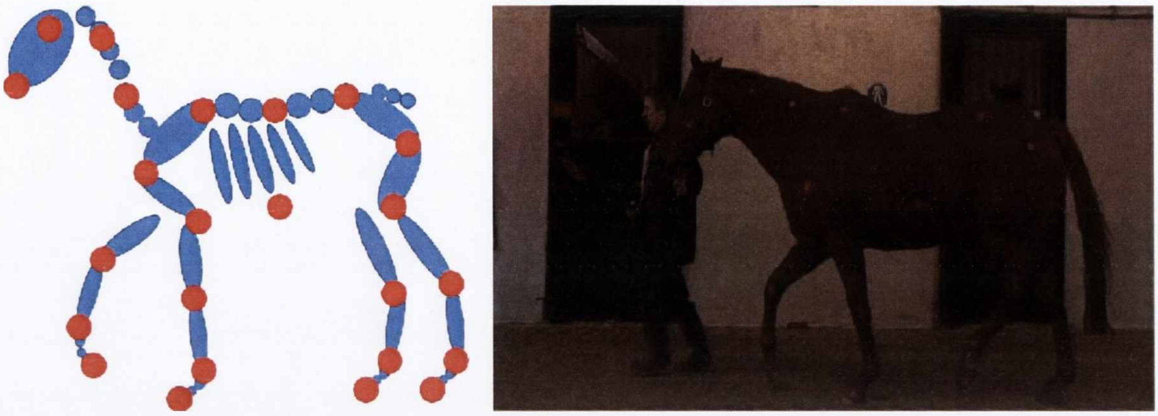

Figure 5.5: On the left is a simple image we used as a guideline for marking the animals, on the right is an image of the horse and how we marked it before filming.

Figure 5.6: An example of the other animals captured, the marked points have been highlighted in order to make them easier to see.

motion of either a horse, cow, sheep or pig. For each animal we tried to have two gaits, a walk and a trot. When recording the pig we were unable to capture two different gaits as the pig either ran or stopped entirely to eat. In this case we slowed down the original motion so that we would have two different speeds of motion like the other animals.

5.2.1 Modified gaits

Taking the animation data we collected from the animals used, we created modified stimuli of animals performing the same gaits. To create these unfamiliar animals we changed the proportions of the limbs and other body parts of the animals. These changes can be grouped into leg length, body length or body width modifications, or a combination of these. In some cases we tried to make one animal have the characteristics of another; by reducing the width of the cow torso and lengthening the legs it looked more like the horse point-light representations. We applied similar changes to other animals too; giving the horse a wider

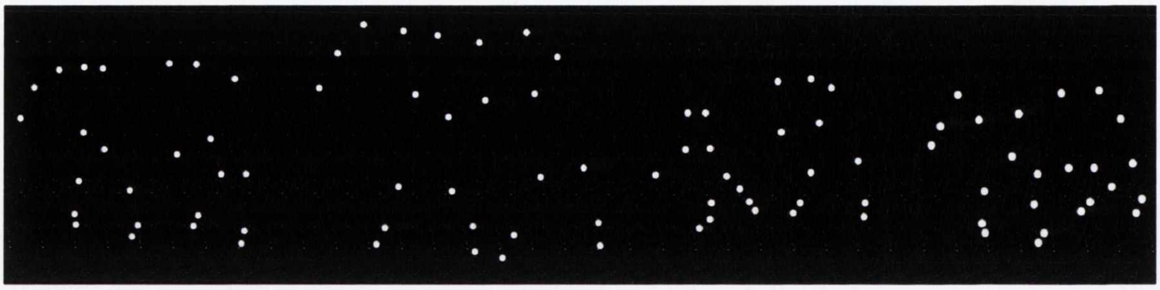

Figure 5.7: Point-light representations of the cow, horse, pig and sheep used in the experiment.

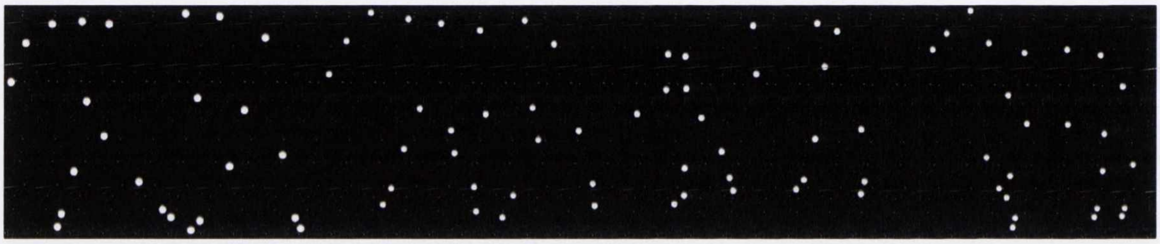

Figure 5.8: From left to right: cow with longer legs and a slimmer torso, horse with shorter legs and a wider torso, pig with longer legs and a longer torso and sheep with longer legs.

torso and shorter legs make it look more similar to a cow or sheep even, see Fig. 5.8. These motions were used to see if people could still recognise the fake motions as the real animal, or whether the changed proportions led to them being perceived as something different or ambiguous.

5.2.2 Synthetic gaits

In addition to using real motion data, we animated a group of 3D models representing the farm animals using traditional hand-animation methods. For each model, we used Muybridge photographs (plates 33, 96, 98 and 100) from [MB57] to create the animations. From these animations we created a set of synthetically generated point-light walkers. As previously discussed, McDonnell et at. found that synthetic motions were perceived to be more neutral, so we included these for completeness. However, when these videos are compared directly to the natural motion, it is obvious that the synthetic movement lacks fluidity and vertical motion.

5.3 Experiment 1

For all our results, we applied a two factor ANalysis Of VAriance (ANOVA) [How99] with repeated measures, followed by Newman Keuls post-hoc tests for significance. Throughout this thesis we only report results that are significant at the 95% level or above, i.e.,$p < .05$. The three conditions presented below were all on-line experiments where we sent the link to students within our university. Each participant did only one randomly assigned trial and we recorded the IP addresses to ensure that they were valid and that nobody did the experiment twice. Each participant was asked for their age and gender. The videos were played in a random order for each participant, and after each video they were asked a multiple choice question.

5.3.1 Original Motion

In this trial we used videos of point-light walkers that were 1 to 3 seconds long. Each video contained the motion of either a horse, cow, sheep or pig. After seeing each video, participants were asked to indicate whether the motion of the point-light walker came from a horse, deer, cow, sheep, dog or pig. We added the dog and deer options as distractions, one being a completely different type of animal (i.e., a carnivore) and the other more similar (i.e., another ungulate). We recorded the results for eleven participants (9 male, 2 female), aged between 22 and 34, who viewed the videos in a random order. For each video shown, a mirror image of the video was also shown, so that the participants did not associate a direction with a particular animal. The results can be seen in Table 5.1.

We found a main effect of response ($F(5, 50) = 4.7409, p < .01$) where the horse response was picked significantly more times than the other choices. There was an interaction between gait and response ($F(5, 50) = 3.2212, p < .05$), where for a trot gait, the horse option was picked significantly more times than a deer or cow. It is also interesting to see that the horse was picked significantly more times for a walk gait than a cow was for a trot, suggesting that humans are more familiar with seeing horses walk than they are with seeing cows trot, and so they would not associate a trot gait as a cow motion. We also found an interaction between gait, animal and response ($F(15, 150) = 1.9382, p < .05$) where the horse trot and the horse walk were identified as a horse significantly more so than any other animal. The cow walk motion was identified to be a cow significantly more times than any other animal. This might be based on viewers' expectations that a cow should walk slowly. For the pig trot, the pig response was selected significantly more times than the horse, deer, cow or sheep option, and the pig walk was identified to belong to the pig significantly more times than any other animal.

	ACTUAL							
	Horse	Horse	Cow	Cow	Sheep	Sheep	Pig	Pig
Response	Walk	Trot	Walk	Trot	Walk	Trot	Walk	Trot
Horse	54.55	77.27	9.09	31.82	31.82	13.64	0.00	0.00
Deer	18.18	13.64	18.18	4.55	27.27	18.18	9.09	0.00
Cow	18.18	0.00	54.55	13.64	0.00	4.55	4.55	0.00
Sheep	0.00	4.55	4.55	22.73	18.18	31.82	13.64	9.09
Dog	0.00	0.00	9.09	18.18	9.09	22.73	13.64	31.82
Pig	4.55	0.00	0.00	4.55	0.00	0.00	54.55	59.09

Table 5.1: Experiment 1, original motion; 'Actual' represents the videos shown to the participants. 'Response' represents the possible answers the participants could have picked. The values are the percentage of times the participants chose a particular animal for videos where the *original* point-lights were shown.

The distractors reveal some interesting results. The participants picked the deer response more often for the sheep than for any other animal, even though the deer's size is closer to that of a cow or a horse. It is possible that this is because the sheep was filmed on very uneven terrain, so its motion may suggest that it is an animal filmed in the "wild". On closer inspection we can see that the speed of the motion had an effect on choices made for the identity of the sheep gaits; the deer response for a sheep walk was 9% higher than for a sheep trot.

5.3.2 Modified motions

In the second trial, we used the modified motions (as described in Section 5.2.1) of real animals and asked the same question as in the first condition. For this experiment we had nine participants (6 male, 3 female) aged between 24 and 40. We had seven videos for each gait (walk, trot), for the four animals (horse, cow, sheep and pig), and we mirrored each video again, resulting in 112 videos. Once again we applied a two factor ANalysis Of VAriance (ANOVA) with repeated measures to the results, followed by Newman Keuls post-hoc tests for significance. The results are shown in Table 5.2. Again we found a main effect of response ($F(5, 40) = 4.7634, p < .01$), where the horse response was selected most often. We also found that there was an interaction between response and gait ($F(5, 40) = 15.768, p < .01$), where the horse trot was identified correctly significantly more times than any other animal. For a trot gait, the sheep response was picked significantly more times than the deer, cow or pig. For a walk gait, the cow response was picked significantly more times than any other animal.

We also found an interaction between animal type and response ($F(15, 120) = 24.493, p <$

	ACTUAL							
	Horse	Horse	Cow	Cow	Sheep	Sheep	Pig	Pig
Response	Walk	Trot	Walk	Trot	Walk	Trot	Walk	Trot
Horse	57.78	81.75	11.90	53.17	6.35	1.59	0.00	1.59
Deer	7.50	5.36	5.56	5.36	46.83	20.74	13.49	10.32
Cow	25.19	3.97	75.40	18.25	6.35	2.68	2.44	3.57
Sheep	2.50	0.79	4.46	13.49	37.30	71.91	12.76	13.39
Dog	4.44	8.73	2.38	4.76	2.38	12.09	16.67	30.95
Pig	2.22	0.00	0.79	5.56	0.00	0.00	53.85	41.27

Table 5.2: Experiment 1, modified motion; the values represent the percentage of time the participants chose a particular animal for videos where the *modified* point-light walkers were shown.

.01). In our analysis we will focus on the results for each type of animal and how they were identified. The horse was identified correctly significantly more times than any other animal. The cow and horse responses were picked significantly more times than any other when participants saw a cow motion on the screen. For a pig motion, the pig response was selected significantly more times than any other. A sheep and deer response were selected significantly more times than any other when the participants were viewing a motion of a sheep.

Finally, there was an interaction between gait, animal type and response ($F(15, 120) = 8.8632, p < .01$). Again we will focus on how the animals were identified in our analysis. The horse response was selected significantly more times than any other response when a horse trot was displayed, and the horse trot motion was identified as belonging to a horse significantly more times than the horse walk. However, a horse response was picked significantly more times than any other when a horse walk was on the screen, so people were able to identify both the horse trot and the horse walk. A horse response was selected significantly more times than any other when a cow trot was on the screen, but when a cow walk was on the screen then the cow response was selected significantly more than the others. A pig response was selected for a pig walk significantly more times than any other response and similarly it was selected significantly more times than all other responses, except for the dog, when a trotting pig was displayed. For the sheep, the sheep response was picked significantly more times than the others when the sheep was trotting and the sheep response was also picked significantly more times than the other responses, except for the deer, when the sheep walk was displayed.

From Table 5.2, we can see that the recognition accuracy for almost all the animals has increased with the modification of the original motions: an unexpected and interesting result. This is true in particular for the cow walk and the sheep trot where the accuracy has increased

	Video		
Animal	Synthetic	Trot	Walk
Horse	11.00	66.83	67.39
Cow	8.00	61.54	57.33
Sheep	1.00	59.00	39.75
Pig	14.00	42.00	46.57

Table 5.3: Experiment 1, synthetic vs. real; a higher value indicates a more realistic motion and a lower value indicates a synthesised animation. The low values for the synthetic motions indicate that for all animals, these motions were obviously unnatural and were never mistaken as a real.

by 20.85% and 39.49% respectively.

5.3.3 Modified vs. Synthetic

In the third trial, we used the modified motions of real animals as in the second condition, but also added the synthetic point-light walkers to the playlist that were created using the technique described in Section 5.2.2. The videos were played as in the first and second trials but this time the participants were asked to rate how real the motion was on a scale from 1 to 5, where 1 indicated realistic motion and 5 indicated synthetic motion. For this trial we had five participants - all male aged between 25 and 30.

The average ratings for each video were calculated. The results are shown in Table 5.3. We found a main effect of animal type ($F(3, 12) = 4.9102, p < .05$), where the synthetic cow and sheep point-lights looked significantly more synthetic than the horse. We found that there was an interaction between animal and video ($F(6, 24) = 4.7030, p < .01$). For all animals the synthetic animations looked less realistic than their modified real motions.

Table 5.3 shows the results for the third condition. Small values indicate that the motion looked very synthetic and large values indicate that the motions looked realistic. We can see that there is a large difference between the values for synthetic motions (all under 15%) and the modified ones (all over 39%).

5.4 Experiment 2

We ran this experiment as part of an event called "Metropolis: Crowd Control" organised within Trinity College Dublin to show the public the kind of research that takes place in our lab. Virtual crowd behaviours, urban planning, pedestrian and traffic modelling and evacuation simulations are just some of the topics researched using the Metropolis system.

During the event the lab organised perceptual experiments that the public took part in, a description of these experiments can be found in Appendix A.2.1.

Our experiment, "What's gotten into ewe?" involved a 3D sheep model that was animated using the original motions of our point-light walkers. We got ethical approval from the university which enabled us to store the results of the experiments for all the adult participants (the ethical form can be found in Appendix A.2.3). This particular experiment proved to be very popular with the public - we had close to 100 participants (we were able to use 69 as the other participants were under 18 years of age), which shows people's keen interest in animal studies. The experiment received particular mention in the national paper as can be seen in Appendix A.2.2.

For this particular experiment we chose to re-target the motions of the original point-light walkers onto a 3D model as we wanted to see whether the appearance of a 3D model had an effect on the perception of motion. By applying the motion to a single model, we made it difficult to identify the proportions of the animals the motion belonged to. We have received much feedback on our work from various reviewers, most claiming that our ideas are useful and have much to contribute to the field. Their suggestions coincided with ours as we planned to test people's sensitivity to animal motions using 3D models. Using a 3D model creates a more realistic test-bed for perception as in real life we cannot see motions of objects only - there is always a visual cue as to what the moving object is. In addition, the trend in the game and film industries is towards more complex and detailed 3D models rather than simple point-light or stick figures. The sheep model was chosen in this instance as it is similar in shape to the horse and cow, but its size is similar to the pig. The sheep model was animated by re-targeting the data from the horse, cow, sheep and pig, as described in Section 4.3.1. Fig: 5.9 shows the animated 3D sheep model.

In this experiment, participants were asked to select whether the motion of a 3D sheep model had come from a horse, deer, cow, sheep, dog or pig. As before, we created mirror images of all the videos and the background to all the animations remained black. Participants sat in front of an 18-inch monitor and indicated their choices using the number keys on the keyboard. Each video was 3 seconds long, we had two videos for each animal gait that we repeated resulting in a total of 32 videos. The videos were played in a random order, and we used Presentation® software to run this and subsequent experiments and to store the results for each participant.

We hypothesized that seeing the upper body would help participants identify the correct animal from the motion. Our eye-tracking experiments described in Section 5.1 showed that humans tend to observe the upper body more than the legs of moving animals and humans. We therefore ran two conditions, one with the full body visible and another where we hid the

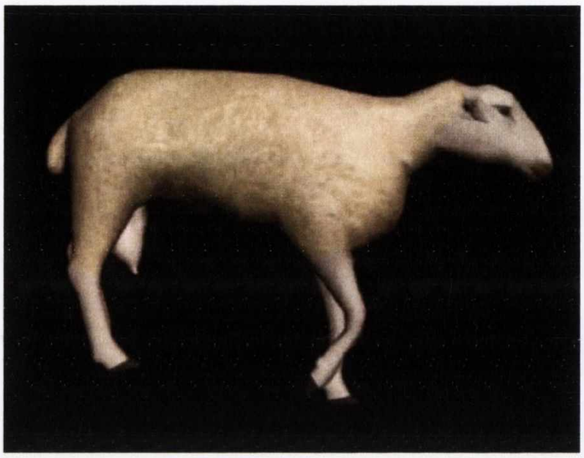

Figure 5.9: Sheep model walking with a cow walk gait.

upper body from view (for each condition we had 32 videos). There were sixty-nine (38 male, 31 female) and thirty (19 male, 11 female) participants for the full body and hidden body conditions respectively. We allowed participants to take part in both experiments if they wished and in this case the legs-only videos were shown first in order to avoid bias. As we described in Section 3.1.1 the spine varies between different ungulates and this is represented on 3D sheep model as the different animal motions are applied to it. For example, when the sheep moves like a cow the spine is very flat, but when the horse motion is applied the spine has a distinctive dip (where the saddle goes on the horse). It is therefore possible that the participants use the visual cue of the spine to decide which motion has been applied to the 3D sheep. Allowing participants to see these full body videos first poses the risk that later, when participating in legs only trial, they would remember a certain leg movement with a spine curve, so even if the spine is not visible they could tell which motion was displayed.

The results for the full body condition are shown in Table 5.4. We applied a two factor ANalysis Of VAriance (ANOVA) with repeated measures, followed by Newman Keuls post-hoc tests for significance. We found a main effect of response, where the horse was identified correctly significantly more times than any other animal ($F(5, 340) = 9.9082, p < .01$), on average 38.4% of the time.

There was an interaction between gait and response ($F(5, 340) = 33.982, p < .01$). All walking gaits were identified to be from a cow more often than from any other animal. From Table 5.4 we can see that on average, for all the walk gaits, the cow was selected 29.17% of the time, where chance performance is at 17%. As in Experiment 1, this suggests that participants associate cows with slower gaits. On the other hand, the dog was selected least often for the walk gait, 4.9% on average, which suggests that, unlike cows, we identify dogs

	ACTUAL							
	Horse	Horse	Cow	Cow	Sheep	Sheep	Pig	Pig
Response	Walk	Trot	Walk	Trot	Walk	Trot	Walk	Trot
Horse	31.16	45.65	21.74	18.48	12.32	28.26	6.88	6.88
Deer	13.77	20.65	11.59	10.14	27.90	27.17	11.59	15.94
Cow	35.51	10.14	40.94	5.79	14.13	4.53	26.09	8.33
Sheep	8.69	11.96	14.13	23.91	20.29	12.68	15.94	19.20
Dog	2.17	6.88	3.62	36.23	7.24	19.93	6.88	20.65
Pig	8.69	4.71	7.97	5.43	18.12	7.97	32.61	28.99

Table 5.4: Experiment 2, 'Actual' represents the videos shown to the participants. 'Response' represents the possible answers the participants could have picked. The values are the percentage of time the participants chose a particular animal. Up to 17% indicates chance performance.

as being more active. For the trotting gaits, the horse was chosen significantly more times than any other animal except the dog, which was chosen equally often.

There was an interaction effect between animal and response ($F(15, 1020) = 22.824, p < .01$). The horse and the pig were identified correctly significantly more times than any other animal. Again, the sheep was identified as a deer significantly more times than any other animal. As in Experiment 1, this may be due to the motion of the sheep over the uneven ground. But again, it suggests that factors other than the innate motion of an animal affect our identification of it.

We also found an interaction effect between gait, animal and response ($F(15, 1020) = 6.3940, p < .01$), which was mainly due to participants reacting differently to walking and trotting gaits for the horse, cow and dog. A possible explanation could be that people are far more used to seeing horses moving quickly. For example, horses are often used in movie action scenes (e.g., westerns and battle scenes). The horse and cow gaits were less often confused with any other animal, perhaps because these gaits were more obviously from larger animals.

The cow motion was identified correctly more times when it was walking rather than trotting, perhaps because cows are typically seen standing or walking around in fields. The cow trot was identified as a dog gait significantly more times than any other animal. From Table 5.4 we can see that a horse walk was chosen to belong to a cow 35% of the time, while a cow trot was seen as a dog motion 36% of the time. We suggest that there are three possible reasons for this: we are not used to seeing cows trot; the dog is a very active animal; and it is obvious that it is not the motion of a horse, which appears to be the most easily identifiable gait. Again, we concluded that other factors played a part in the identification of these gaits, not just the innate feature of the animal motion itself.

There are some similarities between the results we got from the the full-body trial and the legs-only trial. As for the full-body trial, we applied a two factor ANalysis Of VAriance (ANOVA) with repeated measures, followed by Newman Keuls post-hoc tests for significance. We found a main effect of response where participants picked the horse or cow response significantly more times than deer, dog or pig, ($F(5, 145) = 6.2037, p < .01$). We had a similar result for the full-body trial where the horse response was picked significantly more times than any other animal.

There was an interaction effect between gait and response ($F(5, 145) = 15.164, p < .01$), where, as for the full body trial, the cow was picked significantly more times than any other animal for the walking gaits, and the horse was picked significantly more times than any other animal for the trotting gaits. As with the full body trial, we found an interaction between the animal and response ($F(15, 435) = 5.2620, p < .01$), however, in this case we found that for a cow gait, the horse response was picked significantly more times than any other except for the times when participants picked the cow response. In the full body trial the horse response was picked significantly more times than any other for a *horse* gait.

We also found an interaction between gait, animal and response ($F(15, 435) = 2.6621, p < .01$), which as above is mainly due to people reacting differently to walking and trotting gaits. A horse is picked significantly more times than any other animal except the dog, for the cow and sheep trotting gaits. Above the horse is associated with the horse trot the most, however in both trials (full body and legs only) the horse is picked when a faster movement is displayed. In this trial, the cow response is picked significantly more times than horse, dog or pig for the horse walk motion, and it is also picked significantly more times than the sheep, dog or pig for the cow walk motion. As, above the cow is picked more times for the slow motions than for the fast ones, and we can suggest as above that this may be because we are used to seeing cows standing in fields rather than trotting.

5.5 Experiment 3

In this final experiment on the perception of individual animal gaits, we used the cow and horse motions only. Our reason for this is that we found that the results from Experiment 1 and 2 show that the pig gait is consistently recognised and so we decided we did not need to test its perception on different 3D models. The sheep animation on the other hand had a low recognition rate, we suspect that this is due to the noise in the animation and also because it was recoded on uneven ground. Due to the low quality of the sheep motion captured we decided not to investigate the perception of it any further. We applied all four gaits to both

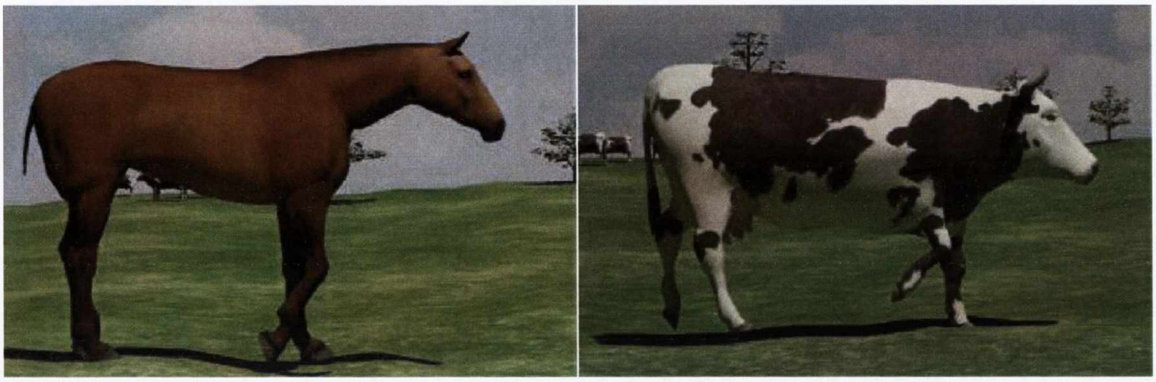

Figure 5.10: Horse and Cow models animated using the captured motion of a horse and cow.

the cow and horse 3D model (cow trot, cow walk, horse trot, horse walk). This time we also added a background condition, to see if a more natural scene had an effect on the perception of motion. The natural scene consisted of a grassy field with trees and cows visible in the distance, as can be seen in Fig. 5.10. The animations for both the natural background and the black background were identical and we included the mirrored videos in the trials as before. We ran two trials, the first where only the legs were visible as before and a second trial where the whole body was on display. Each video lasted approximately 3 seconds and participants were allowed to spend as much time as needed in picking their response. In this experiment we gave the participants a choice of two answers - cow or horse. For this experiment we had twenty-one participants (13 male, 8 female) who took part in both trials, completing the legs-only trial first. As we explained for Experiment 2, in Section 3.1.1 we show how the spine varies between different ungulates, therefore when applying the motion of a cow or a horse to a 3D model the associated spine shape is evident. As in Experiment 2, we prevented the participants from seeing the full body motions first, so that they would not associate a particular leg movement with a specific spine shape. Ethical forms for this and the following experiment can be found in Appendix A.3.

The results can be seen in Tables 5.5 and 5.6. A two factor ANalysis Of VAriance (ANOVA) with repeated measures was applied to the full-body data. A main effect of gait was found ($F(3,60) = 12.097, p < .001$). Post-hoc analysis using a standard Newman-Keuls test showed that the horse trot gait was the easiest to identify. However, the horse walk was the hardest to identify, significantly more so than either of the faster gaits. We can see from Tables 5.5 and 5.6 that, on average, participants performed with an accuracy of over 70% for the horse trot, whereas the horse walk was identified correctly only 41% of the time.

We also found an interaction between model and gait ($F(3,60) = 2.8661, p < .05$). Post-hoc analysis showed that there is a significant difference between participants being able to

	ACTUAL (COW MODEL)			
	Horse	Horse	Cow	Cow
Response	Walk	Trot	Walk	Trot
Horse	21.41	82.14	26.19	33.33
Cow	78.57	17.86	73.80	66.67
	ACTUAL (HORSE MODEL)			
	Horse	Horse	Cow	Cow
Response	Walk	Trot	Walk	Trot
Horse	54.76	72.62	59.52	48.81
Cow	45.23	27.38	40.48	51.19

Table 5.5: Experiment 3, These are the results for videos where the natural background was visible. As before, Actual represents the videos shown to the participants. Predicted represents the possible answers the participants could have picked. The values are the percentage of time the participants chose a particular animal.

	ACTUAL (COW MODEL)			
	Horse	Horse	Cow	Cow
Response	Walk	Trot	Walk	Trot
Horse	42.86	65.48	36.90	41.67
Cow	57.14	34.52	63.09	58.33
	ACTUAL (HORSE MODEL)			
	Horse	Horse	Cow	Cow
Response	Walk	Trot	Walk	Trot
Horse	48.80	69.05	60.71	40.48
Cow	51.19	30.95	39.38	59.52

Table 5.6: Experiment 3, These are the results for videos where the background was black. Actual represents the videos shown to the participants. Predicted represents the possible answers the participants could have picked.

recognize a cow walk on a cow model, and horse walk on a cow model. This can be seen clearly in Tables 5.5 and 5.6, where average recognition accuracy of a cow walk on a cow model is 68% and the accuracy for the horse walk on a cow model is only 32%. The horse trot on the horse model was selected significantly more times than the two walking gaits. These results support our earlier findings, as they indicate that participants did not identify the slow gaits with horses. In addition, as in Experiment 2, participants were able to easily identify a horse trot on the cow model, significantly more so than the horse walk.

We found an interaction between the background, model and gait ($F(3,60) = 5.7143, p < .01$). Due to the number of comparisons, we will only discuss the most interesting findings. In Table 5.5 we can see that in a natural scene, where a cow model was used, the horse trot was recognized 82% of the time, while the horse walk was recognized only 21% of the time. The low accuracy for the walk is a promising result from our perspective, as it implies that a horse walk can be applied to the cow model and still look convincing. Similarly, horse gaits on the horse model reveal that the trot has higher response accuracy than the walk. Surprisingly, the cow trot had higher accuracy in this experiment than in Experiment 1 and 2. We believe that this difference arises due to the participant having only two choices, horse and cow, of which the horse trot is the most distinctive. The average accuracies for the walk gaits on the horse model are 51% for a horse walk and 39% for a cow walk. This implies that the participants were guessing which gait belonged to which animal, suggesting that the gaits are indistinguishable, and that the horse model can be animated using the walk motion of either animal.

As in Experiment 2, we ran a trial where the upper bodies of the animals were hidden. Once again the results were similar to those where the bodies are visible. There was a main effect of gait ($F(3,63) = 5.123, p =< .01$), where post-hoc analysis using standard Newman-Keuls showed that the horse trot was significantly easier to identify than any of the other gaits.

5.6 Discussion

Across the three experiments we can see that the horse trot is the most recognised gait and the results are above chance performance. In Experiment 1 and 2, where chance performance is 16.6%, horse recognition accuracy is between 45% and 81%, and in Experiment 3 where chance performance is 50%, the accuracy is between 65% and 82%. However, the horse walk was better recognised for point-light videos but less so when the motion was applied to 3D models. In Experiment 1 for both condition 1 and 2, the horse walk was identified correctly over 54% of the time which is well above chance performance, whereas in Experiment 2 where the sheep model is used the accuracy dropped down to 31%. In Experiment 3 the accuracy is between 21% and 54%, where chance performance is 50% so this suggest that there is a possibility that the participants guessed whether the 3D models were moving with a horse walk. It is also possible that the motion did not transfer well to the 3D models. The work of Hodgins et al. [HOT98] may provide a possible explanation for this. Hodgins et al. performed an experiment where they use a stick figure and a polygonal model to see if we are more sensitive to the motion of one over the other and found that humans are more aware of changes in motion when represented with the polygonal model. Similarly, the work of McDonnell et al. shows that humans are less sensitive to motion variation on a 3D human model of a low resolution, whereas they are aware of these variations on a model that is of a high resolution [MDO05]. In our case it is possible that the participants noticed noise in the motion when seeing the creases in the polygonal model that would have been missing in a point-light representation.

For the cow motion the results are less clear. The cow trot had a very low recognition rate where the participants were given a large set of possible answers (in Experiment 1 and 2; horse, deer, cow, sheep, dog and pig). In this case, the recognition was between 5.8% and 18.5% and the chance performance was 16%, so it is reasonable to deduce that the participants were guessing or simply unable to recognise the motion. As with the horse model, the accuracy drops when the motion is applied to the 3D sheep model. In Experiment 3 participants were given only a choice of two animals (horse, cow) and the recognition was higher, but one needs to keep in mind that as a horse trot is most distinctive, it is possible that through a process of elimination the participants chose a cow response for a cow trot. The accuracy of the cow trot rates between 51% and 66%, which is close to the chance performance of 50% so it is still possible that some of the participants guessed their answers.

The cow walk on the other hand has a high accuracy rate over the three experiments. In Experiment 1 and 2 it ranges from 40% to 75% which is well above chance performance, and in Experiment 3 it ranges from 39% to 73%. It is interesting that, given only a choice of two answers, the accuracy is as low as 39%. However, this is good news for our perspective as it suggests that in this case the horse walk and the cow walk can be interchanged on the horse and cow models. The opportunity to use only one gait to animate two models means less work for animators.

We will now consider the results from Experiment 1 and 2 for the sheep and the pig gaits (they were not used in Experiment 3). In Experiment 1, first condition and in Experiment 2, the accuracy for the sheep walk was quite low, 18% - 20%, which is close to chance performance. The sheep trot on the other hand was recognised better when it was presented as a point-light (31%) rather than when it was on a 3D sheep model (12%). A possible reason for such a low score for the 3D model representation of the gait may be due to challenges present when capturing the data - the real sheep was filmed on uneven ground, and the sheep was wet which meant that the paint smeared and this introduced a lot of noise. As argued above, it is possible that once again, the motion did not transfer well to the 3D model, and also from the work of Hodgins et al. [HOT98] and McDonnell et al. [MDO05] it can be proposed here that the participants were more sensitive to the noise in the 3D model than the point light. Interestingly, in Experiment 1 condition 2, where the relative proportions of the limbs were modified, the sheep recognition accuracy improved for both the walk and the trot (37% and 71%). However, it is difficult to say which modifications of the point light contributed to this and more experiments would have to be done to investigate this further.

For the pig motion, the recognition rate drops between the point-light representations and the sheep model representation, as has been observed for the rest of the gaits. In Experiment 1 condition 1, the recognition of the pig is second highest compared to the other animals (54% - 59%), and this accuracy is replicated for the second condition of Experiment 1. In Experiment 2 the accuracy drops to 28%/32% for the walk/trot gait, which is still above chance performance and as before it is possible that on a 3D model, the motion is harder to decipher.

We also performed the comparisons between synthetic and modified real motions and found that there is a large gap that indicates it is easy to tell the difference between the two. This confirms our belief that in order to create realistic convincing animation, truly expert animators and artists are needed. While we are not trained artists or animators, we did have considerable experience with the animation tools and hence were not novice. We also had the benefit of reference materials such as Muybridge and videos. However, even with these supports we were unable to create synthetic motion which could identified as genuine animal motion.

Chapter 6

Towards perceptually convincing Herds and Flocks

Creation and improvement of realistic crowd representation has received much attention recently and includes research that focuses on behaviours, motions and appearances of virtual human, animal, and fictional characters. Studying human traffic flow has become very important during the design process of new buildings for safety reasons, for example to test whether the fire exits are placed in optimal locations. In the game and film industries there are many examples where crowd simulation can be seen e.g., in games such as Grand Theft Auto 4 (Rockstar Games) and Assassins Creed 2 (Ubisoft) and in films such as 300 (Warner Bros. Pictures), WALL-E (Pixar Animation Studios), Narnia (Disney) and A Night at the Museum (20^{th} Century-Fox Film Corporation) to name but a few. More recent films like Australia (20^{th} Century-Fox Film Corporation) show a herd of 5000 digital bovines and Kolve et al. from Rising Sun Pictures outline the work flow that allows them to successfully simulate these herds [KBF^{+}09]. However, there are few industry insights into how these crowds are perceived and their overall shape and appearance is controlled by the director or developer.

It is well known that variation adds realism to crowds [Rij09, Kan09], but the animations used need to be of a high quality from this to work. At the start of our research we used hand animated impostors to generate a flock of sheep. We had two animations for the sheep thus adding variation, but due to the low quality of the animations and the foot-skate that occurred during shepherding, the realism of the flock was never successfully portrayed. For this reason we applied the motions of the original point-light walkers to the 3D models of the horse, cow and sheep to create herds and flocks using the trotting and walking motions. We used these simulations to see which animations people preferred to see e.g., a flock of trotting sheep or walking sheep?

The work of McDonnell et. al. [MLD+08] gives insights into the perception of human crowds and how easy it is to spot clones. Through a set of experiments they found that in a crowd, appearance clones are far easier to detect than motion clones. We used the results of this work as a guideline for the generation of our herds and flocks and we added appearance and motion variation. By adding motion variation we could see whether people preferred the more varied herds or not.

In this chapter we will first describe how we used the basic principles of boids to simulate, in real-time, a realistic flock of sheep using impostors (the generation of impostors and the real-time system are described in Chapter 4). We will follow this with a description of our perceptual experiment involving flocks and herds of 3D animals.

6.1 Flocking

In Chapter 2 we reviewed the literature on generating realistic flocking behaviours. We will now describe the technique we used to achieve flocking and shepherding behaviour using several hundred sheep impostors and a dog impostor. As described by Reynolds [Rey87, Rey99], in order to stop a collection of objects from passing through one another in a flock, the distance between each pair of objects is measured and tested to make sure it is not less than the sum of their radii. However, this basic flocking algorithm grows in complexity as the order of the square of the flock's population increases ($O(N^2)$) therefore it is very inefficient for large populations.

6.1.1 Sweep and prune collision detection

In our implementation of the flocking algorithm, we used a Sweep and Prune technique [CLDP95] to keep track of all the animals in the system. This is a very efficient technique as it keeps the positions of all the flock members sorted. In this way, finding pairs of sheep which are close together is very quick and it also helps to keep the view of the flock localised for each sheep. This is important as land animals are not able to see the whole flock at any one time. Using bounding circles along with the Sweep and Prune technique allows for quick collision detection between the objects. Once detected, the animals are directed to walk away from each other. The Sweep and Prune technique also allows for easy implementation of cohesion, where sheep that are found within a certain distance of each other stay together by matching their velocities and orientation angles.

Dogs are often used to control flocks of sheep. In our simulation, we used one dog to move the flock towards a predefined target. In real life dogs need to be able to circle a flock very quickly to ensure that all the sheep are accounted for. Therefore, in our simulation the

Figure 6.1: Bounding spheres of different sizes are used for collision detection. We can see in the foreground that the sheep are turning away from the dog as they are within his bounding sphere.

velocity of the dog is much greater than that of the sheep. As the dog approaches the sheep they turn away from him - a principle on which steering the flock works. However, as the dog moves very quickly, his position changes significantly between successive frames. In order for the sheep to detect where the dog is, his bounding circle is made much bigger, thereby giving the sheep enough time to react, see Fig. 6.1. To steer the flock of sheep, steering points are calculated from the positions of sheep that lie on the perimeter of the flock and the sheep that are the furthermost from the target point, see Fig. 6.2. These points determine a Bézier curve along which the dog zig-zags in order to control the flock. This approach is similar to the one taken by Lien et al. [LBS+04, LRMA04]. As the flock migrates towards the target, the Bézier curve is recalculated to reflect the changes taking place. This ensures that when there are straying sheep, the dog can go after them and direct them back towards the flock.

The position and orientation of each sheep is an average of the vectors that represent the sheep turning away from the dog (if he is close enough) and the vectors that represent the cohesion and avoidance of the nearby sheep. In the case when there are many sheep together, the presence of the dog will not be noticed as each sheep takes into account the position of those around it. This is something that would need further attention, perhaps through the introduction of more dogs as in [LRMA04], or adding a weight to represent the dog's presence. It would also be more realistic if the sheep could run away from the

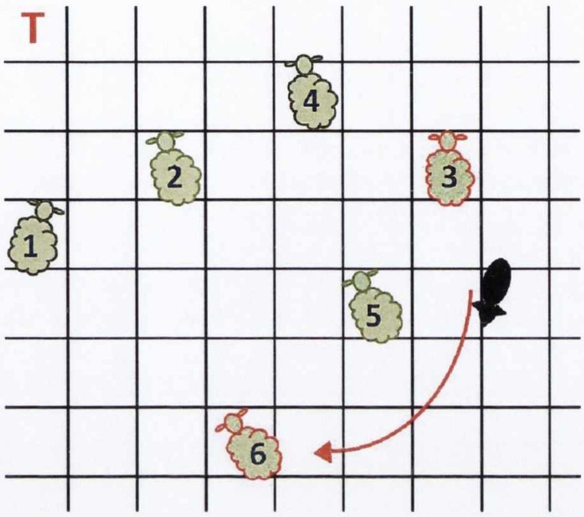

Figure 6.2: The sheep are sorted into two arrays, x: sheep1, sheep2, sheep6, sheep4, sheep5, sheep3, y: sheep6, sheep5, sheep1, sheep2, sheep3, sheep4. The sheep on the perimeter of the flock are sheep1, sheep3, sheep4 and sheep6. The sheep that are furthermost from the Target are sheep6 and sheep3 and so the dog steers the flock from behind them.

Figure 6.3: The dog runs after a sheep that is furthermost from the flock and the "Target". By moving back and forth across the foreground the dog pushes the flock towards the Target point.

dog, thereby creating a natural progression of a stampede formation. For a flock to look realistic the animation needs to be of a high quality. With impostors it was impossible to change the animation, and we were restricted in the movements that could be presented, for example the movement around the root bone (hip) was minimised and so it was difficult to fully represent the extension and flexion of the spine during the dog's gallop, thus it looked very stiff. Speeding up the movement of the walking sheep (as they moved away from the dog) looked comical due to the foot-skate.

These drawbacks encouraged us to apply genuine animal motion to the animals. In this way, we can test whether humans prefer to see a certain type of motion - from earlier experiments we can see that a cow walk is recognised as a cow motion more so than the trot, however does this hold for a herd of cows? It is also interesting to see if the motions of different animals can be used to animate different 3D models. For example, could a cow herd look realistic using the motion of a walking horse? Also, how sensitive are we to the variation of motion for the same gait? For example, is a herd which is animated with two different walking gaits more appealing than one animated using only one walk gait? If this is not the case then it is welcome news for those that have limited resources for simulating flocks and herds of animals.

6.2 Perception of herd and flock animation

In order to validate the results of our gait perception experiments, we ran a series of experiments in which we investigated the perception of gait in a flocking scenario. We applied the real motion from the point-light walkers of the horse and cow to three 3D models (horse, cow and sheep). We tidied up the animations for each model by hand in order to be able to loop the motion seamlessly, and to minimise the hooves intersecting or floating above the ground plane. We used the horse and cow motions to animate the models as our previous experiments illustrated that they looked best when re-targeted motion was applied (no unsightly kinks in the mesh) and they were the easiest to clean up and use. Furthermore, these animals were the easiest to capture the motion from, as they were large enough to paint clear markers on, and very docile and obedient when required to walk or trot. To find out how effective they were for the animation of herds we asked the question: what sort of motions do people prefer to see on animals in a herd or flock?

6.2.1 Flocks and Herds

A few factors have to be taken into consideration when generating herds of animals such as the size of the frame, the size of the animals on the screen, and the size and formation of the herd or flock. As we wanted to test people's sensitivity to motions the animals had to be large enough for the movements to be visible when displayed on a 800x600 pixel frame. We positioned the camera (using perspective projection) so that the sizes of the animals in the foreground were about a third of the frame length and height. Other animals were spread out towards the background in such a way that there would be enough space for their legs to be seen throughout the video. To create a realistic herd, the animals could not be positioned in a neat order (like soldiers) so we moved them around to create a more natural formation. Finally, as the whole herd or flock is moving, we had to ensure that for any animals that went off screen they were replaced by new animals coming into view to ensure that a viewer was getting a consistent flow of information. Taking all the above into consideration we found that using 18 models for each herd gave the acceptable results.

To add variation to the motion of the herds, we created a slowed down and a speeded up version of each gait used. For example, for a cow walk we had the original version, and also a slower and faster version, resulting in 3 different speeds. We animated the models out of step in order make the flock look more realistic. The same procedure was applied to the horse walk motion, and the horse and cow trot motion. Each herd only contained one type of model i.e., only cows, horses or sheep. The animations for each herd consisted of the horse trot, cow trot, horse walk or cow walk only, and we also added a mixture of gaits where a herd was animated using both horse and cow walks (assigned 50-50) or horse and cow trots (assigned 50-50). Therefore, we had 6 different motions for each animal type. We added a combination of horse and cow motions to see if the variation in animation within a herd was preferred over herds that were animated using only one motion type.

We applied 3 different textures to the cow and horse models, and 2 textures to the sheep models, so that the animals did not look too much like clones. Screen shots of the videos can be see in Figures 6.4, 6.5 and 6.6.

6.3 Experiment 4

This experiment was divided into three blocks, because we only wanted to compare gaits between the same model. Participants were asked for preferences between pairs of videos, where each video was 4 seconds long and they were played one after the other. In order to compare across all the gaits, we had 21 comparisons including comparisons between the same gait, which was used as a control. The order in which the blocks were presented and in which the videos were played was random, with 3 repetitions of each video, resulting in a total of 63 comparisons for each animal model.

We had fifteen participants (12 male, 3 female) aged between 18 and 42 (ethical forms for

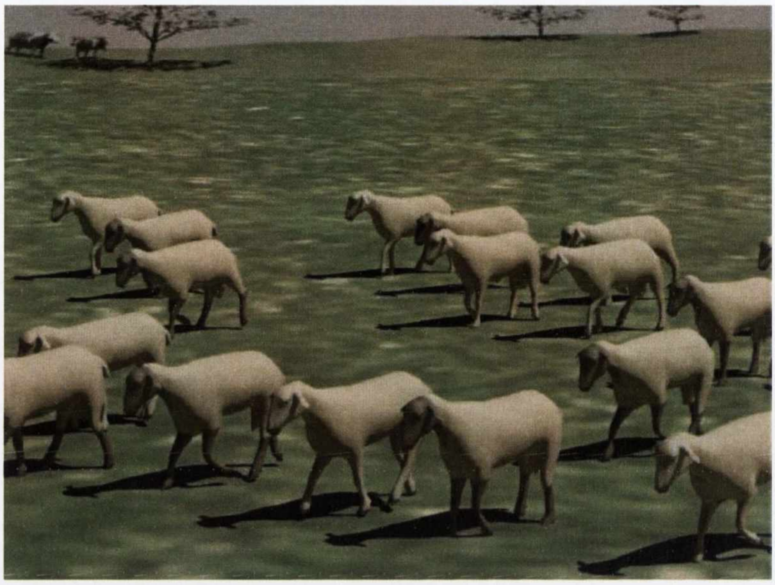

Figure 6.4: A flock of sheep. Only two textures are used, sheep with black faces and legs, and sheep with white faces and legs.

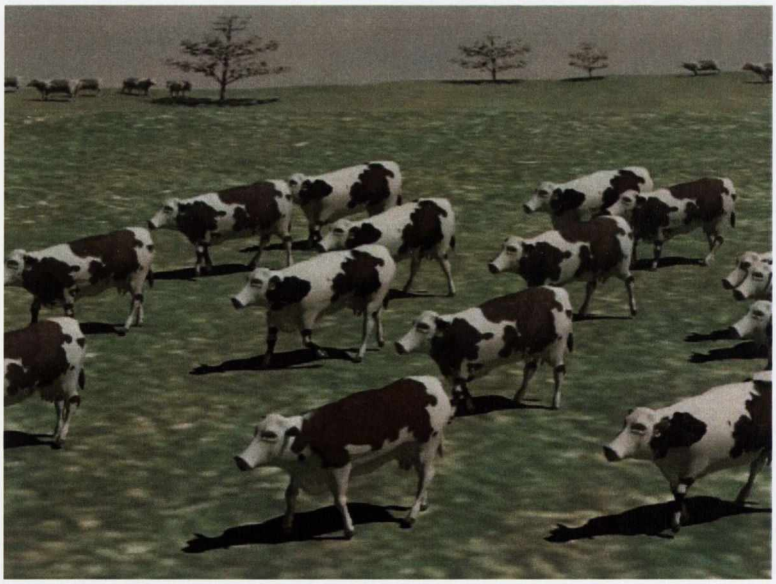

Figure 6.5: A herd of cows. It can be seen that three different textures are used to reduce the clone like appearance of the herd as a whole.

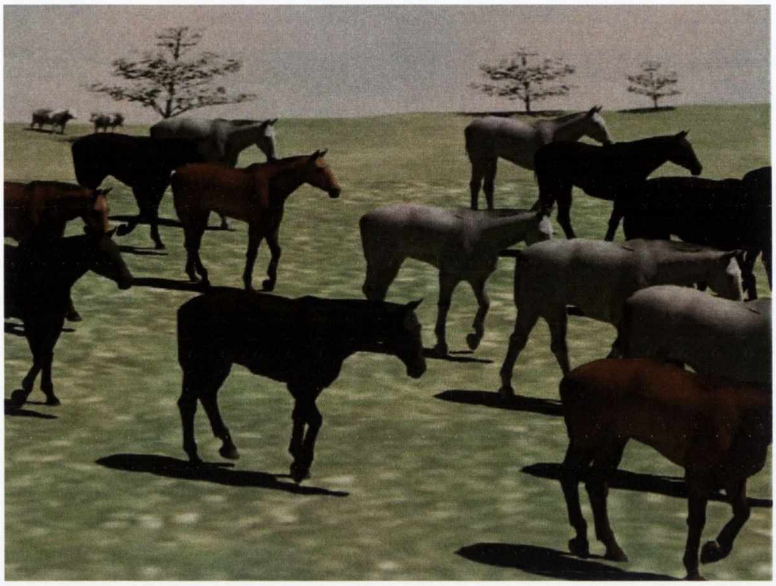

Figure 6.6: A herd of horses.

this experiment can be found in Appendix A.3.3). The participants saw one video followed by another and then they indicated their preference between the two animations seen. They were told to base their answer on the motions of the animals, and to disregard other factors (such as herd formation, or textures). The results are shown in Table 6.1. We were interested in which gaits scored the highest preferences for each of the animal models, and first looked at the results for the gaits according to whether they were walking or trotting. For the walk gaits, although on average the horse walk was selected most often as the preferred gait for all the models, and the cow model was selected most often when there was a walking animation, none of these results are statistically significant. This suggests that horse and cow walks may be interchangeable in practice in flocking applications. For example, on the sheep model there is no statistical difference between the preferences of the sheep flock animated with any of the walk gaits (horse, cow or mixed walk), therefore it can be suggested that they all look equally good, or bad, but in any case it reaffirms that the horse and cow walks can be interchanged.

For the trot gaits, a two factor ANalysis Of VAriance (ANOVA) with repeated measures was applied to the results. A main effect of animal gait was found ($F(2,42) = 4.0064, p < .02$). Post-hoc analysis using a standard Newman-Keuls test showed that the cow trot and the mixed trot (both horse and cow) were preferred significantly more than a horse trot. There was an interaction between model and gait ($F(4,84) = 3.9875, p < .01$). A herd of horses with a horse trot was preferred significantly more than a herd of cows with a horse trot. This

	MODEL		
Gait	Horse	Cow	Sheep
HW	51.85	62.59	54.81
HT	67.04	50.37	55.93
CW	48.15	50.00	45.56
CT	61.48	64.07	72.96
CWHW	54.81	54.44	53.33
CTHT	66.67	68.15	67.41

Table 6.1: Experiment 4, The average % of times a gait was picked for each model, where HW - horse walk, HT - horse trot, CW - cow walk, CT - cow trot, CWHW - both cow and horse walk, CTHT - both cow and horse trot.

also follows from earlier experiments where we found that the horse trot was distinctive and rarely mistaken as a cow motion. It is possible that in a herd, the horse trot motion does not look natural on the cow models. For the cow herd, the mixed trot was preferred significantly more times than the horse trot and the same is true for the sheep flock, where the mixed trot was preferred significantly more times than the horse trot.

When we look at the data for the trot gait in Table 6.1, we can see that on average the cow trot and the mixed trot have a similar score (66% and 67%). As with the walk gait, this implies that their combination did not add to the appeal of the animation of the herd even though both were perceived to be equally applicable to both models. This suggests that, as McDonnell et al. found for humans, motion variation between species may not be very important for realism of flocks. Videos showing horses and sheep trotting were preferred most of the time. This coincides with previous results that showed a high number of participants selecting fast motion to represent the horse and sheep. While the results also show that the cow herd was preferred with a trot rather than a walk, the difference is only minor. This suggest that, when we see cows on their own, we select the slow motion as the more characteristic gait, but in a herd the trot looks as realistic as the walk. The results can be seen in Figs. 6.7 and 6.8.

When we apply a two factor ANalysis Of VAriance (ANOVA) with repeated measures across the 6 gaits, there is a main effect of gait ($F(5,84) = 8.2517, p < .001$). Post-hoc analysis using standard Newman-Keuls tests showed that the cow trot and the mixed trot are significantly different to all the other gaits. We also found an interaction between model and gait ($F(10,168) = 2.6255, p < .001$). A sheep model with a cow trot animation was selected significantly more times than a sheep model with all the other gaits, except the mixed trot. A possible explanation for this is that, when we see a flock of sheep moving, it is usually because they are being moved to a particular place and so they are trotting or

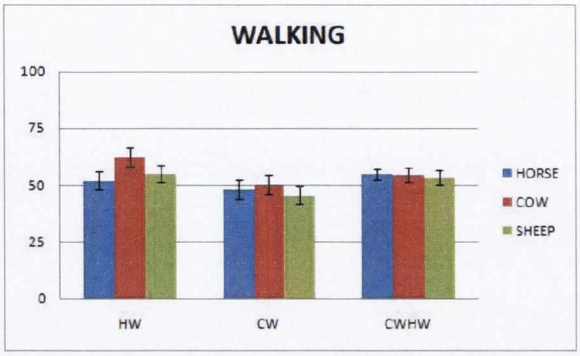

Figure 6.7: Experiment 4, Results for the walk gait for a herd of animals, e.g., second column shows cow herd walking with a horse walk. For all the trials where a cow herd with a horse walk was compared to some other gait, the horse walk was preferred 62.59% of the time.

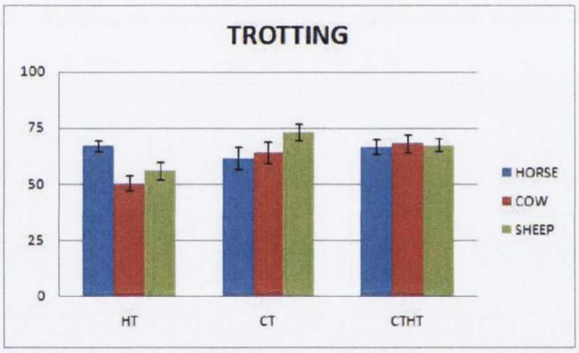

Figure 6.8: Experiment 4, Results for the trot gait for a herd of animals, e.g., first column shows a horse herd with a horse trot. For all the trials where the horse herd with a horse trot was compared to some other gait, the horse trot was preferred 67.04% of the time. For the sheep flock, the cow trot has the highest preference when it is compared against all the other gaits (indicated by the 6^{th} column).

galloping away from a dog or a person.

The trot gaits were picked as favourites for the horse herd, the distinctive horse trot and mixed trot scoring the highest in this case. Once again the difference in preference is 1%, showing that adding the variety of two animal gaits did not have an advantage. This is interesting because like cows, horses can also be seen to stand around in fields, but humans possibly have a preconception about how horses should move. The cow herd was liked best when doing the trot, much like the other two models. However, the cow model with a horse walk also scored relatively high - above 60%. This shows that the motion of cows is versatile and a herd can be convincingly animated using a walk or a trot motion.

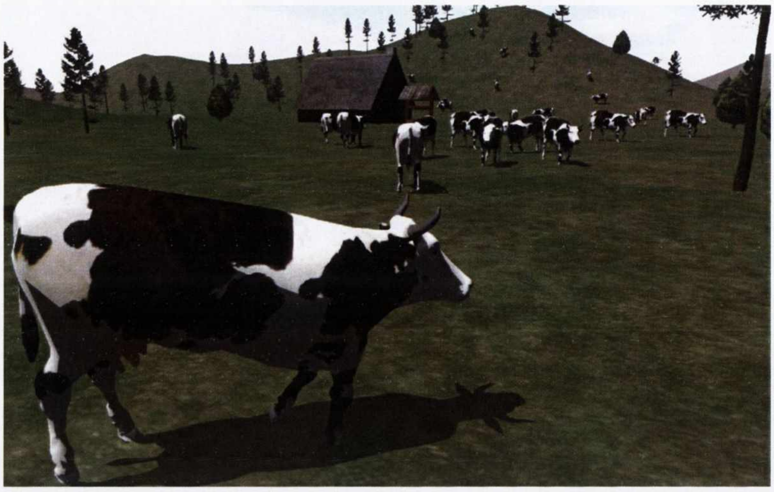

Figure 6.9: A herd of cows in a real-time system that includes static and dynamic obstacle avoidance and basic flocking functions.

6.4 Applications of our work

Finally, we selected two of the most popular animations on the cow model - the horse walk and the cow trot and imported these into a real-time system, Fig. 6.9. This system, based on the Ogre engine, has replaced the original proprietary Virtual Dublin system as used at the beginning of our research. All of the 48 cows perform static and dynamic obstacle avoidance and basic flocking functions based on a fuzzy logic model, that was implemented by our colleague Anton Gerdelan. The machine we are using has a Nvidia 8800 GTX graphics card and an Intel(R) Core(TM)2 Quad CPU at 2.40GHz. When all the cows are on the screen at one time, we can run the simulation at 226.3fps and when we include stencil shadowing and self-shadowing we achieve 20.1fps. The results from the experiment above can be used to help design how a group of animals should look during a shepherding and/or flocking simulation. It can be suggested that when a flock or a herd is in motion (during a steering behaviour) a trot motion looks more natural. However, if one wished to simply display a field that is scattered with animals then a walking animation looks more natural.

As simulating complex crowd behaviours using complex characters takes up a lot of memory, it is necessary to look into new ways of achieving required effects with the least amount of memory.

Chapter 7

Conclusions and Future Work

In this chapter, we provide an overview of the contributions of our work as well as suggestions for possible future work. We investigated the animation of quadrupedal animals while also providing a perceptual analysis of their motions. The techniques that we investigated and implemented, and the results of our experiments, provide useful insights and guidelines for improving the simulation of realistic herds and flocks of animals. There is also potential for future work, as the techniques used can be improved in many ways, and there is much more to be discovered through further perceptual experimentation.

7.1 Summary of Contributions

The work presented here is novel and of interest to both the animation and the perception communities. We were able to capture the 2D motion of real animals without incurring great costs and we applied the captured data to 3D animal models to use in perception experiments, and subsequently in a real-time flocking system. While much research has been conducted using human models and human motions, we contend that, as animals are important in many domains (e.g., entertainment, education and science), studying their motion and their behaviour in detail is a very useful contribution. In this thesis we use a technique which allowed us to recreate the motions of animals realistically and also studied the human perception of animal motion in depth.

More specifically, our contributions include:

- A published STAR at Eurographics 2008 - In collaboration with members of the EVASION group at Inria we published a comprehensive literature review, bringing together many different aspects of Quadruped Animation for the first time. We believe that this will help future students in researching the subject, as previously such literature was spread across a variety of different domains.

- A paper publication at Irish Chapter of Eurographics 2006 - "Animating Dolly" discusses the techniques we used to create sheep and dog impostors that were added to a real time system, where shepherding and flocking behaviours were implemented. This helped us to realise the importance of quality quadruped animation and its perception, and it was the catalyst for our subsequent research.

- An IK system - It became clear that pre-rendered impostor animations were too restrictive for larger scale simulations involving many animals with a variety of motions. Problems such as foot-skate are also very difficult to avoid with this method. For this reason we developed a simple automated system that animates quadrupeds using IK. A user is able to change the size of the legs or else pick from a set of pre-defined shapes representing a horse, cow, sheep and pig. However, as we were unable create natural looking motion using procedural methods alone, we decided to investigate ways of capturing natural motion directly, which could still be varied without introducing disturbing artefacts.

- Capturing the motion of animals - While many sophisticated 3D motion capture systems are available on the market today, using these on animals remains a difficult challenge in most practical situations. We implemented a methodology for the motion capture of animals that is relatively simple and cheap, where we painted the primary joints on the animals we wished to capture and then we filmed them using a video camera. In collaboration with our colleague Ian O'Connell, we built a simple tool that allows a user to track the motion of all the points marked on the animal. This is stored in a text file which can be easily used to visualise the motion in other software packages.

- Re-targeting captured motion to different 3D animal models - We were able to use the motion of the point-light walkers to animate 3D models. As each point-light defines the position of a joint, it was possible to use this information to place the joints of the skeleton of the 3D model into the corresponding positions. Using scaling, it was possible to apply the motion of a horse to a 3D model of a sheep and cow while still retaining the original movement.

- Perception Experiments (single animals) - To gain insights about whether we are sensitive to motions of different animals we designed experiments where we used genuine animal motion to animate point-light walkers and 3D models. From our novel experiments we have found that a horse's trot is distinct and can be spotted easily on point-light walkers and on different 3D models (horse, cow and sheep). Similarly the pig motion is easily recognised on both the point-light and 3D model. Horse and cow walk gaits are difficult to distinguish, suggesting that for this particular gait it is possible to re-target the motion of one animal to the model of another. We found that horses are easier to approach and film, thus the ability to use their walk motion to animate cows is a welcome result. We also found that gaits on 3D models are harder to recognise - another welcome result, as it suggests that the appearance and skinning of the model enables us to convince the audience that the animal is moving with its characteristic motion.

- Perception Experiments (herds of animals) - We found that animations of the sheep impostors were below standard and the resulting flock of sheep depicted was not satisfactory. As a result we created herds of 3D animals (horse, cow and sheep) using the captured motions of the horse and cow to see which animations people preferred. Once again, the horse trot was identified as distinctive, as it was the only motion that was preferred on a horse model. Overall the cow trot was the preferred gait for the cow and sheep models. We added variation to the herds by animating half the herd with a horse gait and the other half with a cow gait. We found that this variation does not have any advantage over using a single gait - a favourable result as it is easier to implement.
- Real-time simulation - It is possible to use our animated models in real-time systems such as games or interactive virtual environments. The results from our perception experiments give insights about what animations to use for simulations of herds of animals, thus minimising the number of resources required to store different animations. From above, we now know that it is possible to use the trot gait of just the cow, along with out-of-step motion to create an aesthetically pleasing herd animation.

We believe that our work has succeeded in bringing together computer science, psychology and veterinary science, thereby allowing all to potentially benefit from the research. In the field of computer graphics, we show a technique for the application of captured data to 3D models. Furthermore, the perception results can be used as a guide for designing the animation of animals and herds of animals, e.g., less time and resources are required if you only need one gait (rather than two or three) to animate a herd of animals realistically. For the perception community, we have highlighted some interesting questions that heretofore have not been addressed in detail e.g., we found that people viewed animals and humans with very similar eye-movement patterns and that factors other than the innate nature of an animal's motion affected perception e.g., the velocity, the terrain on which it walks, the background scene etc. While we have not directly investigated the application of our work in veterinary sciences, we and our collaborators in the UCD veterinary school believe that it can be further developed and used for training and education purposes of students as well as for the simulation and diagnosis of gait abnormalities.

7.2 Future Work

The work performed in this project can facilitate future work in different disciplines such as computer vision, computer animation, veterinary sciences, psychology and physiology.

7.2.1 Reflections on Experimental Design

In our experiments we only used one animal to capture the motion. In future it might be better to capture the motion of a few animals of the same type, for example to capture the walk motion of three or four different horses. However, this is a very big task as we found out from capturing the motion of just one animal. A collection of different animals would however introduce some variety and we would be able to test whether people are sensitive to this and whether it is necessary for the animation of herd of animals. We believe however, that such a variety is not needed, both on the physiological and perceptual grounds. As we explained in Section 3.2, Hildebrand found that variation in gait occurs between different species of animals rather than between the individuals [Hil65]. This implies that unless one wished to animate different species of horses, variation in animation may not be needed. McDonnell et al. found that using out-of-step motion was as effective a strategy for increasing variety in a crowd as applying different motions [MLD+08]. The same principle could apply to animals, so capturing gaits from animals of the same species may not be necessary for adding variations to herds or flocks.

We also think that capturing the motions of more varied animals would give interesting results, for example using animals such as cats and dogs, or elephants and giraffes. Would it be possible to animate large cat models, such as lions and tigers, using the motion of a domestic dog or cat?

It would also be interesting to see if it is possible to create a neutral motion, for example creating a neutral walk gait that is applicable to different models such as the horse, cow, sheep, zebra and maybe even an elephant. It may not be possible to find one base motion that could be used to generate all animals, however, there could be a set of basis animations that could be used to generate all possible animals.

In our herd experiments we used three different textures for the horse and the cow models and two textures for the sheep model. It would be useful to see how many textures are needed to create a realistic herd. For example, if you only need two varied textures to animate a herd

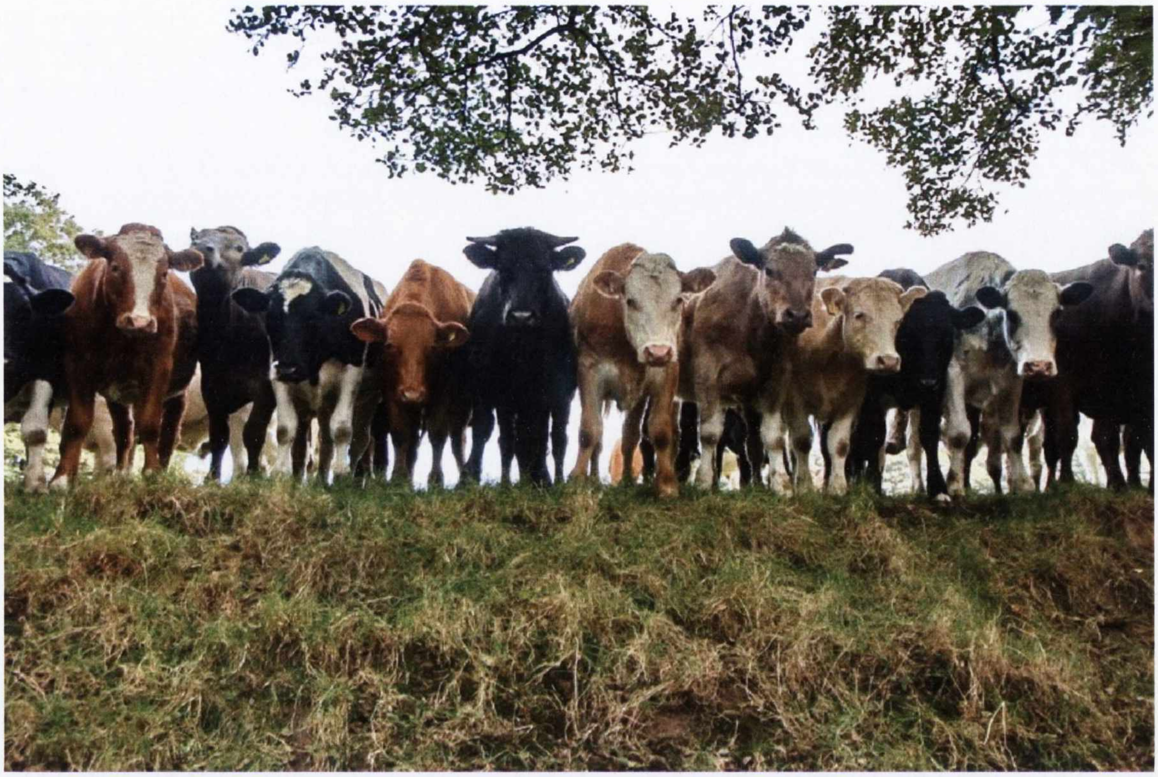

Figure 7.1: "Curious cows" courtesy of grannybuttons.com.

of cows it would save up a lot of memory. In addition, it would be interesting to see if it is better to use textures that are more patterned (e.g., black and white cows where there is a fine pattern) instead of using textures that are a single colour (e.g., only black, brown, red, beige, roan). From observation we have noticed that in fields cattle consist of a mixture of cows, solid colours and patterned, but we wonder whether people realise this without consciously observing it, see Fig. 7.1. Furthermore, shape variation of animals is something that is also worthy of investigation.

7.2.2 Automatic extraction of motion - computer vision

Our technique for the extraction of motion from videos is labour intensive. A user has to manually mark each dot on the animal for every frame of the animation. This can be time consuming where even a second of motion can take hours to track. There are existing techniques that extract the motion of the whole animal from the background of the video [FRDC04]. More specifically, this technique is able to extract the key position of the animal within a cycle of motion, so that a key-frame animator can place their model in certain poses for particular key-frames. For our work it would be useful to have a technique that would just extract the motion of the joints, as this would provide a means of automatically generating point-light walkers for different types of animals. Additionally, it would help with the actual animation of the skeleton for the 3D models, where the joints on the skeleton are controlled by the motion of the tracked points.

7.2.3 Automatic animation - computer animation

Animation of animals is very difficult, in particular for those animals with which we are not very familiar, or have no contact with. Therefore, video recordings of these animals are an essential way to learn about the movement. The extracted motion from these videos results in motion that is almost parallel to the real motion - the drawback is that it is 2D motion most of the time. Nevertheless applying this motion to 3D models directly makes the animation process quicker and easier. It enabled us to produce realistic motions without the help of expert animators.

As with all motion capture, the data used contains noise. In animation it is essential that animations can be looped seamlessly. A technique which is able to smooth out the animation without getting rid of the elementary movements would speed up the animation process. We tidied up the capture motions manually for our work, as we were concentrating on primarily using the animation for experiments and subsequent validation in herds and flocks. Our animation process did include some automation, for example the motion was scaled automatically to fit the different sized bones of the skeletons, but setting up of variables was manual and some tweaking was usually necessary for each model. A method whereby it would possible to apply motion data to different models automatically or with minimum user input would be welcome, as it would save a lot of time during the animation process.

The next step from this could be setting up of motion graphs (as used in films like Australia (20^{th} Century-Fox Film Corporation) and The Chronicles of Narnia (Disney)), where different clips can be joined to create long varied motions. Development of realistic transitions between different motion clips would also be an important contribution to the animation area.

7.2.4 Veterinary Sciences, Psychology and Physiology

The ability to track the motion of animals can give insights into the way that they move and help vets diagnose anomalities in other animals, e.g., by isolating the motions using the point-light motion, students could learn about the motion itself and which areas are prone to injury. Diagnosis of lameness in animals early is important as it allows for immediate treatment and avoids the animals having to be put down. It may be possible to design an "evolutionary" algorithm, using different models and shapes of animals to show how movement has changed as animals adapted to their changing environments.

Further studies of our perception of the motion can further ease the animation process and give a better understanding of how we expect the animals to move. In this work we shed some light on how in some situations people may be inclined to think cows are slow moving animals.

Animals are also used to help people with disabilities, for example dogs are used to hear, see and in therapy of autistic children, other animals such as horses are used for treatment of anger management [Rob09, Mal94]. It may be possible to improve the training processes for this by including studies involving virtual characters where we can find out more about people's perceptions of the behaviour of animals.

From the work of Martin [KSH99] and Kanwisher et al. [Mar07] we find that in neuroscience, animal pictures are commonly used as a visual stimulus for studying different brain activities using fMRI. For example, Martin reports that studies in the field suggest that a certain part of the brain responds more to animal stimulus, irrespective of the stimulus type (pictures, words and associated sounds), similarly, Kanwisher et al. state that particular areas in the brain react to a visual stimulus of human and animal faces.

Humans have always had a close relationship with animals, and further studies on their motions will continue to provide us with insights on how to accurately represent moving characters. In addition, it may also help us understand these animals better, and improve our relationship to them.

Appendix A

Appendix A

A.1 Ethical application forms for filming animals

In order to apply for the ethical approval for filming of animals we had to complete and include the following forms:

- "Animal Research Ethics Subcommittee Application Form: Application to the Animal Research Ethics Subcommittee UCD, for approval for the use of live animals in experiments"
- Include a description of the experiment - Farm Report

A.1.1 Farm report

Aims:
The aim of this project is to capture walking gaits of different farm animals using video cameras. The video data will provide information about the position of the animals' joints during

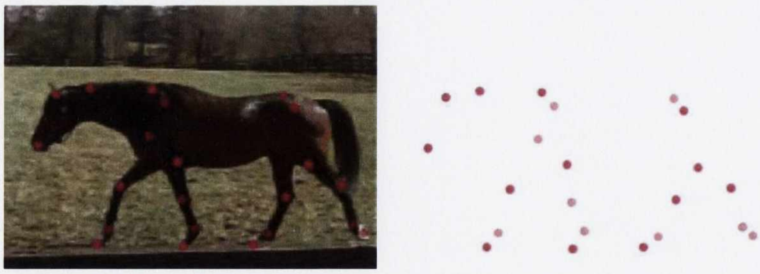

Figure A.1: (a)horse with markers (b)point-like walker

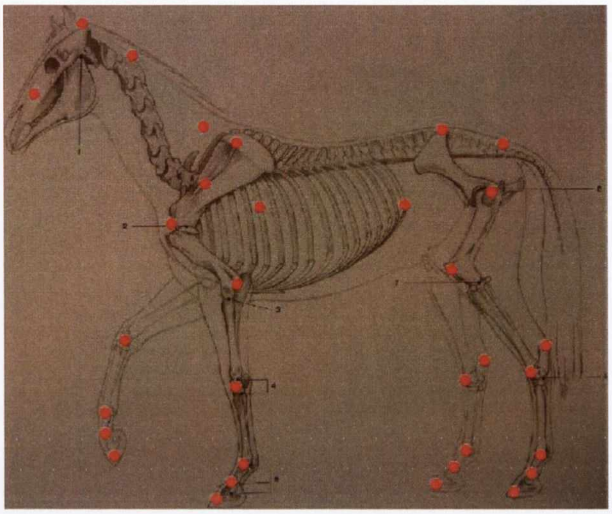

Figure A.2: markings on the animal

the walk cycle. Marking the joints on the animal prior to filming means more accurate results, as opposed to marking the joints after recording the video as in Fig. A.1 (a). By editing the video and extracting the information about the markers we can make point-like-walker videos which can be seen in Fig. A.1(b). By capturing the key motion differences between different species of animals, more realistic simulations can be produced.

Method:
In order to achieve the above

- animals need to be trained to walk in front of the camera, (preferably horizontally in a straight line). The environment should be as enclosed and controllable as possible to prevent them running away while leaving enough space for the video to be captured. It would also be helpful if the animals walked at a steady pace rather than behaving erratically (e.g., switching between not moving at all and galloping).

- It would also be useful capture other motions (e.g., trot, gallop), and to train the animals to transition naturally between motions, e.g., walking to trotting.

- markers need to be attached to all the major moving joints on both sides of the animal (arms, legs, head, neck and spine) including the inner legs as shown in Fig. A.2. Care needs to be taken that they won't bite at the markers etc (can use either luminous card paper stuck on with sellotape or clips or small circles painted on using a stencil).

Figure A.3: Plain Wall

- the animals need to walk in front of a plain background (e.g., the wall shown in Fig. A.3) and on solid ground so that the hooves can be clearly seen. Having a field and trees and other animals in the background can cause difficulties during the video post-processing.

- the animals need to be filmed using (1) a still camera on a tripod, (so the animal just walks in front of it) and (2) a moving camera (have the camera on wheels and move it as the animal walks)

- need about 5 walk cycles of each animal (horse, pony, cow, calf, pig, sheep, goat?).

A.2 Experiment 2

A.2.1 Metropolis launch in the Science Gallery

METROPOLIS: CROWD CONTROL[1] : website to call for participants to take part in Experiment 2
When does a crowd grab your attention? Can you replicate the unique ways people move? What does a crowd sound like? Can you tell the mood of a crowd?
Metropolis is a novel research project, which aims to answer these questions by combining expertise in computer graphics, engineering and cognitive neuroscience to create highly realistic virtual crowds.

At Metropolis: Crowd Control you can check out the demos of the research carried out so far and take part in real scientific studies conducted by Metropolis researchers. Your participation will help us answer more questions about virtual crowds and contribute to the

[1] http://www.sciencegallery.com/metropolis

development of the computer games industry, urban planning & pedestrian and traffic modelling, evacuation simulation and assistive technologies.

Find our more about the experiments you can take part in below.
Metropolis is supported by Science Foundation Ireland.
Prof Carol O' Sullivan, Prof Fiona Newell, Prof Henry Rice
Trinity College Dublin
Find out more at GV2

EMOTIONAL CROWDS
When we look at an angry mob or a group of football fans, we can easily sense the overall emotion of the crowd. In this study we are looking at how accurately we can rate crowds in terms of their emotions.

LOOK WHO'S TALKING
How quickly can you spot a particular person amongst a crowd of people? Can your other senses help you look? In this experiment, you will try to find your "friend" in a crowd of people accompanied by various sounds.

LOOK WHO'S TALKING TOO
Can you pick someone out of a crowd just from their voice? Our brain uses binaural (both ears) cues to figure out positional information about what we hear. This experiment will help determine how good we are at localising the position of a sound.

COUNTING CROWDS
Could you estimate how many people were at St Patrick's Festival or the Dublin City Marathon from just looking? Help us find out the point at which a group becomes a crowd by guessing the number of people in our virtual crowds.

WATCH THIS WAY
If we see others looking in a particular direction we tend to look that way too. We are investigating how many people in a crowd are necessary to look in a particular direction before we follow their gaze.

HIGHWAY PATROL
Traffic is something we're all too familiar with and deal with on a daily basis. In this experiment participants will see how quickly and accurately they can locate a moving object in a busy traffic scene.

NATURAL MOVERS
We can recognise our friends or ourselves from dot patterns of walking gaits alone. But, can we tell whether the motions of virtual humans were captured from real people, or if they were

the synthetic creations of a graphics wizard?

VARIETY - THE SPICE OF LIFE

Have you ever seen a real crowd full of identical clones? Usually, in the real world everyone around us looks unique and we are very good at recognising people from their individual appearance. In this experiment participants will help researchers find out what changes make people in a virtual crowd look as different as possible.

WHAT'S GOTTEN INTO EWE?

Could you spot a wolf in sheep's clothing? To make a virtual human crowd believable you need lots of variety. But what about herds of animated animals do you need the same amount of variety? Watch the motions of our virtual farm animals and tell us whether you think the animals are natural or not.

A.2.2 Irish Times article

Gallery invites public to explore virtual metropolis[2]

Author: Karlin Lillington

NET RESULTS: A virtual Dublin allows research on everything from computer games to urban planning

ALMOST EXACTLY two years ago, I wrote a story about Metropolis, an incredible and precise virtual Dublin under construction by researchers at Trinity College Dublin.

Supported by funding from Science Foundation Ireland, Metropolis was bringing together not just the expected computer scientists and engineers but also neuroscientists interested in how populating such a virtual world could give insight into how the human mind works and how people respond and act in crowd scenarios.

In addition, city planners and the Environmental Protection Agency were interested and involved in the project because of its ability to offer test cases for changing the urban landscape or responding to a natural disaster.

The problem for the average Dubliner, or Dublin visitor, though, was that they had no opportunity to see this virtual world themselves. Well, now anyone can come see a bit of virtual Dublin more precisely, virtual TCD at the latest event at Trinity College Dublin's Science Gallery, which, like the research project, is also called Metropolis.

You can don 3D glasses and view Front Square as an immersive three-dimensional world, or see how computing graduates have manipulated a model of Trinity on an X-Box to show different special effects such as lighting and atmospheric changes. Visitors can participate in research that will contribute to the overall project.

[2] http://www.irishtimes.com/newspaper/finance/2009/0327/1224243526806.html

There is also a fantastic game designed in co-operation with children at Dublin's Central Remedial Clinic that has enabled them to explore the virtual city themselves.

The goal of the Metropolis researchers has been to create the largest simulations of crowds ever achieved, using motion-capture technology for the animated inhabitants who stride through this recreated Dublin as if they own the place. The research is modelled around five main areas of design and investigation: motion, appearance, sound, behaviour and multisensory perception.

Prof Carol O'Sullivan, a computer scientist and one of the principal investigators for Metropolis, says the virtual Dublin is allowing graduate student researchers to conduct experiments they wouldn't be able to do otherwise on crowds, which will benefit computer games development and urban planning. Nine experiments are available for public users, ranging in time from five to 15 minutes.

One of the experiments in which visitors are able to participate involves them donning headphones and trying to pick out an individual in a crowded Front Square while people shout out "hellos" and "over heres" from various corners of the square.

The experiment looks at how much sound influences a viewer and would be very hard to replicate with real people in Front Square as you could hardly get enough volunteers for a big crowd for days on end, or have them shouting out all day long.

Other experiments ask users to gauge the number of cloned figures that are part of a large walking crowd, determine the emotions of a crowd, or pick out an object in busy traffic.

Perhaps the weirdest is a project in which the user tries to guess from body and leg movements which hidden animal is "dressed" as a sheep.

O'Sullivan says trying to use farm animals to get the motion-capture data was a particular challenge especially the woolly body and legs of the sheep. But I can report with certainty that looking at a virtual sheep that is walking like a cow is a bit disconcerting as well.

O'Sullivan says the data gathered from adult visitors will be used by the graduates as part of their research. Children are more than welcome to try their skills on the workstations as well, but their data will not be used.

A number of film screenings and talks are also a part of this week's event, which runs through this Sunday. Wall-E and The Lion, the Witch and the Wardrobe are on today and tomorrow.

There is a talk by the creators of Metropolis at lunchtime today at the gallery, while this evening features a discussion on creating crowd scenes by Paul Kanyuk, a technical director with famed Pixar Studios.

Tomorrow there is a talk on a similar topic with Rhythm and Hues, the company that produced crowd scenes for the Narnia films as well as a Night at the Museum and other films.

Tickets are free but very limited get more information and check availability with the Science Gallery or online, at www.sciencegallery.com/events.

Metropolis is now going into the final two years of its funding, says O'Sullivan, who adds that researchers are delighted to have the chance to show it off a bit to the public.

The next step is to start developing specific commercial applications, most likely in the areas of games development and urban planning.

A.2.3 Application for ethical approval

In order to have the general public take part in Experiment 2 we had to complete and include the following forms:

- "School of Psychology Research Ethics Committee Trinity College Dublin"
- Cover sheet for the experiment
- Consent form each participant signed

A.2.4 Cover sheet for the experiment

Title: What's got into ewe?

Researcher: Ljiljana Skrba

Instructions for running the study:
Participants will be instructed to sit in front of a computer monitor.
They will be asked to enter their id to begin the experiment.
They will see a series of videos of a virtual sheep model.
Different animal motions have been applied to the sheep model.
After each video their task will be to guess if the sheep is moving like a sheep or some other animal (horse, deer, cow, sheep, dog, pig).
The input their answer by pressing one of the keys on the keyboard (1-6).
After the experiment the participants need to be asked what they based their answers on, and if anything in particular looked good or bad.

About the experiment: Humans are very good at identifying different human walks. People are able to recognise their friends even just by looking at a set of moving points. We would like to find out if humans are as sensitive to recognising animal movements. Is it possible to realistically animate a 3d model of a sheep using captured gaits from different

farm animals? Will we notice that a sheep is walking like a cow or a horse? Is there anything in particular that we find wrong in some of the motions?
If we cannot see the difference between the various gaits it means we can retarget motion of one animal onto different models, saving time as customised walks do not need to be generated.
If we can tell the difference then we would like to know what the differences are between the gaits that humans notice most. This way during the animation process an animator can concentrate on these details to produce convincing motions.

A.2.5 Consent form

'What's gotten into Ewe?' experiment
Researcher: Ljiljana Skrba, Graphics, Vision and Visualisation Group, Department of Computer Science, Trinity College Dublin.

Information about the experiment and consent form

Thank you for your interest in this study. In this experiment, you will see a series of videos of virtual sheep. Movements of different animals have been applied to the sheep model. Your task will be to try and guess which animal's movement is on the sheep model. You can choose from a number of animals (horse, deer, cow, sheep, dog and pig). Try to concentrate on the movement of the sheep. Please try to respond as quickly and as accurately as possible by pressing the appropriate keys on the keyboard which will be indicated to you.

Please be aware that members of the public will be able to view the experiment while you are performing the task. However, your final results will be kept strictly confidential. Your data will be stored using your unique ID number and, as such, the experimenter will not be able to identify your data or link your personal details (i.e. name or address) with your data.

You are free to withdraw from this study at any time without any penalty. If you withdraw from this study you can still take part in any of the other ingoing experiments.

If you have any questions, please do not hesitate to ask the mediator.

By signing below:
1.The experimenter has adequately explained the task to you and you do not have any further questions.

2.You are aware that this experiment is conducted in a pubic setting where others can watch you perform the task. However, they will not be able to see your results. Your results will be available only to you upon request.
3.You are free to withdraw at any time without penalty
4.Your data from this study will be kept confidential and stored anonymously. We will not be able to identify your data or associate your data with your name; therefore, the Freedom of Information Act does not apply in this case.
5.Your data will be grouped with other participants and these group results will be presented in scientific papers and conferences.
Name:
Signature: Date:

A.3 Experiment 3 and 4

In order to be able to conduct experiments within our lab in Trinity College Dublin we had to complete and include the following forms:

- "School of Computer Science and Statistics, Research Ethical Approval Form"
- Cover sheet for Experiment 3
- Consent form for Experiment 3
- Consent form for Experiment 4

A.3.1 Cover sheet for the Experiment 3

Title: Gait Recognition

Researcher: Ljiljana Skrba

Instructions for running the study:
Participants will be instructed to sit in front of a computer monitor.
They will be asked to enter their id to begin the experiment.
They will see a series of videos of a virtual sheep model.
Different animal motions have been applied to the sheep model.
After each video their task will be to guess if the sheep is moving like a sheep or some other animal (horse, deer, cow, sheep, dog, pig).
The input their answer by pressing one of the keys on the keyboard (1-6).

After the experiment the participants need to be asked what they based their answers on, and if anything in particular looked good or bad.

About the experiment: Humans are very good at identifying different human walks. People are able to recognise their friends even just by looking at a set of moving points. We would like to find out if humans are as sensitive to recognising animal movements. Is it possible to realistically animate a 3d model of a sheep using captured gaits from different farm animals? Will we notice that a sheep is walking like a cow or a horse? Is there anything in particular that we find wrong in some of the motions?
If we cannot see the difference between the various gaits it means we can retarget motion of one animal onto different models, saving time as customised walks do not need to be generated.
If we can tell the difference then we would like to know what the differences are between the gaits that humans notice most. This way during the animation process an animator can concentrate on these details to produce convincing motions.

A.3.2 Consent form for Experiment 3

Experiment 1
Gait Recognition
Please read this document carefully before you decide to participate in this study
Aim: The aim of this study is to see if humans are good at recognising the difference between a horse's and a cow's gait.
What you will be asked to do in the study: You will be asked to remain in the GV2 lab. You will sit at a computer screen and you will be asked to watch a sequence of videos. After each video your task will be to guess whether the virtual animal on the screen moved with the gait of a cow or a horse. You will answer by pressing the corresponding button on the keyboard.
Before the experiment we will ask you to provide your gender and age, as well as information regarding if you suffer from epilepsy, and if you have any experience dealing with animals (horse riding, farming?). If you do not want to answer any question you are omitted to do so. There will be no recording of audio or any photographs taken that can be identifiable.
Time required: It is expected that the total time required will be around 10min, depending on the time it takes to decide on your answer.
Benefits/Compensation: There will be no financial compensation to you for your participation.
Confidentiality: Your identity will be kept confidential. This consent form, if signed by

you, will be kept locked separately from the data we collect during the experiment. It will not be possible to link the data we collect to your name. The data we collect will be analysed together with the data collected from other participants, and generalised results and conclusions will be drawn from these experiments will be submitted for publication at conferences and/or scientific journals. Your name will not be used in any report or article.

Voluntary participation: Your participation in this study is voluntary. There is no penalty for not participating.

Right to withdraw from the study: You have the right to withdraw from the study at any time without consequence.

Whom to contact if you have questions about the study: Ljiljana Skrba

I understand that my participation in an experiment is entirely voluntary.

I am informed about the purpose of the study.

I understand that I can withdraw from the study at any time without having to give a reason.

I have been give the option of omitting questions that I do not wish to answer

All the information collected will remain confidential.

I agree that my data is used for scientific purposes and therefore have no objection that my data is published in scientific publications in a way that does not reveal my identity.

Neither my participation in the experiment, nor the results obtained, have any bearing on my academic record. If I withdraw, this has no bearing on my academic record.

On request, I can be given a brief explanation after the study.

Date:

Signature:

Questionnaire:

Gender: Female Male

Age:

Profession:

Do you have any experience with animals?(Horse riding, Working on a farm)

Have you or anyone in your family suffered from epilepsy?

A.3.3 Consent form for Experiment 4

Experiment 1

Gait and Appearance Variation in Herds of Animals

Please read this document carefully before you decide to participate in this study

Aim: The aim of this study is to determine what differences humans notice between herds of animals. We would also like to determine how important variation is in creating more realistic or interesting herds.

What you will be asked to do in the study: You will be asked to remain in the GV2 lab. You will sit at a computer screen and you will be asked to watch a sequence of videos. A pair of videos will be played and after the second video finishes you will be asked if you preferred the animations in the first video or the second one.

Before the experiment we will ask you to provide your gender and age, as well as information regarding if you suffer from epilepsy, and if you have any experience dealing with animals (horse riding, farming?). If you do not want to answer any question you are omitted to do so. There will be no recording of audio or any photographs taken that can be identifiable.

Time required: It is expected that the total time required will be around 30min, depending on the time it takes to decide on your answer.

Benefits/Compensation: As a token of our appreciation you will be given a national book voucher.

Confidentiality: Your identity will be kept confidential. This consent form, if signed by you, will be kept locked separately from the data we collect during the experiment. It will not be possible to link the data we collect to your name. The data we collect will be analysed together with the data collected from other participants, and generalised results and conclusions will be drawn from these experiments will be submitted for publication at conferences and/or scientific journals. Your name will not be used in any report or article.

Voluntary participation: Your participation in this study is voluntary. There is no penalty for not participating.

Right to withdraw from the study: You have the right to withdraw from the study at any time without consequence.

Whom to contact if you have questions about the study: Ljiljana Skrba

I understand that my participation in an experiment is entirely voluntary.

I am informed about the purpose of the study.

I understand that I can withdraw from the study at any time without having to give a reason.

I have been give the option of omitting questions that I do not wish to answer

To acknowledge my contribution of time and effort to the experiment a national book voucher will be given to me on completion of the experiment. This reward is not related to my performance in the study.

All the information collected will remain confidential.

I agree that my data is used for scientific purposes and therefore have no objection that my data is published in scientific publications in a way that does not reveal my identity.

Neither my participation in the experiment, nor the results obtained, have any bearing on

my academic record. If I withdraw, this has no bearing on my academic record.
On request, I can be given a brief explanation after the study.
I confirm that I am 18 years or older and competent to supply consent.
I understand that if I or anyone in my family have a history of epilepsy then I am proceeding at my own risk

Date:

Signature:

Questionnaire:

Gender: Female Male

Age:

Profession:

Do you have any experience with animals?(Horse riding, Working on a farm)

Have you or anyone in your family suffered from epilepsy?

Bibliography

[ABHM09] Paul Aichele, Byron Bashforth, Laura Hainke, and Bob Moyer. It's good to be alpha. In *SIGGRAPH '09: SIGGRAPH 2009: Talks*, pages 1–1, New York, NY, USA, 2009. ACM.

[AHLD07] Gregoire Aujay, Franck Hetroy, Francis Lazarus, and Christine Depraz. Harmonic skeleton for realistic character animation. In *SCA '07: Proceedings of the 2007 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 151–160, Aire-la-Ville, Switzerland, Switzerland, 2007. Eurographics Association.

[AHSS04] Aseem Agarwala, Aaron Hertzmann, David H. Salesin, and Steven M. Seitz. Keyframe-based tracking for rotoscoping and animation. In *SIGGRAPH '04: ACM SIGGRAPH 2004 Papers*, pages 584–591, New York, NY, USA, 2004. ACM.

[AJ83] R. McNeill Alexander and A.S. Jayes. A dynamic similarity hypothesis for the gaits of quadrupedal mammals. *Journal of Zoology*, 201:135–152, 1983.

[Ale68] R.McNeill Alexander. *Animal Mechanics.* Sidgwick & Jackson, 1968.

[Ale76] R.M. Alexander. Estimates of speeds of dinosaurs. *Nature*, 261:129–130, May 1976.

[Ale84] R.McNeill Alexander. The gaits of bipedal and quadrupedal animals. In *The International Journal of Robotics Research*, volume 3, pages 49–59. SAGE Publications, 1984.

[Ale96] R.McNeill Alexander. *Optima for Animals.* Princeton University Press, 1996.

[Ale03] R.McNeill Alexander. *Principles of Locomotion.* Princeton University Press, 2003.

[BG95] Bruce M. Blumberg and Tinsley A. Galyean. Multi-level direction of autonomous creatures for real-time virtual environments. In *SIGGRAPH '95: Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 47–54, New York, NY, USA, 1995. ACM.

[Bie08] Alain Bielik. Prince caspian: A darker and grander narnia. *VFXWorlds*, May, 2008.

[BK05] Gladimir Baranoski and Aravind Krishnaswamy. Simulation of light interaction with human skin. In *Eurographics 2005, Tutorial notes*, 2005.

[BMH98] David C. Brogan, Ronald A. Metoyer, and Jessica K. Hodgins. Dynamically simulated characters in virtual environments. *IEEE Comput. Graph. Appl.*, 18(5):58–69, 1998.

[Bos] Boston dynamics: www.bostondynamics.com.

[BP07] Ilya Baran and Jovan Popović. Automatic rigging and animation of 3d characters. *ACM Trans. Graph. (Proceedings of SIGGRAPH)*, 26(3):72, 2007.

[BPB06] J.F. Burn, B. Portus, and C. Brockington. The effect of speed and gradient on hyperextension of the equine carpus. In *Veterinary Journal*, volume 171, pages 169–171, 2006.

[Bur03] Jonathan Burt. *Animals in Film.* Reaktion Books, 2003.

[BW95] Armin Bruderlin and Lance Williams. Motion signal processing. In *SIGGRAPH '95: Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 97–104, New York, NY, USA, 1995. ACM.

[Cap05] Heather Eudora Caprette. Visualisation of canine gait using optical motion capture animation: a study of its effectiveness at helping students indentify lameness. 2005.

[CCC+05] Janko Calic, Neill Campbell, Andrew Calway, Majid Mirmehdi, Tilo Burghardt, Sion Hannuna, Changming Kong, Sarah Porter, Nishan Canagarajah, and David Bull. Towards intelligent content based retrieval of wildlife videos. In *Proceedings of the 6th International Workshop on Image Analysis for Multimedia Interactive Services (WIAMIS 2005)*, April 2005.

[CDNDMN85] M.S. Charles D. Newton, D.V.M. and V.M.D David M. Nunamaker. *Textbook of Small Animal Orthopedics*. J.B. Lippincott Company, 1985. Normal and Abnormal Gait (Ch.91).

[CK77] J. Cutting and L. Kozlowski. Recognising friends by their walk: gait perception without familiarity cues. In *Bulletin of the Psychonomic Society*, volume 9, pages 353–356, 1977.

[Cla97] Hilary Clayton. Biomechanics of the horse. 1997. Presented by Dr. Hilary Clayton at the United States Dressage Federation Convention, December 6, 1997.

[Cla98] Hilary Clayton. Gait analysis: Past, present and future. 1998. The McPhail Chair Lecture given by Dr. Clayton at the annual Michigan Veterinary Conference on January 23rd, 1998.

[Cla01] Hilary Clayton. Gaits and movements of dressage horses. 2001. Presentations by Dr. Clayton at the Tufts Animal Expo. Boston, Massachusetts. 2001.

[CLDP95] J.D. Cohen, M.C. Lin, D.Manocha, and M. K. Ponamgi. I-collide: An interactive and exact collision detection system for large-scale environments. In *Symposium on Interactive 3D Graphics*, pages 189–196, 1995.

[Cra82] Donald Crafton. *Before Mickey - The Animated Film 1898 - 1928*. The MIT Press, 1982.

[Cra99] Michael Crandol. The history of animation: Advantages and disadvantages of the studio system in the production of art form. 1999.

[dAST+08] Edilson de Aguiar, Carsten Stoll, Christian Theobalt, Naveed Ahmed, Hans-Peter Seidel, and Sebastian Thrun. Performance capture from sparse multi-view video. In *ACM SIGGRAPH*, page To appear, 2008.

[dATSS07] Edilson de Aguiar, Christian Theobalt, Carsten Stoll, and Hans-Peter Seidel. Marker-less deformable mesh tracking for human shape and motion capture. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages XX–XX, Minneapolis, USA, June 2007. IEEE, IEEE.

[dATTS08] Edilson de Aguiar, Christian Theobalt, Sebastian Thrun, and Hans-Peter Seidel. Automated conversion of mesh animations into skeleton-based animations. *Computer Graphics Forum (Proc. Eurographics)*, 27(2):To appear, 2008.

[DHOO05] Simon Dobbyn, John Hamill, Keith O'Conor, and Carol O'Sullivan. Geopostors: a real-time geometry / impostor crowd rendering system. In *Proc. Symp. on Interactive 3D graphics and games (SI3D)*, pages 95–102, 2005.

[DSP06] Kevin G. Der, Robert W. Sumner, and Jovan Popovic. Inverse kinematics for reduced deformable models. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Papers*, pages 1174–1179, New York, NY, USA, 2006. ACM Press.

[EKB95] Dimitri Metaxas Evangelos Kokkevis and Normal I. Badler. Autonomous animation and control of four-legged animals. In *Graphics Interface '95*, pages 10–17. ACM Press / ACM SIGGRAPH, 1995.

[FRDC04] Laurent Favreau, Lionel Reveret, Christine Depraz, and Marie-Paule Cani. Animal gaits from video. In *SCA '04: Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 277–286, Aire-la-Ville, Switzerland, Switzerland, 2004. Eurographics Association.

[Fro09] Arran Frood. Running with greyhounds, horses, cheetahs... and dinosaurs. 2009.

[FTF03] J. Francik and K. Trybicka-Francik. A framework for program control of animation of human and animal characters. *Studia Informatica, Silesian University of Technology*, 24(4):55–65, 2003.

[GAD+07] Philippe Gluckman, Guillaume Aretos, David Doepp, Scott Peterson, Jason Waltman, Lucia Modesto, Larry Cutler, and Bill Seneshen. From "shrek" to "shrek the third": evolution of cg characters in the "shrek" films. *ACM SIGGRAPH courses*, 2007.

[Gam74] P. P. Gambaryan. *How Mammals Run: Anatomical Adaptations.* John Wiley and Sons, 1974.

[GCT03] D.P. Gibson, N.W. Campbell, and B.T. Thomas. Quadruped gait analysis using sparse motion information. In *Image Processing, 2003. ICIP 2003. Proceedings. 2003 International Conference on*, volume 3, pages 333–6, 2003.

[Gil02] Robert L. Gillette. Caring for canines, 2002.

[Gir87] Michael Girard. Interactive design of 3-d computer-animated legged animal motion. In *SI3D '86: Proceedings of the 1986 workshop on Interactive 3D graphics*, pages 131–150, New York, NY, USA, 1987. ACM Press.

[GLH04] Dragos Golubovic, Bo Li, and Huosheng Hu. A hybrid software platform for sony aibo robots. In *RoboCup 2003: Robot Soccer World Cup VII*, pages 478–486. Springer Berlin / Heidelberg, 2004.

[GMC+05] Philippe Gluckman, Denise Minter, Kendal Chronkhite, Cassidy Curtis, Milana Huang, Rob Vogt, and Scott Singer. Madagascar: bringing a new visual style to the screen. *ACM SIGGRAPH Courses*, 2005.

[GODC05] David Gibson, David Oziem, Colin Dalton, and Neill Campbell. Capture and synthesis of insect motion. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 39–48, 2005.

[HB94] Jessica K. Hodgins and David Brogan. Robot herds: Group behaviors for systems with significant dynamics. In *Proc. Artificial Life IV*, pages 319–324, 1994.

[HBaANS04] S. Hochstein, A. Barlasov, O. Herschler adn A. Nitzan, and S. Shneor. Rapid vision is holistic. In *Journal of Vision*, volume 4, 2004.

[HCG05] Sion Hannuna, Neill Campbell, and David Gibson. Identifying quadruped gait in wildlife video. In *International Conference on Image Processing*, pages 713–716. IEEE, September 2005.

[HCN+09] Gabor Horvath, Adelinda Csapo, Annamaria Nyeste, Balazs Gerics, Gabor Csorbe, and Gyorgy Kriska. Erroneous quadruped walking depictions in natural history museums. *Current Biology*, 19:R61 – R62, 2009.

[HDK+06] Brad Hiebert, Jubin Dave, Tae-Yong Kim, Ivan Neulander, Hans Rijpkema, and Will Telford. The chronicles of narnia: the lion, the crowds and rhythm and hues. *ACM SIGGRAPH Courses*, 2006.

[HG01] Milton Hildebrand and George Goslow. *Analysis of Vertebrate Structure*. 5 edition, 2001.

[HG07] Rachel Heck and Michael Gleicher. Parametric motion graphs. In *I3D '07: Proceedings of the 2007 symposium on Interactive 3D graphics and games*, pages 129–136, New York, NY, USA, 2007. ACM.

[HHO04] Sarah Howlett, John Hamill, and Carol O'Sullivan. An experimental approach to predicting saliency for simplified polygonal models. In *ACM SIGGRAPH Symposium on Applied Perception in Graphics and Visualization (APGV'04)*, pages pp 57–64, 2004.

[HHO05] Sarah Howlett, John Hamill, and Carol O'Sullivan. Predicting and evaluating saliency for simplified polygonal models. In *ACM Transactions on Applied Perception*, pages pp. 286 – 308, 2005.

[Hil65] Milton Hildebrand. Symmetrical gaits of horses. 150:701–8, 1965.

[Hil89] Milton Hildebrand. The quadrupedal gaits of vertebrates. 39:766–5, 1989.

[HO03] John Hamill and Carol O'Sullivan. Virtual dublin - a framework for real-time urban simulation. pages 221 – 225, 2003.

[HOT98] Jessica K. Hodgins, James F. O'Brien, and Jack Tumblin. Perception of human motion with different geometric models. *IEEE Transactions on Visualization and Computer Graphics*, 4(4):307–316, 1998.

[How99] D. Howell. *Fundamental Statistics for Behavioral Sciences.* Duxbury Press, 1999.

[HRE+08] Chris Hecker, Bernd Raabe, Ryan W. Enslow, John de Weese, Jordan Maynard, and Kees van Prooijen. Real-time motion retargeting to highly varied user-created morphologies. In *SIGGRAPH '08*, New York, NY, USA, 2008. ACM.

[ICE] Proceedings of the international conference on equine locomotion. Elsevier.

[JA78] A. S. Jayes and R. M. Alexander. Mechanics of locomotion of dogs (canis familiaris) and sheep (ovis aries). *J Zool*, pages 289–308, 1978.

[JEB01] James J. Little Jeffrey E. Boyd. Phase models in gait analysis. 2001.

[Joh73] G. Johansson. Visual perception of biological motion and a model for its analysis. In *Perception and Psychophysics*, volume 14, pages 201–211, 1973.

[JPFPB01] Mark S. Nixson Jeff P. Foster and Adam Prugel-Bennett. Recognising movement and gait by masking functions. 2001.

[JT05] Doug L. James and Christopher D. Twigg. Skinning mesh animations. *ACM Trans. Graph.*, 24(3):399–407, 2005.

[Kan09] Paul Kanyuk. Rivers of rodents & lots of bots - crowd simulation at pixar, 2009.

[KBF+09] Carsten Kolve, Dan Bethell, Rogier Fransen, Johannes Saam, and Premamurti Paetsch. Venomous cattle for australia. In *SIGGRAPH '09: SIGGRAPH 2009: Talks*, pages 1–1, New York, NY, USA, 2009. ACM.

[kHD03] kevin H. Deal. *Wildlife & natural resource management.* 2003.

[KRFC07] Paul Kry, Lionel Reveret, Francois Faure, and Marie-Paule Cani. Modal locomotion: Controlling passive elastic dynamics. In *Siggraph Technical Sketches and Applications.* ACM Press, 2007.

[KSH99] Nancy Kanwisher, Damian Stanley, and Alison Harris. The fusiform face area is selective for faces but not animals. pages 183–187, 1999.

[KT03] Sagi Katz and Ayellet Tal. Hierarchical mesh decomposition using fuzzy clustering and cuts. *ACM Trans. Graph. (Proceedings of SIGGRAPH)*, 22(3):954–961, 2003.

[Kuh03] Jennifer Kuhnel. Rigging a horse and rider: Simulating predictable and repetitive movement of the rider. Master's thesis, Texas A & M University, 2003.

[Las87] John Lasseter. Principles of traditional animation applied to 3d computer animation. In *SIGGRAPH '87: Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, pages 35 – 44, 1987.

[Law] Sian E.M. Lawson. Equine mechanics: www.equinemechanics.com.

[LBS+04] Jyh-Ming Lien, O. Burchan Bayazit, Ross T. Sowell, Samuel Rodriguez, and Nancy M. Amato. Shepherding behaviours. In *Proc. IEEE Int. Conf. Robot. Autom.*, pages 4159–4164, 2004.

[LKA06] Jyh-Ming Lien, John Keyser, and Nancy Amato. Simultaneous shape decomposition and skeletonization. In *ACM Symposium on Solid and Physical Modeling*, pages 219–228, 2006.

[LRMA04] Jyh-Ming Lien, Samuel Rodriguez, Jean-Phillipe Malric, and Nancy M. Amato. Shepherding behaviours with multiple shepherds, 2004.

[LvdPF00] Joseph Laszlo, Michiel van de Panne, and Eugene Fiume. Interactive control for physically-based animation. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 201–208, New York, NY, USA, 2000. ACM Press/Addison-Wesley Publishing Co.

[Mal94] Gerald P. Mallon. Cow as co-therapist: Utilization of farm animals as therapeutic aides with children in residential treatment. In *Child and Adolescent Social Work Journal*, pages 455–474, 1994.

[Mar07] Alex Martin. The representation of object concepts in the brain. pages 25 - 45, 2007.

[MB57] Eadweard Muybridge and Lewis S. Brown. *Animals in Motion.* Dover Publications, 1957.

[MC03] J. Martineau and S. Cochin. Visual perception in children: human, animal and virtual movement activates different cortical areas. 51(1):37–44, 2003.

[MDMT+04] Laurent Moccozet, Fabien Dellas, Nadia Magnenat-Thalmann, Sivlia Biasotti, Michela Mortara, Bianca Falcidieno, Patrick Min, and Remco Veltkamp. Animatable human body model reconstruction from 3d scan data using templates. In *Workshop on Modeling and Motion Capture Techniques for Virtual Environments (CAPTECH'04)*, pages 73–79, 2004.

[MDO05] Rachel McDonnell, Simon Dobbyn, and Carol O'Sullivan. Lod human representations: A comparative study. In *Proceedings of the First International Workshop on Crowd Simulation*, pages 101 - 115, 2005.

[Mil98] Mike Milne. Entertaining the future: Real animators don't rotoscope, or do they? *SIGGRAPH Comput. Graph.*, 32(2):5–7, 1998.

[MJH+08] Rachel McDonnell, Sophie Joerg, Jessica K. Hodgins, Fiona Newell, and Carol OSullivan. Evaluating the effect of motion and body shape on the perceived sex of virtual characters. In *ACM Transactions on Applied Perception,*, volume 5, pages 21:1 - 21:14, 2008.

[MJM+08] Rachel McDonnell, Sophie Joerg, Joanna McHugh, Fiona Newell, and Carol OSullivan. Evaluating the emotional content of human motions on real and virtual characters. In *APGV 08: Proceedings of the symposium on Applied Perception in Graphics and Visualisation*, pages 67–74, 2008.

[ML05] S.A Marsland and R.J Lapeer. Physics-based animation of a trotting horse in a virtual environment. In *Proceedings of the 9th International Conference on Information Visualisation (IV'05)*, pages 398–403, 2005.

[MLD+08] Rachel McDonnell, Micheal Larkin, Simon Dobbyn, Steven Collins, and Carol O'Sullivan. Clone attack! perception of crowd variety. In *SIGGRAPH '08: ACM SIGGRAPH 2008 papers*, pages 1–8, New York, NY, USA, 2008. ACM.

[MLH+09] Rachel McDonnell, Micheal Larkin, Benjamin Hernandez, Isaac Rudomin, and Carol O'Sullivan. Eye-catching crowds: saliency based selective variation. *ACM Trans. Graph.*, 28(3):1–10, 2009.

[MMMD09] Meredith McLendon, Ann McNamara, Tim McLaughlin, and Ravindra Dwivedi. Connecting the dots: discovering what's important for creature motion. In *SIGGRAPH '09: SIGGRAPH 2009: Talks*, pages 1–1, New York, NY, USA, 2009. ACM.

[MP98] Maja J. Mataric and Marc Pomplun. What do people look at when watching human movement?, 1998.

[MW93] G. Mather and S. West. Recognition of animal locomotion from dynamic point-light displays. *Perception*, 22:759–766, 1993.

[QCVT06] Heinrich H. Bulthoff Quoc C. Vuong, Andries F. Hof and Ian M. Thornton. An advantage for detecting dynamic targets in natural scenes. volume 6, pages 87–96, 2006.

[RBR+] Lionel Reveret, Vincent Bels, Sabine Renous, Franck Multon, and Pierre-Paul Vidal. Projet anr - kameleon.

[Rey87] Craig W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *SIGGRAPH '87: Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, pages 25–34, New York, NY, USA, 1987. ACM Press.

[Rey99] Craig W. Reynolds. Steering behaviors for autonomous characters. In *Proc. Game Developers Conf.*, pages 763–782, 1999.

[RF03] Deva Ramanan and D. A. Forsyth. Using temporal coherence to build models of animals. In *ICCV '03: Proceedings of the Ninth IEEE International Conference on Computer Vision*, page 338, Washington, DC, USA, 2003. IEEE Computer Society.

[RFDC05] Lionel Reveret, Laurent Favreau, Christine Depraz, and Marie-Paule Cani. Morphable model of quadrupeds skeletons for animating 3d animals. In *SCA*

'05: Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation, pages 135–142, New York, NY, USA, 2005. ACM.

[RH91a] Marc H. Raibert and Jessica K. Hodgins. Animation of dynamic legged locomotion. In *SIGGRAPH '91: Proceedings of the 18th annual conference on Computer graphics and interactive techniques*, pages 349–358, New York, NY, USA, 1991. ACM Press.

[RH91b] Marc H. Raibert and Jessica K. Hodgins. On the run. *ACM SIGGRAPH Electronic Theater*, 1991.

[Rij09] Hans Rijpkema. The evolution of crowds - inside the crowd pipeline at rhythm and hues studios, 2009.

[Rob09] Barbara Robertson. Wonder dogs - new jobs tap into their many talents. *BARK magazine*, pages 74–79, 2009.

[Sey07] M. Seymour. Spi is legend. 2007.

[SF06] Andras Szunyoghy and Gyorgy Feher. *Anatomia Scuola di Disegno*. Konemann, 2006.

[SH05] David M. Shackleton and Alton S. Harestad. *Grzimek's Animal Life Encyclopedia*. Thomson Gale, 2005.

[Sil04] Matt Silverman. *History of Rotoscoping*. Phoenix Editorial: http://www.measurand.com/motion-capture-resources/motion-capture-history-rotoscope.htm, 2004.

[Smi02] Lindsay I. Smith. A tutorial on principal components analysis, 2002.

[SMT03] Hyewon Seo and Nadia Magnenat-Thalmann. An automatic modeling of human bodies from sizing parameters. In *ACM SIGGRAPH Symposium on Interactive 3D Graphics*, pages 19–26, 2003.

[SO09] Ljiljana Skrba and Carol O'Sullivan. Human perception of quadruped motion. In *APGV '09: Proceedings of the 6th Symposium on Applied Perception in Graphics and Visualization*, pages 130–130, New York, NY, USA, 2009. ACM.

[SOO08] Ljiljana Skrba, Ian O'Connell, and Carol O'Sullivan. Eye-tracking dynamic scenes with humans and animals. In *APGV '08: Proceedings of the 5th symposium on Applied perception in graphics and visualization*, pages 199–199, New York, NY, USA, 2008. ACM.

[SRH+08] Ljiljana Skrba, Lionel Reveret, Franck Hetroy, Marie-Paule Cani, and Carol O'Sullivan. Quadruped animation. In *Eurographics State-of-the-Art Report (EG-STAR)*, pages 1–17. Eurographics Association, Eurographics Association, 2008.

[Stu94] David J. Sturman. Character motion systems: A brief history of motion capture for computer character animation. *ACM SIGGRAPH Courses*, 1994.

[SWG02] Maryann Simmons, Jane Wilhelms, and Allen Van Gelder. Model-based reconstruction for creature animation. In *ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 139–146, New York, NY, USA, 2002. ACM Press.

[SY07] Scott Schaefer and Can Yuksel. Example-based skeleton extraction. In *ACM SIGGRAPH/Eurographics Symposium on Geometry Processing*, pages 153–162, 2007.

[SZT+07] Xiaohan Shi, Kun Zhou, Yiying Tong, Mathieu Desbrun, Hujun Bao, and Baining Guo. Mesh puppetry: cascading optimization of mesh deformation with inverse kinematics. In *SIGGRAPH '07: ACM SIGGRAPH 2007 papers*, page 81, New York, NY, USA, 2007. ACM Press.

[TB03] Bill Tomlinson and Bruce Blumberg. Alphawolf: Social learning, emotion and development in autonomous virtual agents. *Lecture Notes in Computer Science*, 2564:35–45, 2003.

[TJ95] Frank Thomas and Ollie Johnston. *The illusion of life : Disney animation.* New York:Hyperion, 1995.

[Tor97] Nick Torkos. Footprint based quadruped animation. Master's thesis, University of Toronto, 1997.

[TP91] M. Turk and Alex Pentland. Eigen faces for recognition. *Journal of Cognitive Neuroscience*, 3(1):71–86, 1991.

[TT94] Xiaoyuan Tu and Demetri Terzopoulos. Artificial fishes: physics, locomotion, perception, behavior. In *SIGGRAPH '94: Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 43–50, New York, NY, USA, 1994. ACM Press.

[TT04] Nadia Magnenat Thalmann and Daniel Thalmann. An overview of virtual humans. In *Handbook of Virtual Humans*, pages 1–25. John Wiley, 2004.

[TvdP98] Nick Torkos and Michiel van de Panne. Footprint-based quadruped motion synthesis. In *Proceedings of Graphics Interface '98*, pages 151–160, 1998.

[VCG+01] H. Valladas, J. Clottes, J.-M. Geneste, M. A. Garcia, M. Arnold, H. Cachier, and N. Tisnrat-Laborde. Palaeolithic paintings: Evolution of prehistoric cave art. *Nature*, 2001.

[vdP96] Michiel van de Panne. Parameterized gait synthesis. In *Computer Graphics and Applications, IEEE*, volume 16, pages 40–49, 1996.

[vdP97] Michiel van de Panne. Making them move: Motor control for animated humans and animals. In *European Control Conference*, pages 191–210, 1997.

[vdP98] Michiel van de Panne. Control for simulated human and animal motion. In *Atelier IFAC International Workshop on Motion Control*, pages 425–435, 1998.

[VSH+98] Richard Vaughan, Neil Sumpter, Jane Henderson, Andy Frost, and Stephen Cameron. Robot control of animal flocks. 1998.

[WB05] Bill Westenhofer and Jim Berney. Chronicles of narnia: The lion, the witch and the wardrobe diaries: production diaries. *VFXWorlds*, December, 2005.

[WBK+07] Kelly Ward, Florence Bertails, Tae-Yong Kim, Stephen R. Marschner, Marie-Paule Cani, and Ming Lin. A survey on hair modeling: Styling, simulation, and rendering. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 13(2):213–34, Mar-Apr 2007.

[Wel09] Real time animation of virtual humans: A trade-off between naturalness and control. In *Eurographics State-of-the-Art Report (EG-STAR)*, page 4572. Eurographics Association, Eurographics Association, 2009.

[WFM01] Marcelo Walter, Alain Fournier, and Daniel Menevaux. Integrating shape and pattern in mammalian models. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 317–326, New York, NY, USA, 2001. ACM Press.

[WG97] Jane Wilhelms and Allen Van Gelder. Anatomically based modeling. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 173–180, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co.

[WG03] Jane Wilhelms and Allen Van Gelder. Combining vision and computer graphics for video motion capture. *The Visual Computer*, 19(6):360–376, October 2003.

[Wid] Richard Widgery. Kinetic impulse, horse locomotion: www.kinetic-impulse.com.

[Wik] Wikipedia. Films featuring anthromorphic characters.

[Wil86] Jane Wilhelms. Virya a motion control editor for kinematic and dynamic animation. In *Proceedings on Graphics Interface '86/Vision Interface '86*, pages 141–146, Toronto, Ont., Canada, Canada, 1986. Canadian Information Processing Society.

[Wil94] Josh Wilmes. The lion king. 1994.

[WML+06] Fu-Che Wu, Wan-Chun Ma, Rung-Huei Liang, Bing-Yu Chen, and Ming Ouhyoung. Domain connected graph: the skeleton of a closed 3d shape for animation. *The Visual Computer*, 22(2):117–135, 2006.

[WP00] Lawson Wade and Richard E. Parent. Fast, fully-automated generation of control skeletons for use in animation. In *CA '00: Proceedings of the Computer Animation*, page 164, Washington, DC, USA, 2000. IEEE Computer Society.

[WP02] Lawson Wade and Rick Parent. Automated generation of control skeletons for use in animation. *The Visual Computer*, 18(2):97–110, 2002.

[WP09] Kevin Wampler and Zoran Popovic. Optimal gait and form for animal locomotion. volume 28, pages 1–8, New York, NY, USA, 2009. ACM.

[ZS09] Liming Zhao and Alla Safonova. Achieving good connectivity in motion graphs. *Graph. Models*, 71(4):139–152, 2009.