Institute of Perception, Action and Behaviour

School of informatics

Overview

At IPAB, we study a variety of problems in perception, action and behaviour from a formal modelling perspective. We use methods such as Bayesian inference, hidden Markov models and statistical learning to model and investigate a broad variety of topics from insect sensing and behaviour to motor control in robotic and biological systems, and computer vision.

3D Computer Vision

3D computer vision involves reasoning about the nature of the 3D world based on scanned 3D data. We currently have 3 systems for acquiring this data: a very high precision laser scanner, a high resolution stereo system and a state-of-the-art video-rate stereo rig.

Plausible 3D Surface Completion

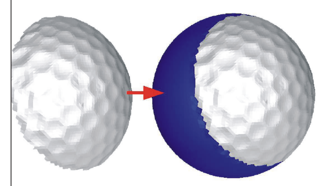

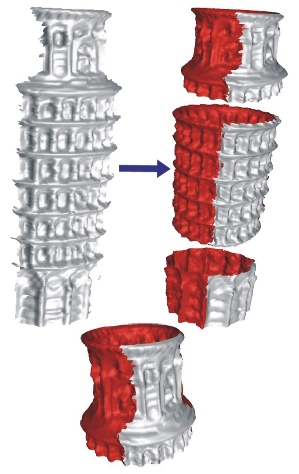

- Uni-directional 3D capture devices can only recover the visible portion of an object's surface
- But we wish to capture the entire object including backfacing and occluded surfaces

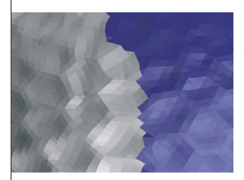

- We have been investigating methods for plausible 3D surface completion using geometric completion and surface relief propagation

Video-rate 3D Analysis

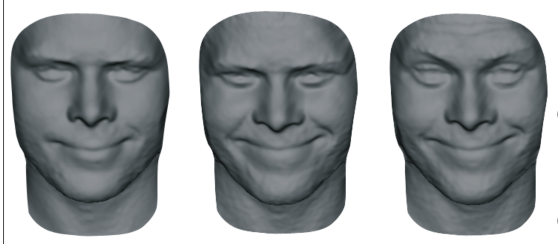

Using surface modelling and statistical techniques, we are attempting to answer questions such as:

- Can we compactly encode the dynamics of an individual's face?
- How unique is this, how does it vary with injury/surgery and can we use it to recover the underlying muscle structure?

Capturing Human Nuance

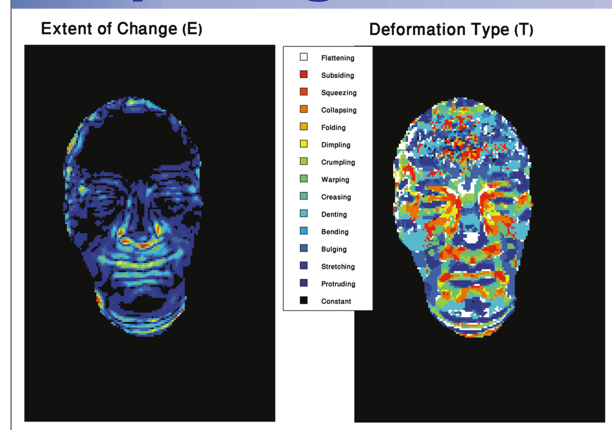

Temporal Characteristics of a smile

- Can human nuance be captured by analysing the local curvature of a face?
- We are currently extending classification schemes for static surface curvature into the temporal domain
- We seek to identify regions that exhibit variations in shape change (such as the face undergoing expressions), and to characterise the nature of the deformation

Animating Human Avatars

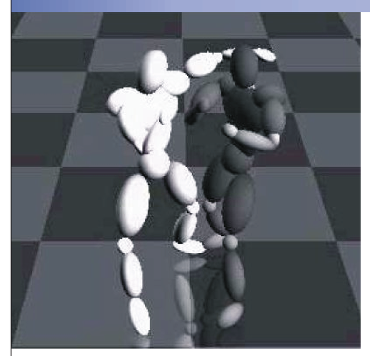

We have developed methods for simulating dense interactions of human avatars, such as pushing, fighting and wrestling. The aim is to generate realistic behaviour in novel contexts based on a small library of motion capture data.

Left: Simulating a fight scene by expanding the game tree
Below: Simulating an avatar being hit while walking

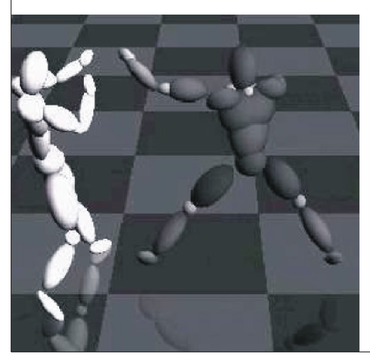

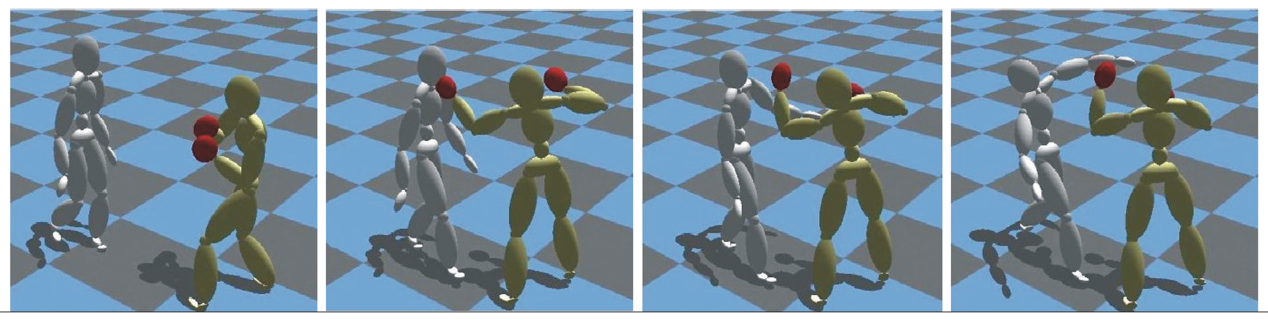

Statistical Learning for Motor Control

The SLMC group develops statistical machine learning methods, particularly in the context of control of high-dimensional movement systems such as humanoid robots.

Sensorimotor Learning in High Dimensional Sys.

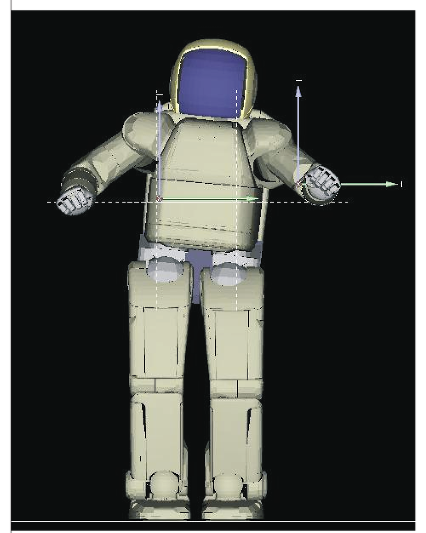

Current research projects include:

- Robust algorithms for real-time internal model learning in high dimensional movement system
- Learning cost functions for generating realistic human-like movement
- Biological motor learning: Identifying mechanisms and capabilities of learning and representation in the cerebellum.

Left: Learning cost functions for the Honda ASIMO robot (simulated kinematic model is shown here)
Right: DLR LW-III dextrous arm with hand

Learning under Varying Contexts

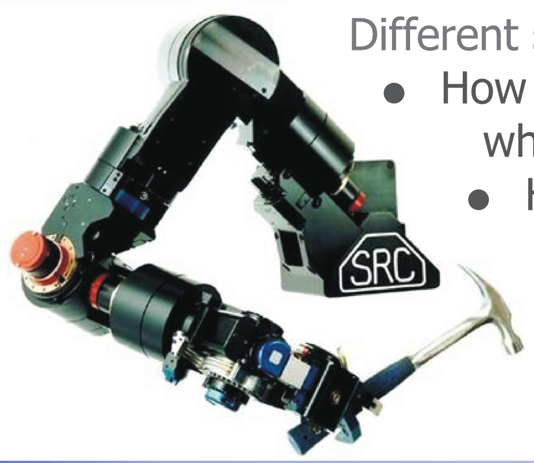

Different situations or contexts require different control strategies:

- How can we learn control in a new situation without unlearning what we already know?
  - How can force sensors on the fingertips help to determine what the context is?
    - For example: Is the bottle full or empty?

Left: 7 DOF SARCOS anthropomorphic arm holding a tool

Insect Sensing and Behaviour

The Biorobotics group is interested in studying the behaviour of invertebrates, motivated by the relative simplicity of their nervous systems.

Computational and Robotic Modelling

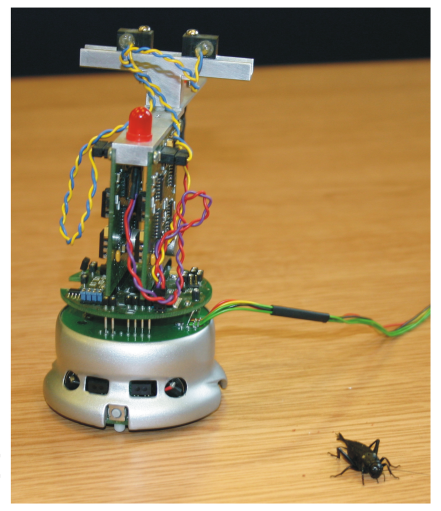

Models, either physical or simulated, are a valuable tool in investigating theories of behaviour. These models can vary in their level of description from single spiking neurons to complete systems, and are being applied to:

- Phonotaxis for mate localisation in crickets (right)
- Neural network models of insect learning in a classical conditioning paradigm

- Multimodal integration: vision and audition in crickets and vision and olfaction in fruit flies (left)

Left: automatically tracked trajectory of flying fly
Right: robot with two 'ears', for investigating cricket phonotaxis

Biologically inspired robotics

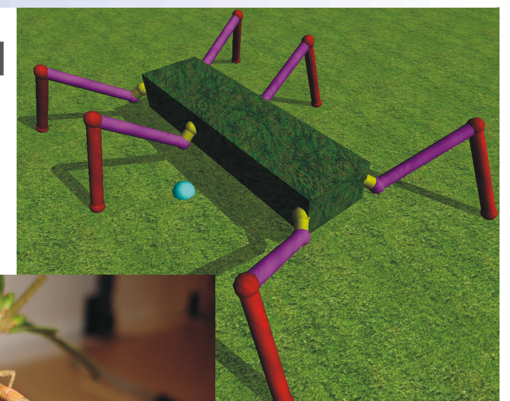

Insights gained from insects can be used to inspire novel algorithms or hardware for robots:

- Landmark-based homing, as seen in the desert ant
- Wind sensing in crickets
- Six-legged locomotion on uneven terrain, inspired by the stick insect (right)

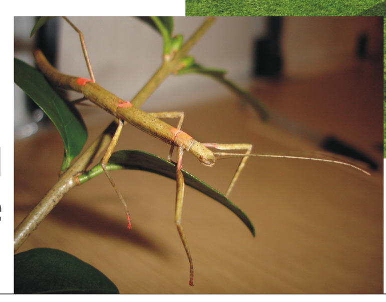

Right: Actual stick insect with coloured markers for visual tracking. Far right: simulated hexapod robot based on the stick insect

Visual Surveillance

Given the increased demands on security and surveillance system capabilities, and the vast quantity of associated video data, it is essential to automate the various processes involved; from integrating many views of a scene into a single 'superimage', to informing supervisors of interesting events, crowd analysis and behaviour classification.

Identifying Group Interactions

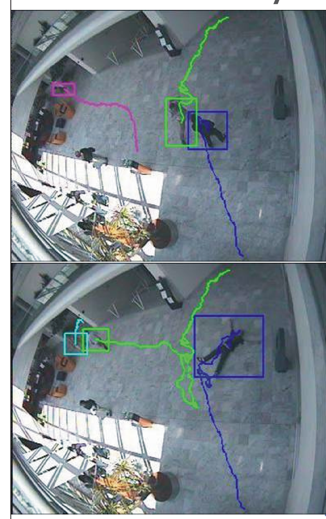

- Aim to establish who is interacting with who and the nature of interaction
- Identification of interaction class. Are these people: following, meeting, walking together, fighting etc.?

Left: Analysis of an assault captured on a surveillance camera

Detecting Abnormal Crowd Behaviour

- Flow based analysis of the behaviour of many individuals using crowd-flux statistics and congestion analysis
- We achieve this using explicit models of normal flow using optical flow with Hidden Markov Models used to filter normal crowds

Human Multimodal Perception

We perform psychophysics experiments to find out what people perceive when they receive conflicting information from multiple sources, and we compare the results with a formal Bayesian model of multimodal perception and cue integration.

- Do you trust what you see or what you hear?
- Integration vs Segregation: When should you combine information from different senses and when should you treat them as independent?

Right: Psychophysics experimental setup with eye tracker

Reinforcement learning

We are concerned with investigating how autonomous agents, either simulated or robotic, can learn via interaction with their environment.

Extending learning algorithms

Work is being undertaken to extend traditional reinforcement learning algorithms, applying them to a number of complex tasks and environments:

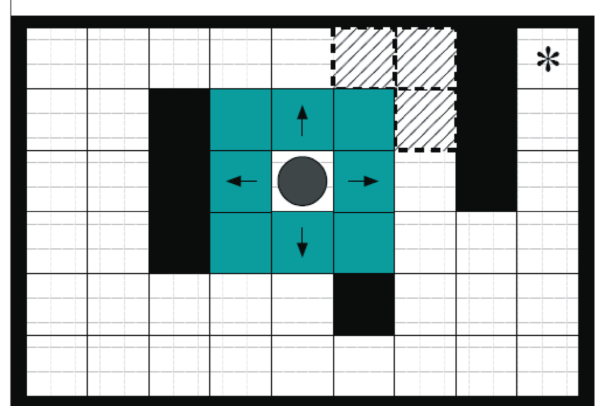

- Dynamic segmentation of the state/action space
- Active perception: including actions for sensing in ambiguous worlds (left)
- Internal motivations or 'drives' for learning

Left: perceptual aliasing in a gridworld environment

Multiagent systems

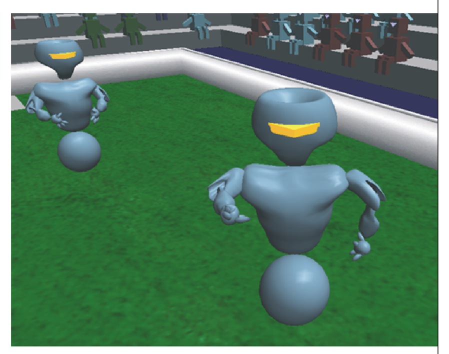

We are investigating ways in which agents can learn to co-operate with each other to produce cohesive team behaviour:

- Rewarding individuals for group behaviour (left)

- Learning to use actions based on inter-agent communication
- Learning by imitation
- Assembling robots from autonomous modules

Left: Multiagent reinforcement learning for commercial computer games
Right: teleo-reactive programming in an educational soccer game