



26th British Machine Vision Conference

Workshop Programme

PM Thursday 10 September



Swansea University
Prifysgol Abertawe

7th UK Computer Vision Student Workshop <i>Taliesin</i>	1
Machine Vision of Animals and their Behaviour <i>Faraday K & Taliesin Annexe</i>	4
Differential Geometry in Computer Vision for Analysis of Shapes, Images and Trajectories <i>Faraday Lecture Theatre & Taliesin Annexe</i>	9
Computer Vision Problems in Plant Phenotyping <i>Faraday L</i>	15
Internal Back Cover	Map

7th UK Computer Vision Student Workshop (BMVW)

Location: Taliesin

13:15 - 13:30 Setup

13:30 - 13:40 Opening Remarks

13:35 - 14:20 Keynote – Frontiers of Identification: Biometrics and Computer Vision

Mark Nixon, University of Southampton

14:20 – 15:00 Oral Presentations

Generating Local Temporal Poses from Gestures with Aligned Cluster Analysis for Human Action Recognition 1

Mike Edwards, Xianghua Xie (Swansea University)

Advancements in Contact-free Heart Rate Measurements Using Human Face Videos 1

Muhammad Waqar, Bernard Tiddeman (Aberystwyth University)

15:00 – 15:30 Coffee Break (Taliesin Annexe)

15:30 – 16:00 Invited Post-doc Talk – Computer Vision for Health: Classical Applications and New Perspectives... and My Life as a Post-doc

Dr. Adeline Paiement (Bristol University)

16:00 – 17:00 Oral Presentations

Analysis of Face and Segment Level Descriptors for Robust 3D Co-Segmentation 2

David George, Gary Tam, Xianghua Xie (Swansea University)

Using 3D Representations of the Nasal Region for Improved Landmarking and Expression Robust Recognition 2

Jiangning Gao, Adrian Evans (University of Bath)

Evaluating the Resilience Face Recognition Systems against Malicious Attacks 2

Luma Omar, Ioannis Ivrissimtzis (Durham University)

17:00 – 17:10 Closing Remarks

Machine Vision of Animals and their Behaviour (MVAB 2015)

Location: Faraday K

13:15 - 13:30 Setup

13:30 - 13:35 Welcome

13:35 - 14:20 Keynote – What can machine vision do for livestock?

Ilias Kyriazakis, Newcastle University

14:20 – 15:00 Livestock

Non-intrusive automated measurement of dairy cow body condition using 3D video 4

Mark Hansen (University of the West of England), Melvyn Smith (University of the West of England), Lyndon Smith (University of the West of England) and Duncan Forbes (Kingshay Conservation and Farming Ltd.)

Detecting Gallbladders in Chicken Livers using Spectral Imaging 4

Anders Jørgensen (Aalborg University), Eigil Mølviig Jensen (IHFood) and Thomas B. Moeslund (Aalborg University)

15:00 – 15:50 Coffee Break and Poster Session (Taliesin Annexe)

A Computer Vision Approach to Classification of Birds in Flight from Video Sequences 5

John Atanbori (University of Lincoln), John Murray (University of Lincoln), Wenting Duan (University of Lincoln), Kofi Appiah (Nottingham Trent University) and Patrick Dickinson (University of Lincoln)

Microfluidic environment and tracking analysis for the observation of Artemia Franciscana 5

Jose Alonso Solis-Lemus (City University London), Yushi Huang (RMIT University), Donald Wlodkowic (RMIT University) and Constantino Carlos Reyes-Aldasoro (City University London)

Invariant Image-Based Species Classification of Butterflies and Reef Fish 6

Hafeez Anwar (Vienna University of Technology), Sebastian Zambanini (Vienna University of Technology) and Martin Kampel (Vienna University of Technology)

15:50 – 16:35 Keynote – Fish detection, tracking, recognition, and analysis with the Fish4Knowledge dataset

Robert Fisher, University of Edinburgh

16:35 – 18:00 Fish and Macaque

Croatian Fish Dataset: Fine-grained classification of fish species in their natural habitat 6

Jonas Jäger (Hochschule Fulda), Marcel Simon (Friedrich Schiller University Jena), Joachim Denzler (Friedrich Schiller University Jena), Viviane Wolff (Hochschule Fulda), Klaus Fricke-Neudert (Hochschule Fulda) and Claudia Kruschel (University of Zadar)

Convolutional Neural Networks for Counting Fish in Fisheries Surveillance Video 7

Geoffrey French (University of East Anglia), Mark Fisher (University of East Anglia), Michal Mackiewicz (University of East Anglia) and Coby Needle (Marine Scotland)

Affinity Matting for Pixel-accurate Fin Shape Recovery from Great White Shark Imagery 7

Benjamin Hughes (University of Bristol) and Tilo Burghardt (University of Bristol)

Detecting and Tracking Bottoms and Faces of the Crested Black Macaque in the Wild 7

John Chiverton (University of Portsmouth), Jerome Micheletta (University of Portsmouth) and Bridget Waller (University of Portsmouth)

1st International Workshop on DIFFerential Geometry in Computer Vision for Analysis of Shapes, Images and Trajectories (DIFF CV)

Location: Faraday Lecture Theatre

13:15 - 13:30 Setup

13:30 - 13:35 Opening Remarks

13:35 - 14:35 Keynote – Some Thoughts on the Role of Geometry in Computer Vision

Anuj Srivastava, Florida State University

14:35 – 15:00 Poster Session (Taliesin Annexe)

Zero-Shot Domain Adaptation via Kernel Regression on the Grassmannian 9

Yongxin Yang, Timothy Hospedales (Queen Mary, University of London)

Karcher Mean in Elastic Shape Analysis 9

Wen Huang (Universite Catholique de Louvain), Yaqing You (Florida State University), Pierre-Antoine Absil (Universite Catholique de Louvain), Kyle Gallivan (Florida State University)

A Novel Riemannian Framework for Shape Analysis of Annotated Surfaces 10

Jiaqi Zaetz, Sebastian Kurtek (The Ohio State University)

A Generalized Lyapunov Feature for Dynamical Systems on Riemannian Manifolds 10

Rushil Anirudh, Vinay Venkataraman, Pavan Turaga (Arizona State University)

Improving 3D Facial Action Unit Detection with Intrinsic Normalization 11

Eric Yudin, Aaron Wetzler, Matan Sela, Ron Kimmel (Technion)

Gauge Invariant Framework for Trajectories Analysis 11

Hassen Drira, Barbara Tumpach, Mohamed Daoudi (University Lille1)

15:00 – 15:30 Coffee Break (Taliesin Annexe)

15:30 – 16:00 Poster Session Continued (Taliesin Annexe)

16:00 – 18:00 Oral Session

Distance Metric Learning by Optimization on the Stiefel Manifold 11

Ankita Shukla, Saket Anand (IIIT- Delhi)

Geometric Analysis of Axonal Tree Structures 12

Adam Duncan, Anuj Srivastava (Florida State University), Xavier Descombes (INRIA, Sophia Antipolis), Eric Klassen (Florida State University)

Second Order Elastic Metrics on the Shape Space of Curves 12

(University of Vienna), Martins Bruveris, Philipp Harms (ETH Zurich), Jakob Møller-Andersen (Technical University of Denmark)

FFT-based Alignment of 2D Closed Curves with Application to Elastic Shape Analysis 13

Gunay Dogan (Theiss Research), Javier Bernal, Charles Hagwood (National Institute of Standards and Technology)

Temporal Reflection Symmetry of Human Actions: A Riemannian Analysis 13

Qiao Wang, Rushil Anirudh, Pavan Turaga (Arizona State University)

Cylindrical Surface Reconstruction by Fitting Paths on Shape Space 14

Chafik Samir (University of Clermont), Pierre-Yves Gousenbourger (Universite Catholique de Louvain), Shantanu Joshi (University of California, Los Angeles)

18:00 Adjourn

Computer Vision Problems in Plant Phenotyping (CVPPP 2015)

Location: Faraday L

13:15 - 13:30 Setup

13:30 - 13:40 Opening Remarks

13:40 - 14:20 Keynote - Image Segmentation using Statistical Shape and Appearance Models

Tim Cootes, University of Manchester

14:20 – 15:00 Counting Leaves – Results from the LCC Challenge

Learning to Count Leaves in Rosette Plants 15

Mario Valerio Giuffrida, Massimo Minervini, Sotirios Tsafaris (IMT Lucca)

Counting leaves without "finger-counting" by supervised multiscale frequency analysis of depth images from top view 15

David Rousseau (Université Lyon 1), Henricus J. Van de Zedde (Wageningen)

15:00 – 15:30 Coffee Break (Taliesin Annexe)

15:30 – 17:10 Segmentation, Annotation and Reconstruction - Results from the LSC Challenge

Utilizing machine learning approaches to improve prediction of leaf counts and individual leaf segmentation of rosette plants..... 16

Jean-Michel Pape (IPK-Gatersleben), Christian Klukas (LemnaTec GmbH)

Gaussian process shape models for Bayesian leaf segmentation 16

Kyle Simek, Kobus Barnard (University of Arizona)

Fuzzy c-means based plant segmentation with distance dependent threshold 17

Mads Dyrmann (University of Southern Denmark)

An interactive tool for semi-automated leaf annotation 17

Massimo Minervini, Mario Valerio Giuffrida, Sotirios Tsafaris (IMT Lucca)

3D Surface Reconstruction of Plant Seeds by Volume Carving 18

Johanna Roussel, Andreas Fischbach, Siegfried Jahnke, Hanno Scharf (Forschungszentrum Jülich)

17:10 - 17:20 Closing Remarks

17:20 Adjourn

7th UK Computer Vision Student Workshop (BMVW)

Generating Local Temporal Poses from Gestures with Aligned Cluster Analysis for Human Action Recognition

Mike Edwards, Xianghua Xie (Swansea University)

The use of pose estimation for human action recognition has seen a resurgence in previous years, due in part to the natural representation of the activity as a sequence of key poses and gestures. The use of sequence alignment techniques has aided the process of comparing between sequences of differing temporal rates, with aligned cluster analysis segmenting an observation into lower level action primitives. We suggest that the representation of a given action class via its lower level gestures can help to identify the higher-level action class label. We therefore present a method for the generation of key poses via the initial segmentation of an action class into gestures that are similar across numerous observations. We treat all training observations as a single observation in which there are repetitions of the same action class. By applying segmentation, we then identify common gestures across the class, which are used to generate the key poses we optimize via evolutionary programming. Global recognition rates of 97.4\% are achieved using a subset of the MSR Action3D dataset. We then expand the method to recognize interaction events between two individuals using the SBU Kinect Interaction dataset, achieving recognition rates of 83.9\% and over 96.4\% when observing the first 6 classes.

Advancements in Contact-free Heart Rate Measurements Using Human Face Videos

Muhammad Waqar, Bernard Tiddeman (Aberystwyth University)

Recently, several methods have been reported for the contact-free measurement of heart rate from human face videos. Typically these are based on applying independent component analysis to sequences of colour values from tracked face patches and selecting the peak value within a fixed range in the frequency spectrum. Those methods performed well enough under controlled conditions but in realistic situations their performance degraded significantly. In this paper, we have highlighted issues such as light reflection and participant's movement, and proposed solutions to those factors that reduces the performance of existing methods. We experiment with calculating a robust mean of the skin colour pixels values using a mean shift algorithm and estimating an optimal frequency selection function in the Fourier domain. Preliminary results show that these methods can improve the results of existing methods.

Analysis of Face and Segment Level Descriptors for Robust 3D Co-Segmentation

David George, Gary Tam, Xianghua Xie (Swansea University)

3D shape co-segmentation is an important topic in computer graphics. The idea of co-analysis brings new insights into understanding a collection of shapes. Rather than analysing individual shapes, an entire set is looked at, giving much more information about the class of shape in the set. Existing co-segmentation techniques use both face and segment level descriptors in order to provide enough information to give an accurate co-segmentation result. In the literature, a lot of these descriptors are proposed but there is limited empirical studies to compare which would perform well. In this paper, we have two aims: (a) propose new useful face and segment level descriptors and (b) analyse the effectiveness of them. Our experiment indicates that smoothly varying descriptors (Average Euclidean Distance) that respects geometry would improve the segmentation results.

Using 3D Representations of the Nasal Region for Improved Landmarking and Expression Robust Recognition

Jiangning Gao, Adrian Evans (University of Bath)

This paper investigates the performance of different representations of 3D human nasal region for expression robust recognition. By performing evaluations on the depth and surface normal components of the facial surface, the nasal region is shown to be relatively consistent over various expressions, providing motivation for using the nasal region as a biometric. A new efficient landmarking algorithm that thresholds the local surface normal components is proposed and demonstrated to produce an improved recognition performance for nasal curves from both the depth and surface normal components. The use of the Shape Index for feature extraction is also investigated and shown to produce a good recognition performance.

Evaluating the Resilience Face Recognition Systems against Malicious Attacks

Luma Omar, Ioannis Ivrissimtzis (Durham University)

This paper presents an experiment designed to test the resilience of several user verification systems based on face recognition technology against simple identity spoofing methods, such as trying to gain access to the system by using mobile camera shots of the users, their ID cards, or social media photos of them that are available online. We also aim at identifying the compression threshold above which a photo can be used to gain access to the system. Four major user verification tools

BMVW

were tested: Keyemon and Luxand Blink on Windows and Android Face Unlock and FaceLock on Android. The results show all tested systems to be vulnerable to even very crude attacks, indicating that the technology is not ready yet for adoption in applications where security rather than user convenience is the main concern.

Machine Vision of Animals and their Behaviour (MVAB 2015)

Non-intrusive automated measurement of dairy cow body condition using 3D video

Mark Hansen (University of the West of England), Melvyn Smith (University of the West of England), Lyndon Smith (University of the West of England) and Duncan Forbes (Kingshay Conservation and Farming Ltd.)

Regular scoring of a dairy herd in terms of various physical metrics such as Body Condition Score (BCS), mobility and weight are essential for maintaining high animal welfare. This paper presents preliminary results of an automated system capable of nonintrusively measuring BCS automatically as the cow walks uninhibited beneath a 3D camera. The system uses a 'rolling ball' algorithm on the depth map which simulates how well a ball of a set radius fits the surface. In this way a measure of angularity is generated which is shown to be inversely related to BCS on 95 cows. The measurements are shown to be highly repeatable with 14 out of 15 cows being scored within one quarter BCS score repeatedly and seven of those being scored within an eighth of a BCS score.

Detecting Gallbladders in Chicken Livers using Spectral Imaging

Anders Jørgensen (Aalborg University), Eigil Mølviig Jensen (IHFood) and Thomas B. Moeslund (Aalborg University)

This paper presents a method for detecting gallbladders attached to chicken livers using spectral imaging. Gallbladders can contaminate good livers, making them unfit for human consumption. A data set consisting of chicken livers with and without gallbladders, has been captured using 33 wavelengths within the visible spectrum. This work shows how to reduce the high number of wavelengths while maintaining a high accuracy. A classification tree has been trained to evaluate if a gallbladder is present and whether it is suitable for automatic removal, which could increase profits for the processing plants. As a preliminary study this shows good results with a classification accuracy of 91.7%.

A Computer Vision Approach to Classification of Birds in Flight from Video Sequences

John Atanbori (University of Lincoln), John Murray (University of Lincoln), Wenting Duan (University of Lincoln), Kofi Appiah (Nottingham Trent University) and Patrick Dickinson (University of Lincoln)

Bird populations are an important bio-indicator; so collecting reliable data is useful for ecologists helping conserve and manage fragile ecosystems. However, existing manual monitoring methods are labour-intensive, time-consuming, and error-prone. The aim of our work is to develop a reliable system, capable of automatically classifying individual bird species in flight from videos. This is challenging, but appropriate for use in the field, since there is often a requirement to identify in flight, rather than when stationary. We present our work in progress, which uses combined appearance and motion features to classify and present experimental results across seven species using Normal Bayes classifier with majority voting and achieving a classification rate of 86%.

Microfluidic environment and tracking analysis for the observation of Artemia Franciscana

Jose Alonso Solis-Lemus (City University London), Yushi Huang (RMIT University), Donald Wlodkowic (RMIT University) and Constantino Carlos Reyes-Aldasoro (City University London)

The movement of small animals in well-defined environments is increasingly used in ecotoxicology as a sensitive indicator of water contamination with toxicants. This work describes a framework for the analysis of *Artemia franciscana* within miniaturised chambers. This work combines: (a) a proof-of-concept miniaturised Lab-on-a-Chip (LOC) platform where any changes in swimming patterns exhibited by *Artemia franciscana* can be related to the toxicity of the environment with (b) a pipeline of image processing steps for the movement analysis of video sequences of test specimens. The analysis pipeline is semiautomatic as it allows a user to define a region of interest and confirm intensity levels that are used for two-level hysteresis segmentation. Tracking is automatic, and a large number of statistical measurements and visualisation tools are provided. In addition, proofreading tools for validation of tracks are provided to link, break or delete tracks. The software consists of Matlab® m-files and is available in the website <http://www.objecttrack.org.uk/>.

Invariant Image-Based Species Classification of Butterflies and Reef Fish

Hafeez Anwar (Vienna University of Technology), Sebastian Zambanini (Vienna University of Technology) and Martin Kampel (Vienna University of Technology)

We propose a framework for species-based image classification of butterflies and reef fish. To support such image-based classification, we use an image representation which enriches the famous bag-of-visual words (BoVWs) model with spatial information. This image representation is developed by encoding the global geometric relationships of visual words in the 2D image plane in a scale- and rotation-invariant manner. In this way, invariance is achieved to the most common variations found in the images of these animals as they can be imaged at different image locations, exhibit various in-plane orientations and have various scales in the images. The images in our butterfly and reef fish datasets belong to 30 species of each animal. We achieve better classification rates on both the datasets than the ordinary BoVWs model while still being invariant to the mentioned image variations. Our proposed image-based classification framework for butterfly and reef fish species can be considered as a helpful tool for scientific research, conversation and education.

Croatian Fish Dataset: Fine-grained classification of fish species in their natural habitat

Jonas Jäger (Hochschule Fulda), Marcel Simon (Friedrich Schiller University Jena), Joachim Denzler (Friedrich Schiller University Jena), Viviane Wolff (Hochschule Fulda), Klaus Fricke-Neuderth (Hochschule Fulda) and Claudia Kruschel (University of Zadar)

This paper presents a new dataset for fine-grained visual classification (FGVC) of fish species in their natural environment. It contains 794 images of 12 different fish species collected at the Adriatic sea in Croatia. All images show fishes in real live situations, recorded by high definition cameras. Remote and diver-based videography is used by a growing number of marine researchers to understand spatial and temporal variability of habitats and species. The required large numbers of independent observations necessitate the development of computer vision tools for an automated processing of high volumes of videos featuring high fish richness and density. As baseline experiment, we are using CNN features [1] and a linear SVM classifier and achieve an accuracy of 66:78% on our dataset.

Convolutional Neural Networks for Counting Fish in Fisheries Surveillance Video

Geoffrey French (University of East Anglia), Mark Fisher (University of East Anglia), Michal Mackiewicz (University of East Anglia) and Coby Needle (Marine Scotland)

We present a computer vision tool that analyses video from a CCTV system installed on fishing trawlers to monitor discarded fish catch. The system aims to support expert observers who review the footage and verify numbers, species and sizes of discarded fish. The operational environment presents a significant challenge for these tasks. Fish are processed below deck under fluorescent lights, they are randomly oriented and there are multiple occlusions. The scene is unstructured and complicated by the presence of fishermen processing the catch. We describe an approach to segmenting the scene and counting fish that exploits the N4-Fields algorithm. We performed extensive tests of the algorithm on a data set comprising 443 frames from 6 belts. Results indicate the relative count error (for individual fish) ranges from 2% to 16%. We believe this is the first system that is able to handle footage from operational trawlers.

Affinity Matting for Pixel-accurate Fin Shape Recovery from Great White Shark Imagery

Benjamin Hughes (University of Bristol) and Tilo Burghardt (University of Bristol)

The objective of this paper is to obtain pixel-accurate reconstructions of white shark fins given automatically generated coarse pre-segmentations. Reconstruction performance is compared for affinity matting, colour matting and GrabCut against expert annotated ground truth for a test-set of 120 fin images taken in the wild. For the present domain, we find affinity matting able to most accurately recover fine shape details, whilst being robust to wide baseline trimap initialisations as needed to reconstruct prominent notches on the fin edge.

Detecting and Tracking Bottoms and Faces of the Crested Black Macaque in the Wild

John Chiverton (University of Portsmouth), Jerome Micheletta (University of Portsmouth) and Bridget Waller (University of Portsmouth)

Monkeys are important to many areas of science and ecology. The study of monkeys and their welfare are important components requiring complex observational studies. This work is therefore concerned with the development of computer vision

MVAB

techniques for the purposes of detecting and tracking monkeys with the ultimate aim to help in such studies. Monkeys are complex creatures for the purposes of tracking because of complex deformations. This complexity is further compounded by an in the wild setting where forest conditions result in frequent occlusions and changes in lighting. Despite these complexities monkeys present some interesting features that can make detection and tracking possible: their bottoms and faces. A system is thus described consisting of detectors trained to detect faces and bottoms of monkeys which are used within a tracking framework to initialise a system of tracklet construction. Steps are also described to enable disparate but coincident tracklets to be merged thus enabling longer run analysis of individual monkey movements. Experiments are performed using image data taken from video footage of Crested Black Macaques in natural forest surroundings. Results demonstrate relatively successful detection of monkey bottoms where the correspondence analysis and tracking process helps to reduce false positives.

1st International Workshop on DIFFerential Geometry in Computer Vision for Analysis of Shapes, Images and Trajectories (DIFF CV)

Zero-Shot Domain Adaptation via Kernel Regression on the Grassmannian

Yongxin Yang, Timothy Hospedales (Queen Mary, University of London)

Most visual recognition methods implicitly assume the data distribution remains unchanged from training to testing. However, in practice domain shift often exists, where real-world factors such as lighting and sensor type change between train and test, and classifiers do not generalise from source to target domains. It is impractical to train separate models for all possible situations because collecting and labelling the data is expensive. Domain adaptation algorithms aim to ameliorate domain shift, allowing a model trained on a source to perform well on a different target domain. However, even for the setting of unsupervised domain adaptation, where the target domain is unlabelled, collecting data for every possible target domain is still costly. In this paper, we propose a new domain adaptation method that has no need to access either data or labels of the target domain when it can be described by a parametrised vector and there exists several related source domains within the same parametric space. It greatly reduces the burden of data collection and annotation, and our experiments show some promising results.

Karcher Mean in Elastic Shape Analysis

Wen Huang (Universite Catholique de Louvain), Yaqing You (Florida State University), Pierre-Antoine Absil (Universite Catholique de Louvain), Kyle Gallivan (Florida State University)

In the framework of elastic shape analysis, a shape is invariant to scaling, translation, rotation and reparameterization. Since this framework does not yield a closed form of geodesic between two shapes, iterative methods have been proposed. In particular, path straightening methods have been proposed and used for computing a geodesic that is invariant to curve scaling and translation. Path straightening can then be exploited within a coordinate-descent algorithm that computes the best rotation and reparameterization of the end point curves [13]. A Riemannian quasi-Newton method to compute a geodesic invariant to scaling, translation, rotation and reparameterization has been given in [15] and shown to be more efficient than the coordinate-descent/path-straightening approach. This paper extends [15] by

showing that using the new approach to the geodesic when computing the Karcher mean yields a faster algorithm.

A Novel Riemannian Framework for Shape Analysis of Annotated Surfaces

Jiaqi Zaetz, Sebastian Kurtek (The Ohio State University)

We present a novel, parameterization-invariant method for shape analysis of annotated surfaces. While the method can handle various types of annotation including color and texture, in this paper we focus on soft landmark annotations. Landmark annotations are commonly provided in various applications including medical imaging where an expert marks points of interest on the objects. Most methods in current literature either study shapes using landmarks only or surfaces only. In either case, the analyst is forced to ignore a lot of useful information. We propose a novel representation of surfaces that can jointly incorporate shape and landmark annotation. Our framework properly removes all shape preserving transformations from the representation space including translation, scale, rotation and re-parameterization. We present results of comparing, averaging and classifying annotated shapes on toy and real data.

A Generalized Lyapunov Feature for Dynamical Systems on Riemannian Manifolds

Rushil Anirudh, Vinay Venkataraman, Pavan Turaga (Arizona State University)

Dynamic phenomena such as human activities, dynamic scenes, and moving crowds are commonly observed through visual sensors, resulting in feature trajectories sampled in time. Such phenomena can be accurately modeled by taking the temporal variations and changes into account. For problems where the trajectories are sufficiently different, elastic metrics can provide distances that are invariant to speed, but for more complex problems such as fine grained activity classification, one needs to exploit higher order dynamical properties. For features in the Euclidean space, applications such as crowd monitoring, dynamic scene recognition and human movement quality analysis have found a lot of success this way. In this paper we propose the largest Riemannian Lyapunov exponent (L-RLE), which is the first generalization of the largest Lyapunov exponent to Riemannian manifolds. The largest Lyapunov exponent is a classic measure to quantify the amount of chaos within signals in the Euclidean space, and allows us to exploit higher order dynamics for various applications. We show the effectiveness of the L-RLE on two manifolds - the Grassmann and the $SO(3)$ lie group. By modeling human actions as dynamic processes evolving on Riemannian manifolds, we show that L-RLE can measure the amount of chaos within each action accurately. We show that our measure is a good

generalization of largest Euclidean Lyapunov exponent (L-ELE), and is less susceptible to arbitrary distortions.

Improving 3D Facial Action Unit Detection with Intrinsic Normalization

Eric Yudin, Aaron Wetzler, Matan Sela, Ron Kimmel (Technion)

Data normalization techniques are commonly used to reduce intra-class variance while preserving inter-class differences that aid in classification and regression tasks. Such practices bring all data examples to similar scale and range, and help to decrease the dimensionality of categorization problems. In this paper we describe a novel use of a geometric framework to extend the concept of data normalization to the domain of functions that lie on surfaces. In this context, normalization is posed as an embedding of all example functions into Riemannian manifolds nearly isometric to one another. Using geometric tools, we propose an implementation for the case of discretized functions on triangulated meshes. To demonstrate the proposed framework we apply it as a preprocessing step to the task of automatic facial action unit detection for depth maps from a 3D scanner. We show how this subject-independent representation produces more accurate action unit classification results. We also empirically demonstrate that utilizing this intrinsic normalization technique indeed improves action unit recognition performance.

Gauge Invariant Framework for Trajectories Analysis

Hassen Drira, Barbara Tumpach, Mohamed Daoudi (University Lille1)

Distance Metric Learning by Optimization on the Stiefel Manifold

Ankita Shukla, Saket Anand (IIT- Delhi)

Distance metric learning has proven to be very successful in various problem domains. Most techniques learn a global metric in the form of a $n \times n$ symmetric positive semidefinite (PSD) Mahalanobis distance matrix, which has $O(n^2)$ unknowns. The PSD constraint makes solving the metric learning problem even harder making it computationally intractable for high dimensions. In this work, we propose a flexible formulation that can employ different regularization functions, while implicitly maintaining the positive semidefiniteness constraint. We achieve this by eigendecomposition of the rank p Mahalanobis distance matrix followed by a joint optimization on the Stiefel manifold $S_{n;p}$ and the positive orthant R_{p+} . The resulting nonconvex optimization problem is solved by employing an alternating strategy. We use a recently proposed projection free approach for efficient optimization over the Stiefel manifold. Even though the problem is nonconvex, we empirically show competitive classification accuracy on UCI and USPS digits datasets.

Geometric Analysis of Axonal Tree Structures

Adam Duncan, Anuj Srivastava (Florida State University), Xavier Descombes (INRIA, Sophia Antipolis), Eric Klassen (Florida State University)

This paper develops a framework for shape analysis of tree-like structures with the following common features: (1) a main branch viewed as a parameterized curve in \mathbb{R}^3 , and (2) a random number of secondary branches, each one of them a parameterized curve in \mathbb{R}^3 , emanating from the main branch at arbitrary points. In this framework, comparisons of objects is based on shapes-scales-orientations of the curves involved, and locations and number of the side branches. The objects are represented as composite curves made up of: a main branch and a continuum of side branches along the main branch with each branch being a curve in \mathbb{R}^3 itself (including the null curve, or zero curve). Extending the previous work on elastic shape analysis of Euclidean curves, the space of these composite curves is endowed with a natural Riemannian metric, using the SRVF representation, and one computes geodesic paths in the quotient space of this representation modulo the re-parameterization function. As a result, appropriate geometric structures are optimally matched across trees, and geodesic paths show deformations of main branches into each other while either deforming/sliding/creating/destroying the side branches. We present some preliminary results using axonal trees taken from the Neuromorpho database.

Second Order Elastic Metrics on the Shape Space of Curves

Martin Bauer (University of Vienna), Martins Bruveris, Philipp Harms (ETH Zurich), Jakob Møller-Andersen (Technical University of Denmark)

Second order Sobolev metrics on the space of regular unparametrized planar curves have several desirable completeness properties not present in lower order metrics, but numeric are still largely missing. In this paper, we present algorithms to numerically solve the initial and boundary value problems for geodesics. The combination of these algorithms allows to compute Karcher means in a Riemannian gradient-based optimization scheme. Our framework has the advantage that the constants determining the weights of the zero, first, and second order terms of the metric can be chosen freely. Moreover, due to its generality, it could be applied to more general spaces of mapping. We demonstrate the effectiveness of our approach by analyzing a collection of shapes representing physical objects.

FFT-based Alignment of 2D Closed Curves with Application to Elastic Shape Analysis

Gunay Dogan (Theiss Research), Javier Bernal, Charles Hagwood (National Institute of Standards and Technology)

For many shape analysis problems in computer vision and scientific imaging (e.g., computational anatomy, morphological cytometry), the ability to align two closed curves in the plane is crucial. If the curves have the same length and are centered at the origin, the critical steps to an optimal alignment are finding the best rotation for one curve to match the other and redefining the starting point of the rotated curve so that the starting points of the two curves match. Unlike open curves, closed curves do not have fixed starting points, and this introduces an additional degree of freedom in the alignment. Hence the common naive method to find the best rotation and starting point for optimal alignment has $O(N^2)$ time complexity, N the number of nodes per curve. This can be slow for curves with large numbers of nodes. In this paper, we propose a new $O(N \log N)$ algorithm for this problem based on the Fast Fourier Transform. Together with uniform resampling of the curves with respect to arc length, the new algorithm results in an order of magnitude speed-up in our experiments. Additionally, we describe how we can use our new algorithm as part of elastic shape distance computations between closed curves to obtain accurate shape distance values at a fraction of the cost of previous approaches.

Temporal Reflection Symmetry of Human Actions: A Riemannian Analysis

Qiao Wang, Rushil Anirudh, Pavan Turaga (Arizona State University)

While spatial symmetry of objects is studied frequently, the problem of measuring temporal symmetry has rarely been explored. In this paper we present a mathematical framework that utilizes a recently developed differential geometric approach to quantify temporal reflection symmetry in movements from video data, in which human movements are modeled as trajectories on Riemannian manifolds. Using the recently proposed transport square-root velocity functions (TSRVF) representation, the amounts of temporal asymmetry in movements are decomposed into a spatial component and a temporal one, each of which is quantified by an asymmetry score.

Cylindrical Surface Reconstruction by Fitting Paths on Shape Space

Chafik Samir (University of Clermont), Pierre-Yves Gousenbourger (Universite Catholique de Louvain), Shantanu Joshi (University of California, Los Angeles)

We present a differential geometric approach for cylindrical anatomical surface reconstruction from 3D volumetric data that may have missing slices or discontinuities. We extract planar boundaries from the 2D image slices, and parameterize them by an indexed set of curves. Under the SRVF framework, the curves are represented as invariant elements of a nonlinear shape space. Differently from standard approaches, we use tools such as exponential maps and geodesics from Riemannian geometry and solve the problem of surface reconstruction by fitting paths through the given curves. Experimental results show the surface reconstruction of smooth endometrial tissue shapes generated from MRI slices.

Computer Vision Problems in Plant Phenotyping (CVPPP 2015)

Learning to Count Leaves in Rosette Plants

Mario Valerio Giuffrida, Massimo Minervini, Sotirios Tsafaris (IMT Lucca)

Counting the number of leaves in plants is important for plant phenotyping, since it can be used to assess plant growth stages. We propose a learning-based approach for counting leaves in rosette (model) plants. We relate image-based descriptors learned in an unsupervised fashion to leaf counts using a supervised regression model. To take advantage of the circular and coplanar arrangement of leaves and also to introduce scale and rotation invariance, we learn features in a log-polar representation. Image patches extracted in this log-polar domain are provided to K-means, which builds a codebook in an unsupervised manner. Feature codes are obtained by projecting patches on the codebook using the triangle encoding, introducing both sparsity and specifically designed representation. A global, per-plant image descriptor is obtained by pooling local features in specific regions of the image. Finally, we provide the global descriptors to a support vector regression framework to estimate the number of leaves in a plant. We evaluate our method on datasets of the Leaf Counting Challenge (LCC), containing images of Arabidopsis and tobacco plants. Experimental results show that on average we reduce absolute counting error by 40% w.r.t. the winner of the 2014 edition of the challenge –a counting via segmentation method. When compared to state-of-the-art density-based approaches to counting, on Arabidopsis image data 75% less counting errors are observed. Our findings suggest that it is possible to treat leaf counting as a regression problem, requiring as input only the total leaf count per training image.

Counting leaves without ‘finger-counting’ by supervised multiscale frequency analysis of depth images from top view

David Rousseau (Université Lyon 1), Henricus J. Van de Zedde (Wageningen)

Depth imaging is applied to characterize the shoot of seedlings from top-view. We demonstrate how quantitative informations of biological interest, such as leaves counting can be extracted from such images without performing 3D reconstruction of the shoot. This is obtained from 2D Fourier multiscale analysis without any requirement to segment nor detect leaves one by one numerically. We discuss the robustness and limitations of this approach and present possible extension with 3D Fourier analysis applied to estimate the plant plastochrone or 3D+T Fourier analysis in the estimation of circadian rhythms.

Utilizing machine learning approaches to improve prediction of leaf counts and individual leaf segmentation of rosette plants

Jean-Michel Pape (IPK-Gatersleben), Christian Klukas (LemnaTec GmbH)

The segmentation of individual leaves in plant images is still a challenging task, especially in case of leaf overlaps. The exact determination of individual leaf areas could improve the biomass estimation which is a good indicator for plant performance. In addition, the number of leaves is directly related to plant development, leaf counts give insight into changing plant development stages. Machine learning is a powerful tool in vision tasks. Here we propose an approach including image analysis (based on the software IAP) for extraction of a comprehensive set of image features to predict the number of leaves for *Arabidopsis thaliana* and tobacco plants supplied by the organizers of the Leaf Counting Challenge (LCC) of the Computer Vision Problems in Plant Phenotyping (CVPPP) workshop in conjunction with the British Machine Vision Conference (BMVC) in 2015. In addition, we developed a method to detect exact leaf borders for resolving inaccurate leaf segmentation in case of leaf overlaps. For classifier training we evaluate a broad set of different colour and texture features. The predicted leaf borders are used as input for further image processing methods to complete the leaf segmentation. The results show the methods ability for improved leaf count estimations and for predicting leaf overlap borders, which helps to improve the segmentation of individual leaves.

Gaussian process shape models for Bayesian leaf segmentation

Kyle Simek, Kobus Barnard (University of Arizona)

We develop a novel probabilistic model for multi-part shapes based on Gaussian processes, which we apply to model rosette leaves of *Arabidopsis* plants. Our model incorporates domain knowledge of *Arabidopsis* leaves in two ways. First, leaves are modeled using two anatomical parts: a blade and a petiole. We model the two regions with separate Gaussian processes, with a smoothness constraint at the boundary. Second, we constrain all leaf petioles to initiate at the rosette center, which is also modeled. This Bayesian prior is combined with a simple likelihood function over foreground pixels to perform image segmentation by optimizing a posterior distribution. A simple data driven approach is used to over-segment the image, then excess leaves are pruned using a Bayesian model selection criterion. We show that our approach is effective, even with minimal training data.

Fuzzy c-means based plant segmentation with distance dependent threshold

Mads Dyrmann (University of Southern Denmark)

An important element in weed control using machine vision is the ability to identify plant species based on shape. For this to be done, it is often necessary to segment the plants from the soil. This may cause problems, if the colour of a plant is not consistent, since plants are then at risk of being separated into several objects. This study presents a plant segmentation method based on fuzzy c-means and a distance transform. This segmentation method is compared with four other plant segmentation methods based on various parameters, including the ability to maintain the plants as whole, connected components. The method presented here is found to be better at preserving plants as connected objects, while keeping the false positive rate low compared to commonly used segmentations techniques.

An interactive tool for semi-automated leaf annotation

Massimo Minervini, Mario Valerio Giuffrida, Sotirios Tsaftaris (IMT Lucca)

High throughput plant phenotyping is emerging as a necessary step towards meeting agricultural demands of the future. Central to its success is the development of robust computer vision algorithms that analyze images and extract phenotyping information to be associated with genotypes and environmental conditions for identifying traits suitable for further development. Obtaining leaf level quantitative data is important towards understanding better this interaction. While certain efforts have been made to obtain such information in an automated fashion, further innovations are necessary. In this paper we present an annotation tool that can be used to semi-automatically segment leaves in images of rosette plants. This tool, which is designed to exist in a stand-alone fashion but also in cloud based environments, can be used to annotate data directly for the study of plant and leaf growth or to provide annotated datasets for learning-based approaches to extracting phenotypes from images. It relies on an interactive graph-based segmentation algorithm to propagate expert provided priors (in the form of pixels) to the rest of the image, using the random walk formulation to find a good per-leaf segmentation. To evaluate the tool we use standardized datasets available from the LSC and LCC 2015 challenges, achieving an average leaf segmentation accuracy of almost 97% using scribbles as annotations. The tool and source code are publicly available at <http://www.phenotiki.com> and as a GitHub repository at <https://github.com/phenotiki/LeafAnnotationTool>.

3D Surface Reconstruction of Plant Seeds by Volume Carving

Johanna Roussel, Andreas Fischbach, Siegfried Jahnke, Hanno Scharf
(Forschungszentrum Jülich)

We describe a method for 3D reconstruction of plant seed surfaces, focusing on small seeds with diameters as small as 200 μ m. The method considers robotized systems allowing single seed handling in order to rotate a single seed in front of a camera. Even though such systems feature high position repeatability, at sub-millimeter object scales, camera pose variations have to be compensated. We do this by robustly estimating the tool center point from each acquired image. 3D reconstruction can then be performed by a simple shape-from-silhouette approach. In experiments we investigate runtimes, the achieved accuracy, and show as a proof of principle that the proposed method is well sufficient for 3D seed phenotyping purposes.

Swansea University Campus



- Visitors Car Park (Pay and Display)
- Staff Car Parks
- Bus Stops
- Taxis
- Public Telephones
- Catering Facilities
- Baby Changing Facilities

Preseli

Fulton House

Poster: Taliesin Annex

Main Venue: Taliesin

Workshop: Faraday

Singleton Hospital

Singleton Park

Abbey Meadow

Singleton Park

To Sports Village (pedestrian access)

Bus, Cycle & Pedestrian Access

Porters' Traffic Control Lodge

Traffic Control Barriers

Main Entrance

To City Centre

To Sports Village and Mumbles Gower

BMVC2015

Swansea, UK



BMVA
British Machine
Vision Association