CRiSM workshop on

Contemporary Issues in Hypothesis Testing

15–16 September 2016, University of Warwick

Organisers:
Dalia Chakrabarty (University of Leicester) dc252@leicester.ac.uk
Sara Wade (University of Warwick) s.wade@warwick.ac.uk
Julia Brettschneider (University of Warwick) julia.brettschneider@warwick.ac.uk

Administrator:
Olivia Garcia-Hernandez (University of Warwick) olivia.garcia@warwick.ac.uk

Contents

1 Administrative Details

Webpage

Any further information and announcements will be placed on the workshop webpage: warwick.ac.uk/crism/workshops/hypothesistesting

Getting Here

- Information on getting to the University of Warwick can be found at warwick.ac.uk/about/visiting
- Parking permits can be acquired at no further cost at https://carparking.warwick.ac.uk/events/contemporary-issues-in-hypothesis-training

Registration and Locations

- **Registration** is open 9:00–9:15 and 10:30–11:00, Thursday 15th September, in the main atrium of the Zeeman building. Tea and coffee will be available.
- **Talks** will be held in room MS.04, Zeeman Building. MS.04 is on the second floor, up both sets of main stairs from the atrium.
- **Breakfast** is provided for those with on-campus accommodation.
- **Lunch** is provided on Thursday and Friday in the atrium of the Zeeman building (and undergraduate workroom).
- **Dinner** is not provided (with the exception of the workshop dinner on Thursday, for those who have registered for this). There are several restaurants on campus, see the facilities section below.
- The **poster session and wine reception** is on Thursday from 16:30 to 18:30, in the atrium of the Zeeman building.
- The **workshop dinner** is on Thursday at 19:00, at the Radcliffe restaurant, for those who have registered.
- The **workshop ends** at 17:00 on Friday 16th September.

Accommodation

Accommodation is in en-suite rooms on campus. Keys can be collected from the conference reception in the student union atrium, with the exception of invited speakers who should collect their keys from Radcliffe reception. All rooms have linen and toiletries. Rooms will be available after 15:00 for check in. All bedrooms must be vacated by 9:30am on the day of departure.

Internet Access

- **Campus:** Wireless access is most easily available via eduroam, which is supported across most of the Warwick campus. See eduroam.org.
- **Accommodation:** Wireless access is available, ask for log-in details whenever you check-in to your accommodation.

Facilities

- **Pubs and Resaurants:**
  - Xananas, Students' Union: warwicksu.com/xananas
  - Le Gusta, Arts Centre: warwick.ac.uk/services/retail/legusta
  - The Dirty Duck, Students' Union: warwicksu.com/thedirtyduck
  - For other options, see warwick.ac.uk/services/retail/openingtimes
- **Shop:** Rootes Grocery Store, next to the Students' Union. Open 8am - 8pm.
- **Arts Centre:** warwickartscentre.co.uk
- **Sports Centre:** warwick.ac.uk/sport
- **Health Centre:** uwhc.org.uk
- **Pharmacy:** Students Union Atrium. Open 9am - 6pm.

Telephone Numbers

- **Emergency:** Internal - 22222; External - 024 7652 2222
- **Security:** Internal - 22083; External - 024 7652 2083
- **Department of Statistics:** Internal - 574812; External - 024 7657 4812

Taxis

- **Swift Cabs** 024 7777 7777
- **Trinity Street Taxis** 024 7699 9999

2 Timetable

All talks will take place in room MS.04 in the Mathematics & Statistics Building.

Thursday, 15th September

9:00 – 9:15:	*Registration and coffee*
Morning session	chaired by Joris Mulder
9:15 – 9:30:	Opening remarks by Dalia Chakrabarty
9:30 – 10:30:	Christian Robert
10:30 – 11:00:	*Coffee break and registration*
11:00 – 11:30:	Nick Chater
11:30 – 12:00:	Tom Nichols
12:00 – 12:30:	Sudip Bose (remote)
12:30 – 14:00:	*Lunch*
Afternoon session	chaired by Christian Robert
14:00 – 15:00:	Andrew Gelman (remote)
15:00 – 15:30:	*Coffee break*
15:30 – 16:30:	David Draper
16:30 – 18:30:	Poster session, cheese and wine
19:00 – 22:00:	*Workshop dinner, Radcliffe restaurant*

Friday, 16th September

Morning session	chaired by David Draper
9:30 – 10:30:	Jim Berger
10:30 – 11:00:	*Coffee break*
11:00 – 11:30:	Susan Ellenberg
11:30 – 12:00:	Anne-Laure Bouslateix
12:00 – 12:30:	Alexandra Carpentier
12:30 – 12:45:	Louis Lyons
12:45 – 14:30:	*Lunch*
Afternoon session	chaired by Jim Berger
14:30 – 15:30:	Joris Mulder
15:30 – 16:00:	*Coffee break*
16:00 – 17:00:	Floor opened for discussions.

3 Invited talks, in order of appearance

Testing hypotheses as a mixture estimation problem

Christian P. Robert, Université Paris-Dauphine, University of Warwick, Institut Universitaire de France

We consider a novel paradigm for Bayesian testing of hypotheses and Bayesian model comparison. Our alternative to the traditional construction of posterior probabilities that a given hypothesis is true or that the data originates from a specific model is to consider the models under comparison as components of a mixture model. We therefore replace the original testing problem with an estimation one that focus on the probability weight of a given model within a mixture model. We analyse the sensitivity on the resulting posterior distribution on the weights of various prior modelling on the weights. We stress that a major appeal in using this novel perspective is that generic improper priors are acceptable, while not putting convergence in jeopardy. Among other features, this allows for a resolution of the Lindley-Jeffreys paradox. When using a reference Beta B(a, a) prior on the mixture weights, we note that the sensitivity of the posterior estimations of the weights to the choice of a vanishes with the sample size increasing and advocate the default choice $a = 0.5$, derived from Rousseau and Mengersen (2011). Another feature of this easily implemented alternative to the classical Bayesian solution is that the speeds of convergence of the posterior mean of the weight and of the corresponding posterior probability are quite similar.
Joint work with K. Kamary, K. Mengersen and J. Rousseau
arxiv.org/abs/1412.2044

The psychology of explanation

Nick Chater, Warwick Business School

How does the brain explain the world around us? And how does the brain's approach contrast with scientific and statistical methodology? I suggest that cognition operates using a "cycle of thought:" a slow, sequential process, each step of which involves "focusing" on a very narrow pattern complete problem (e.g., recognizing a face, or a word, or planning an action). The sequentiality of the process arises because each pattern completion draws on a parallel memory search over prior memory traces; such searches can only be conducted one at a time without interference. According to this picture of cognition, the brain (1) considers data sequentially ; (2) considers one hypothesis at a time; (3) but can propose and assess "local" adjustments to that hypothesis; (4) cannot compare very different hypotheses directly, but only via qualitative arguments; (5) has no memory for past data, but just the "explanation" of that data. The memory-based computational approach typically generates "shallow" explanations—each new set of data is "explained" by mapping it on to explanations of prior data. Constructing explicit statistical and scientific methodologies is required precisely because the brain does not operate according to these principles by default.

Large scale evaluation of random field theory inference in fMRI

Thomas E. Nichols, University of Warwick

A fundamental goal in "brain mapping" with functional Magnetic Resonance Imaging (fMRI) is localising the parts of the brain activated by a task. The standard tool for making this inference has been Random Field Theory (RFT), a collection of results for Gaussian Processes of the null statistic image. RFT provides inference on individual voxels (voxel-wise) and sets of contiguous suprathreshold voxels (cluster-wise) while controlling the familywise error rate, the chance of one or more false positives over the brain. I have spent much of my career developing RFT methods as well as complementary resampling-based inference method, always carefully evaluating the methods with simulated Gaussian Process realisations. The focus of the talk will be new, large-scale evaluations with real data. Usually, convincing scientists to spend time collecting null data is a challenge, but in fMRI an entire discipline has evolved around the "resting state". In resting state fMRI, instead of comparing brain activity between states, subjects are asked to lie in the scanner in a state of "resting wakefulness", and the pattern of connectivity between regions is studied. We exploited 1000's of publicly available resting state fMRI datasets, putting them to work as real data realisations of the null hypothesis for putative task fMRI experiments. These massive real data evaluations show that, even with $n = 20$ or 40 subjects, RFT suffers from slightly conservative voxel-wise inferences and sometimes catastrophically liberal cluster-wise inferences. I will discuss the reasons for these failures of RFT and practical solutions going forward.
Joint work with A. Eklund and H. Knutsson.

Frequentist vs. Bayesian testing: when are data more extreme?

Sudip Bose, George Washington University

In comparing Bayesian hypothesis testing with frequentist, several authors, notably Berger, Sellke and Sivaganesan have calculated lower bounds on posterior probabilities and on Bayes factors and compared them with p-values. We consider a different sort of comparison – what sets of data are more or less extreme in Bayesian and frequentist analyses? We present some results for the exponential family.

Hypothesis testing is a bad idea

Andrew Gelman, Columbia University

Through a series of examples, we consider problems with classical hypothesis testing, whether performed using classical p-values or confidence intervals, Bayes factors, or Bayesian inference using noninformative priors. We locate the problem not in the use of any particular statistical method but rather with larger problems of deterministic thinking and a misguided version of Popperianism in which the rejection of a straw-man null hypothesis is taken as confirmation of a preferred alternative. We suggest solutions involving multilevel modeling and informative Bayesian inference.

The Jaynes information criterion (JIC), the role of parsimony in Bayes Factors, and a comparison of predictive and structural Bayesian model choice criteria

David Draper, University of California at Santa Cruz

Bayes factors are one way to compare two models, M_1 (simpler) and M_2, in settings in which M_1 contains a *structural singleton*: a single point in a continuous parameter space that, e.g., is uniquely specified by a scientific theory of interest. The role of parsimony in making such comparisons via Bayes factors is in general implicit, as opposed to the situation with approximate Bayes factors such as *BIC*, in which an explicit tradeoff between goodness of fit and parsimony is evident. In the first part of this talk I'll introduce the *Jaynes Information Criterion (JIC)*, which is a method — related to a proposal by Jaynes (2003) — for deriving Bayes factors that generalizes *BIC*, in the sense that the fit/parsimony tradeoff is always explicit with *JIC* (across all possible parametric prior choices, not just with the unit-information prior specific to *BIC*). In the second part of the talk I'll present a comparison of Bayesian and non-Bayesian model selection criteria based on *predictive* accuracy (AIC, log scores and [interestingly] DIC) with criteria based on identification of correct model *structure* (BIC and more general Bayes factors).

The use of rejection odds and rejection ratios in testing hypotheses

Jim Berger, Duke University

Much of science is (rightly or wrongly) driven by hypothesis testing. Even in situations where the hypothesis testing paradigm is correct, the common practice of basing inferences solely on p-values has been under intense criticism for over 60 years. We discuss, as an alternative, the use of the odds of a correct rejection of the null hypothesis to incorrect rejection. Both pre-experimental versions (involving the power and Type I error) and post-experimental versions (depending on the actual data) are considered. Implementations are discussed that range from depending only on the p-value to consideration of full Bayesian analysis. A surprise is that all implementations even the full Bayesian analysis have a strong frequentist justification. Versions of these techniques can be implemented that require only minor modifications to existing practices, yet overcome some of their most severe shortcomings.

The role of hypothesis testing in clinical trials

Susan S. Ellenberg, University of Pennsylvania

Many statistical methods and tools have been developed to guide scientific decision-making. A tool widely used in medical research, particularly in the area of clinical trials in which one investigates whether one medical approach is superior to another, is hypothesis testing. Methods for hypothesis testing were first developed early in the 20th century, and they have taken hold in medical research because they address the straightforward question that medical researchers have when they undertake a comparative study: does approach A have any advantage over approach B with respect to a particular outcome? A hypothesis test allows us to quantitate the uncertainty about this comparison. A small p-value suggests that our observed outcome would be unlikely under the assumption that neither treatment has any advantage over the other. Of course, there are many other factors to consider in making a decision to choose A or B. One needs to consider the magnitude of the difference; a trivial difference might produce a very low p-value in a very large study. One needs also to consider other relevant outcomes, which may or may not favor the approach shown to have an advantage for the primary outcome. One needs to consider whether the quality of the study conduct supports the credibility of the finding. Most of the objections to hypothesis testing have to do with the tendency of many researchers, journal editors and regulators to view the p-value as the absolute determinant of the study conclusion, rather than as a tool to guide interpretation of the study results.

Can and should the choice of statistical methods be more "evidence-based"?

Anne-Laure Boulesteix, LMU Munich

The goal of medical research is to develop interventions that are, with respect to patient outcome, superior to existing ones in some sense. Similarly, the goal of research in methodological computational statistics is to develop data analysis tools that are superior to existing ones in some sense. Methodological aspects of the evaluation of medical interventions have been devoted a lot of attention in the literature and it is now well-accepted that medicine should be at least partly "evidence-based". Although statisticians (including ourselves) are convinced of the importance of good study designs and evidence-based approaches (in particular, statistical inference) in the context of clinical research, they often tend to ignore these principles when designing their own studies for evaluating statistical methods in the context of their methodological research. In this paper, we draw analogies between clinical trials and real data based benchmarking experiments in methodological statistical science, with datasets playing the role of patients and methods playing the role of medical interventions. Based on this analogy, we suggest directions for potential improvements of study designs for the evaluation of statistical methods based on real data and for better interpretation of these studies, in particular with respect to statistical testing and sample size issues, inclusion criteria for datasets and various types of bias. More generally, we discuss the concept of "evidence-based" statistical research, its limitations and its impact for the design and interpretation of benchmark experiments.
Joint work with A. Hapfelmeier

Rank testing and confidence sets for matrix completion

Alexandra Carpentier, University of Potsdam

In general, there is a strong connection between the problem of model testing, and the problem of constructing adaptive and honest confidence sets. This presentation will be about rank model testing/estimation and adaptive and honest confidence sets for high dimensional, bounded and low rank matrix completion. Two design assumptions will be considered, namely a) that the (noisy) entries of the matrix are sampled uniformly at random and b) that each (noisy) entry of the matrix has a given probability of being revealed. If an additional information on the noise that is added to the entries, e.g. its variance, is not available, then one can prove that although adaptive and honest confidence sets exist in model a), they do not exist in model b). This highlight a fundamental difference between models a) and b), which does not exist in the case of optimal and adaptive estimation of the low rank matrix (where the optimal rates of estimation are the same up to logarithmic factors in both models).
Joint work with O.Klopp, M.Loeffler and R.Nickl

The particle physicists' approach to hypothesis testing

Louis Lyons, High Energy Physics, Imperial College

In searching for new phenomena in Particle Physics,we compare the null hypothesis (just standard known physics) with an alternative (standard physics plus some specific form of new physics, such as the production of supersymmetric particles). This involves several statistical issues. These include our use of p-values, and why we avoid Bayesian methods in searches; our extreme criterion (5 sigma) for discovery claims; the 'modified frequentist method' ($\mathrm{CL}_s = p_1/(1-p_0)$) that is used for excluding alternative hypotheses; the 'Look Elsewhere Effect', etc. These topics are discussed, and are illustrated by the analysis resulting in the discovery of the Higgs Boson at the CERN Large Hadron Collider.

Bayesian hypothesis testing in the social sciences

Joris Mulder, Tilburg University

Researchers in the social and behavioral sciences often formulate competing hypotheses with equality and/or order constraints on the parameters of interest. The goal is then to test these hypotheses using the observed data. Bayes factors have proven useful for this testing problem because (i) Bayes factors can be straightforwardly used for testing multiple nonnested hypotheses in a direct manner; (ii) Bayes factors automatically balance between fit and complexity; and (iii) Bayes factors have an intuitive interpretation as the relative evidence in the data between two hypotheses. All these properties are not shared by the Fisherian p-value, the dominant testing criterion in social research. This talk consists of two parts. In the first part, a Bayes factor is proposed for a multiple hypothesis test on bivariate, ordinal and partial correlations. In the second part, information consistency is investigated for various Bayes factor tests in normal linear models.

4 Posters

Bayesian model selection for the validation of computer codes

Pierre Barbillon, AgroParisTech / INRA

Poster board 1

Complex physical systems are increasingly modeled by computer codes which aim at predicting the reality as accurately as possible. During the last decade, code validation has benefited from a large interest within the scientific community because of the requirement to assess the uncertainty affecting the code outputs. Inspiring from past contributions to this task, a testing procedure is proposed in this paper to decide either a pure code prediction or a discrepancy-corrected one should be used to provide the best approximation of the physical system. In a particular case where the computer code depends on uncertain parameters, this problem of model selection can be carried out in a Bayesian setting. It requires the specification of proper prior distributions that are well known as having a strong impact on the results. Another way consists in specifying non-informative priors. However, they are sometimes improper, which is a major barrier for computing the Bayes factor. A way to overcome this issue is to use the so-called intrinsic Bayes factor (IBF) in order to replace the ill-defined Bayes factor when improper priors are used. For computer codes which depend linearly on their parameters, the computation of the IBF is made easier, thanks to some explicit marginalization. In the paper, we present a special case where the IBF is equal to the standard Bayes factor when the right-Haar prior is specified on the code parameters and the scale of the code discrepancy. On simulated data, the IBF has been computed for several prior distributions. A confounding effect between the code discrepancy and the linear code is pointed out. Finally, the IBF is computed for an industrial computer code used for monitoring power plant production.

Distinguishing distributions with interpretable features

Wittawat Jitkrittum, Gatsby Unit, UCL

Poster board 2

Two semimetrics on probability distributions are proposed, given as the sum of differences of expectations of analytic functions evaluated at spatial or frequency locations (i.e, features). The features are chosen so as to maximize the distinguishability of the distributions, by optimizing a lower bound on test power for a statistical test using these features. The result is a parsimonious and interpretable indication of how and where two distributions differ locally. An empirical estimate of the test power criterion converges with increasing sample size, ensuring the quality of the returned features. In real-world benchmarks on high-dimensional text and image data, linear-time tests using the proposed semimetrics achieve comparable performance to the state-of-the-art quadratic-time maximum mean discrepancy test, while returning human-interpretable features that explain the test results.

Pragmatic evaluation of the performance of valid and invalid stopping rules for hypothesis tests

Michael J. Lew, University of Melbourne

Poster board 3

In the current debates about the reliability of scientific results it is common to nominate inappropriate statistical testing procedures among the problems and while it is true that statistical practices could be improved, usual discussions of the possible contribution of hypothesis testing to the problem focus mainly on false positive errors. However, most of the analytical practices that increase false positive errors concurrently reduce false negative errors and or the sample size and so any evaluation of the performance of hypothesis testing procedures should include considerations of both types of error as well as sample size. In this study the performance of Student's t-test with a variety of stopping rules was evaluated pragmatically at a range of sample sizes, effect sizes and, in contrast to common practice, using a variety of loss functions. The results using dichotomous outcomes of each run indicate that, under many reasonable circumstances, informal procedures with optional stopping are superior to the standard procedure where sample size is fixed in advance of seeing the data, despite yielding false positive errors more frequently than the nominal rate. In other words, the increased power from optional stopping more than compensated for the increased rate of false positives for a range of loss functions and effect sizes. A formal sequential test that yields exactly the nominal false positive error rate was also found to be superior, and, in contrast to the informal optional stopping rule procedures, it was never inferior to the fixed sample size rule test. Advice that might follow from a more complete evaluation of the performance of hypothesis testing differs from the usual advice to apply frequentist type I error control procedures more rigorously or to eschew hypothesis testing altogether.

Automatic Bayes factors for testing equality and inequality constrained hypotheses on variances

J. Mulder, Tilburg University

Poster board 4

In comparing characteristics of independent populations, researchers frequently expect a certain structure of the population variances. These expectations can be formulated as hypotheses with equality and/or inequality constraints on the variances. In this article we consider the Bayes factor for testing such (in)equality constrained hypotheses on variances. Application of Bayes factors requires specification of a prior under every hypothesis to be tested. However, specifying subjective priors for variances based on prior information is a difficult task. We therefore consider so-called automatic or default Bayes factors. These methods avoid the need for the user to specify priors by using information from the sample data. We discuss three automatic Bayes factors for testing variances. The first is a Bayes factor with equal priors on all variances, where the priors are specified automatically using a small share of the information in the sample data. The second is the fractional Bayes factor, where a fraction of the likelihood is used for automatic prior specification. The third is an adjustment of the fractional Bayes factor such that the parsimony of inequality constrained hypotheses is properly taken into account. Results from a simulation study indicate that the adjusted fractional Bayes factor converges fastest to the true hypothesis.
Joint work with F. Böing-Messing.

Confidence sets - Going beyond voxel-level and cluster-level null hypothesis testing

Thomas Nichols, University of Warwick

Poster board 5

Null hypothesis testing is the foundation of brain mapping but critics have often raised the issue of the "null hypothesis fallacy", that the null hypothesis is never true. Recent studies using 100 fMRI sessions per subject have illustrated this problem, finding effects everywhere in the brain. While limited sample size prevents discovery of such universal activation (or deactivation), more "big data" studies are providing such power (e.g. Human Connectome Project, $N = 1,200$; UK Biobank final $N = 100,000$). With such sample sizes, traditional null hypothesis testing is of limited value. In this work we apply recent work (Sommerfeld, Sain, Schwartzman, 2015) to develop confidence sets (CSs) on clusters. Whereas traditional voxel- or cluster- wise inferences indicate where the null, i.e. an effect size of 0, can be rejected, the CSs are statements about non-zero effect sizes analogous to confidence intervals. They operate on either raw units or (standardised) effect sizes (e.g. Cohen's d). For a cluster constructed with cluster-forming threshold c, the CSs comprise two sets of voxels: The upper CS is smaller, giving the voxels we infer to be truly larger than c; the larger lower CS is best described by its complement – all voxels outside this set we infer to be truly smaller than c. We describe the method and apply it to datasets from the Human Connectome Project, demonstrating the value of this spatial inference method.
Joint work with A. Bowring, A. Schwartzman, and M. Sommerfeld.

A new Bayesian test to test for the intractibility-countering null

Kangrui Wang, University of Leicester

Poster board 6

This poster discusses a new Bayesian test of hypothesis, that tests for the null, when the likelihood is intractable outside the null model. Thus, the null can be considered to be the simplifying assumption that can counter intractability of the more complex model, (that the simpler null model is nested within). Bayes Factors are shown to be known up to a ratio of unknown data-dependent constants in such a situation. The main instrument of use in this new test is the "generated data", that is generated from the model in which the null is true. Thus, the support in the measured data for the null can be given by the ratio of the marginalised posteriors of the model parameter given the measured, and that given the generated data. However, when we are asked to compare support for a null in one data, to another data of a different size, or compare supports in a given data for null models in which model parameters have different dimensionalities, the ratio mentioned above is seen to confound interpretation. In such applications, we define support in a measured data for a null by identifying parameter values that are as or more consistent with the measured data than is minimally possible given the generated data, and marginalising the posterior of the model parameter, over such values. Application to galactic data is undertaken to illustrate differential support in two sets of such data for the hypothesis that the galactic state space is isotropic in shape, thus explaining the difference in the results of performing unsupervised learning of the galactic gravitational mass density using such data sets.
Joint work with C. Spire and D. Chakrabarty.

5 Participant List

Name	Affiliation	Email
Rodrigo Abrunhosa Collazo	University of Warwick	R.A.Collazo@warwick.ac.uk
Alejandra Avalos Pacheco	University of Warwick	A.Avalos-Pacheco@warwick.ac.uk
Pierre Barbillon	AgroParisTEch	pierre.barbillon@agroparistech.fr
Jim Berger	Duke University	berger@stat.duke.edu
Concha Bielza	ETSI Informaticos - UPM	mcbielza@fi.upm.es
Anne-Laure Boulesteix	LMU Munich	boulesteix@ibe.med.uni-muenchen.de
Julia Brettschneider	University of Warwick	Julia.Brettschneider@warwick.ac.uk
Franois-Xavier Briol	University of Warwick	F-X.Briol@warwick.ac.uk
Terry Brown	University of Warwick	t.brown.1@warwick.ac.uk
Alexandra Carpentier	University of Potsdam	alexandra.carpentier5@gmail.com
Dalia Chakrabarty	University of Leicester	d.chakrabarty@warwick.ac.uk
Nick Chater	University of Warwick	Nick.Chater@wbs.ac.uk
David Draper	University of California	draper@ucsc.edu
Susan Ellenberg	University of Pennslyvania	sellenbe@upenn.edu
John Fenlon	University of Warwick	J.Fenlon@warwick.ac.uk
Habib Ganjgahi	University of Warwick	H.Ganjgahi@warwick.ac.uk
Simon Gates	University of Warwick	s.gates@warwick.ac.uk
Helen Hodges	Swansea University	h.r.hodges.807913@swansea.ac.uk
Yue Hu	University of Leicester	yh214@le.ac.uk
Wittawat Jitkrittum	University College London	wittawat@gatsby.ucl.ac.uk
Pedro Larranaga	ETSI Informaticos - UPM	pedro.larranaga@fi.upm.es
Fabrizio Leisen	University of Kent	fl79@kent.ac.uk
Michael Lew	University of Melbourne	michaell@unimelb.edu.au
Mike Lonergan	University of Dundee	m.lonergan@dundee.ac.uk
Louis Lyons	High Energy Physics, Imperial College	l.lyons@physics.ox.ac.uk
Jason Madan	University of Warwick	j.j.madan@warwick.ac.uk
Joris Mulder	Tilburg University	j.mulder3@uvt.nl
Tom Nichols	University of Warwick	t.e.nichols@warwick.ac.uk
Eric Parent	AgroParisTEch	eric.parent@agroparistech.fr
Rui Paulo	University of Lisboa/CEMAPRE	rui@iseg.ulisboa.pt
Christian Robert	University of Warwick	xian@ceremade.dauphine.fr
Adam Sanborn	University of Warwick	a.n.sanborn@warwick.ac.uk

Iqbal Shamsudheen	University College London	iqbal.shamsudheen.14@ucl.ac.uk
Cedric Spire	University of Leicester	dc252@leicester.ac.uk
Nigel Stallard	University of Warwick	N.Stallard@warwick.ac.uk
Simone Tiberi	University of Warwick	S.Tiberi@warwick.ac.uk
Cristiano Villa	University of Kent	cv88@kent.ac.uk
Sara Wade	University of Warwick	S.Wade@warwick.ac.uk
Kangrui Wang	University of Leicester	kw202@le.ac.uk
Samuel Watson	University of Warwick	s.i.watson@warwick.ac.uk


CENTRAL CAMPUS
Cryfield Sports Pavilion
Radcliffe
Scarman
Warwick Business School (WBS)
Gatehouse
University House
Tennis Centre
Games Hall
Arden
Ramphal
Physics
Physical Sciences
Social Sciences
Chemistry
Sciences
IARC
Engineering
Bus Interchange
Library
Modern Records Centre
Arts Centre
Chaplaincy
Rootes Building
Rootes
Senate House
Humanities
IMC
Students' Union
ATM
Psychology
Computer Science
Mathematics
Zeeman
IDL
IIPSI
Millburn House
EMB
Bluebell
Rootes
Sports Centre
Health Centre
Jack Martin
Tocil
Arthur Vick
Claycroft
Argent Court
WESTWOOD WAY
GIBBET HILL ROAD
SCARMAN ROAD
UNIVERSITY ROAD
KIRBY CORNER ROAD
LIBRARY ROAD
SERVICE ROAD ONLY
LEIGHFIELD ROAD
HEALTH CENTRE ROAD
MILLBURN HILL ROAD
SIR WILLIAM LYONS ROAD
LYNCHGATE ROAD
A45, M42, M5, M6 BIRMINGHAM
M6, M69, M1 LONDON
A46, M40 WARWICK, KENILWORTH, LEAMINGTON SPA
KEY TO SYMBOLS
Accommodation
Bus stop
Construction zones
Information totem
Male, female, disabled
Male, female, disabled (inc. baby change)
Pay & display parking
Multi-storey pay & display parking
Pay & display parking (Without accessibility parking)
Shopping on campus
Sustrans cycle route
Taxi rank