Final report: Modern statistical approaches to off-equilibrium modelling for nonlinear system control (GR/M76379/01)

Roderick Murray-Smith,[1] **Mike Titterington,**[2] **Ken Hunt**[3]
Departments of Computing Science,[1] Statistics[2] and Mech. Eng.,[3], University of Glasgow
E-mail: `rod@dcs.gla.ac.uk` WWW: `http://www.dcs.gla.ac.uk/~ rod/MSAFCproject.htm`

We believe that this has been a successful project, bringing together work by statisticians, engineers and computer scientists, and applying their diverse expertise to challenging application areas. Major achievements include: mixture of Gaussian process (GP) models for efficiently representing heterogeneity, use of Gaussian process models in control, development of user interfaces which can present changing conditional distributions to the user in a natural manner in real time. 18 papers were published as a result of this project, and 1 patent relating to the audio feedback mechanism has been filed.

1 Background/Context

The improvement in available computing power and the recent growth in understanding of practical Markov-Chain Monte Carlo (MCMC) methods for performing the high-dimensional integrations needed in complex models has made the use of sophisticated statistical tools far more feasible in engineering applications. The use of such approaches in the engineering literature, however, is remarkably rare, due probably to lack of exposure to the modern methods. In terms of model structure selection, an area which has caused many problems in the application of multiple model methods, the model structure identification approach (how many local models should be used, of which dynamic order, and how should the state-space be partitioned) is currently being solved using classical statistical methods such as generalised cross-validation, which are far from optimal. The MCMC approach offers an alternative to selecting a single 'best' model. Instead of this we create a probability distribution covering a range of model structures. Related work using MCMC methods with locally linear models can be found in Holmes and Mallick (1999). Further, recent developments in, or rediscovery of, the use of Gaussian process priors for classification and regression, e.g. O'Hagan (1978); Williams (1998), have interesting ramifications for the multiple model approach.

Multiple model approaches are increasingly popular in the engineering community; see Murray-Smith and Johansen (1997) for examples coming from neural network, statistical and control communities. However, there are a number of open problems. In Shorten *et al.* (1999) we show that the methods have serious problems when modelling *transient* or *off-equilibrium* conditions in dynamic systems, and that other claimed interpretability benefits are often far less clear in practice. The problem of dealing with off-equilibrium data in control problems is awakening a great deal of interest Murray-Smith *et al.* (1999); Leith and Leithead (1999). Current engineering implementations of multiple models tend to use very basic estimates of model uncertainty, from a statistical viewpoint, with the notable exception of Waterhouse *et al.* (1996). This has had detrimental effects on the validity of the modelling process itself, but has been even more damaging to the use of such models in control system development. It should be noted that these theoretical results are backed by experience in our industrial work, where the problems led to many time-consuming and expensive experiments with the physical system, and many iterations in the modelling phase. Putting the modelling aspects of this framework on to a firmer footing would have an immediate impact on the scientific literature, and would quickly be applied in industry, especially by the industrial collaborators in this project.

In Murray-Smith *et al.* (1999) we outline how the multiple model approach can be fruitfully compared and combined with computationally-intensive statistical approaches such as Gaussian process priors. The main advantage of the use of Gaussian processes for modelling dynamic systems is that we have better estimates of model prediction variance (and variance of derivatives of the model, which is very useful in control engineering contexts), without necessarily resorting to MCMC simulation. This provides a useful bridge to engineers who find it easier to relate to models obtained by likelihood maximisation, than by integration. The computational cost associated with Gaussian process priors is less important in transient regions, which typically have very sparse data.

2 Key Advances and Supporting Methodology

The practical aim of this project was to apply advanced, data- and computationally-intensive statistical tools to improve the utility of widely-used engineering tools. The research involved a number of strands, as follows.

1. Development of novel nonparametric modelling approaches inspired by challenges faced in rehabilitation applications
2. Use of Gaussian process priors in control system design
3. Experimental evaluation of suitability of Gaussian process models for modelling dynamic systems found in engineering.

4. Extension of Gaussian process models to include information commonly used in engineering applications, and to make implementation more efficient.
5. Gaussian process priors in control system design
6. Development of user-interfaces, which can cope with uncertainty, based on probabilistic models.

2.1 Development of novel nonparametric modelling approaches inspired by challenges faced in rehabilitation applications

The rehabilitation engineering application area chosen was that of modelling standing-up behaviour supported by Functional Electrical Stimulation (FES) of the lower limbs, in spinal-cord injury patients. The rehabilitation engineering application provides a number of challenges for modelling work. The systems modelled are strongly time-varying, at a range of different timescales, are nonlinear, and show strong inter-patient and inter-session (even for the same patient) variation. This variation among batches of data is described as *heterogeneity* in the statistics literature. A model composed of hierarchical mixture of Gaussian Process priors was developed, using Markov-Chain Monte Carlo (MCMC) algorithms to perform inference, and described in Shi *et al.* (2002*c*). This represents an interesting development of the multiple-model approach to modelling for application to batch data. This has significant computational efficiency advantages, and also copes well with the heterogeneity present in data from different sources, such as from multiple patients. We also develop a method to give a better estimate of prediction variance and then give a better prediction interval for data from different sources. This was extended to include classification with hierarchical mixtures of GP's Shi *et al.* (2002*a*, 2003). This classification work is promising for classification of motion phases, which can guide design of hybrid control algorithms.

A development of birth-death approaches to MCMC methods for unknown number of components in mixture models and Hidden Markov Models is described in Shi *et al.* (2002*b*). Software for the above methods was implemented in a mix of MATLAB and C to allow us to perform experimental evaluation of the models.

Some further new topics emerged as the project progressed, for example, the mixture of Gaussian processes with an allocation model, which can use patients' personal information, such as their age, weight, level of injury etc; and functional regression analysis with Gaussian process can address the mean and correlation relationship in a single model. Follow-on research projects based on this work are in preparation.

2.2 Experimental evaluation of suitability of Gaussian process models for modelling dynamic systems found in engineering

Roman Kamnik, a visiting scientist at the CRE in Glasgow, from the University of Ljubljana, who took over the rehabilitation component of the research, generously provided FES-supported paraplegic standing data from experiments in the group in Ljubljana, one of the world's leading rehabilitation groups. We used these data to make an experimental comparison of the new hierarchical Gaussian process prior models with more established multi-layer perceptron artificial neural network models used previously within the group in Ljubljana. The Gaussian process prior models produced better results than the neural networks. Another aim of the study was to compare multiple model structures for the evaluation of minimal sensor requirements for field use of feedback FES schemes. This was published in Kamnik *et al.* (2003).

Henrik Gollee, the co-director of the Centre for Rehabilitation Engineering at Glasgow, also advised us on aspects of these experiments, and implemented the audio feedback mechanisms (described below) in the CRE lab, on an arm-cranking device.

Ongoing research involves the extension of the hierarchical mixture model to include an allocation model, which incorporates further information about the patient, and importantly to use the hierarchical nature of the model to help with more effective on-line experiment design, to reduce the amount of data acquisition needed for new patients. This is especially important for the rehabilitation research side, as patients will not be keen to cooperate if the data acquisition phase is unnecessarily stressful.

2.3 Extension of Gaussian process models to include information commonly used in engineering applications, and to make implementation more efficient.

The covariance functions typically used in previous work were not always suitable for use in dynamic systems models, and the amount of data used was typically too great to use with GP models, where the computational load increases as $O(N^3)$, for N data points. We developed a number of extensions, which made it clear to engineers that they could still use much of their existing prior knowledge. Covariance functions for the correlated noise typical of closed-loop control problems were developed - an example with ARMA noise models is presented in Murray-Smith and Girard (2001).

k-step-ahead predictions are inherent to simulation of dynamic systems. We developed methods for GPs which allowed us to make predictions with uncertain inputs, assuming a Normal distribution Girard *et al.* (2002, 2003), which allowed us to approximately propagate the uncertainty in simulation. This was then extended to include training GP priors with uncertain inputs Girard and Murray-Smith (2003).

Engineers often think about systems in terms of the properties of local linearisations of those systems. We used the fact that the derivative of a Gaussian process remains a Gaussian process to provide a way of fusing data points with derivative 'observations'. This allows engineers to incorporate existing models or prior knowledge, and also absorbs the large proportion of data which are close to equilibrium conditions, leaving the off-equilibrium data which tend to be much more sparse. Very importantly, this also provides a consistent way of interpolating between multiple linearisations, with guarantees on integrability, something which other approaches such as Tagaki-Sugeno fuzzy models or gain-scheduled controllers have not achieved. This is described in Solak *et al.* (2003) and Leith *et al.* (2002).

The work in Solak *et al.* (2003) and Leith *et al.* (2002) builds on our original work Murray-Smith *et al.* (1999), showing how problems in modelling *transient* or *off-equilibrium* conditions are reduced when using GP models, instead of multiple-model approaches.

A generalisation of such transformations of latent variables, described by GPs, is presented in Murray-Smith and Pearlmutter (2003) - this generalises the derivative example to general transformations which once discretised become linear mappings. It provides an elegant, general framework for sensor fusion, and it can also allow the computational load to be mitigated from $O(N^3)$ to $O(N^2)$.

Work is in progress on application of the GP priors with derivative observations to helicopter dynamics models - this is a challenging test for the approach.

2.4 Gaussian process priors in control system design

Use of Gaussian process models in control system design was a major area of interest for this project, and we made significant process in this area, developing and investigating the first control systems based on GPs. This work was done with colleagues abroad, in Chile, Germany and Slovenia. The initial development of a one-step-ahead adaptive, nonlinear controller based on a GP is described in Murray-Smith and Sbarbaro (2002), which provided a good example of the robust adaptation, and implicit regularisation which comes with the GP. We extended this to include priors on hyperparameters, and multi-step-ahead control, in Murray-Smith *et al.* (2003), with an expanded form submitted for publication Sbarbaro and Murray-Smith (2003). Together with colleagues at the Institut Josef Stefan, we implemented Model Predictive Control with GPs and tested this on simulated process plants Kocijan *et al.* (2003*a*).

An exciting feature of the GP-based controller is the way the effective order of the controller can vary in a continuous fashion, throughout the state-space, depending on the conditional density of the GP model - a novel idea within control theory. Such state-dependent estimates of uncertainty are especially important in transient, off-equilibrium areas.

A tutorial introduction to modelling of dynamic systems, together with experiments with a real twin-tank set-up, was published in Gray *et al.* (2003)

2.5 Development of user-interfaces which can cope with uncertainty

In this part of the project we developed user-interfaces for systems incorporating probabilistic models which allow natural feedback of the uncertainty of any variables to the user. This is potentially of great use in training users how to make certain moves as part of rehabilitation. It is important as a tool for providing body state estimates for spinal cord injury patients who have greatly reduced proprioception, and it will provide a novel way for disabled users to interact with computing devices.

The basic audio-feedback idea is presented in Williamson and Murray-Smith (2002), and was implemented on an arm-cranking rehabilitation device, a paper describing the results of which is in preparation. The ideas were further developed to create a dynamic systems approach to text entry using a mobile computer with accelerometers. The handling qualities of the device vary with the probability of interpreted intentions conditioned on recent actions Williamson and Murray-Smith (2003). This is an important alternative to the standard keyboard and screen used on current mobile phones, and is potentially a very exciting technology for disabled users. The dynamics model used in the device can be configured to suit their individual behaviour, which often is well outside distributions for the typical user population, using the models developed in this project.

We are also using the feedback technique in a follow-on project (EPSRC grant GR/R79234/01) on *Integrated voluntary control of standing in paraplegia* to provide interesting experiment designs.

While suited to feedback from motor control models, this is a very general approach, which we are now developing in diverse areas such as gestural interaction with mobile phones, and display of context-of-use for mobile devices.

2.6 Milestones

Publications associated with each of the milestones in the original grant application are as follows:

- MS1.1 - see Solak *et al.* (2003); Murray-Smith *et al.* (2002)
- MS1.2 - see Shi *et al.* (2002*b*)
- MS1.3 - see Murray-Smith and Girard (2001); Solak *et al.* (2003)
- MS1.4 - see Murray-Smith and Sbarbaro (2002); Sbarbaro and Murray-Smith (2003); Murray-Smith *et al.* (2003); Kocijan *et al.* (2003*b*)
- MS2.1 The models have been implemented in software in C and Matlab and used to provide the results in the publications. We are currently making them available as stand-alone code on the Internet.
- MS2.2 The original plan to incorporate the code in Dr. Johansen's ORBIT package was dropped, as the focus of this project moved from multiple-model approaches to GP approaches and Dr. Johansen moved from SINTEF, who owned the software. We therefore felt that stand-alone packages made more sense.
- MS2.3 The feedback from the experimental stages occurred throughout the project and led to developments such as the hierarchical GP models for batch and large data sets Shi *et al.* (2002*c*,*a*), the birth-death MCMC for integrating over model structures Shi *et al.* (2002*b*), the transformations of GPs for sensor fusion and large data sets Murray-Smith and Pearlmutter (2003) and the audio feedback mechanisms for improving patient awareness Williamson and Murray-Smith (2002).
- MS3.1 Shi *et al.* (2002*c*) described early work on the rehabilitation engineering application.
- MS3.2 Kamnik *et al.* (2003); Murray-Smith and Sbarbaro (2002); Sbarbaro and Murray-Smith (2003). A tutorial example summing up this experience is presented in Gray *et al.* (2003) Further publications in the area of helicopter dynamics will be completed this year.

3 Project Plan Review

Participants from Prof. Hunt's group included Roman Kamnik, a visiting scientist from the University of Ljubljana, who took over the rehabilitation component of the research, and Henrik Gollee, the co-director of the Centre for Rehabilitation Engineering, who interacted in some of the experimental set-ups and provided supervision on the rehabilitation aspects.

Use of staff varied slightly from the initial plan, because of the early departure of the main post-doctoral researcher, Jian Qing Shi (a statistician), after 26 months, to a permanent lecturing position at Newcastle. He continues to interact strongly with all investigators, and is frequently in Glasgow, so this did not disadvantage the project but provided opportunities to augment the main research in the original proposal, by employing two researchers with different skills for the remaining time.

John Williamson, a computing science graduate, focused on making the experimental software developed by Dr. Shi more robust (in terms of software implementation), and easier to use by new users. He also developed some innovative approaches to audio display of uncertainty in model behaviour, which we expect to be an excellent tool in providing feedback for training rehabilitation patients - especially those with spinal cord injuries who have no proprioception.

Dr. Gary Gray, a post-doctoral researcher with experience in helicopter system identification provided some very useful experience in applying the theoretical ideas to engineering practice. He tested the basic software on an experimental set-up in the Dept. of Electrical Engineering (the twin-tanks experiment), and wrote a tutorial case-study for readers with an engineering background Gray *et al.* (2003). He also started the implementation of the use of Gaussian processes and derivative observations and off-equilibrium data, for modelling helicopter dynamics, together with colleagues in the Dept. Aerospace engineering. This work is ongoing, with results expected by the end of the summer.

The original proposal had also been planned to fit with requested money for a Ph.D. student, but this was not funded. To make up the difference, we interacted strongly with colleagues at other institutes. Significant interaction included that with colleagues at Strathclyde University and the Hamilton Institute in Ireland, (D. Leith, W. Leithead, B. Pearlmutter), the University of Concepcion, Chile (D. Sbarbaro), the Gatsby Computational Neuroscience Unit/Max-Planck Gesellschaft (C. Rasmussen), and the Josef Stefan Institute (J. Kocijan).

4 Explanation of Expenditure

Expenditure was much as expected. Use of money allocated for staff computing equipment and software, and for travel to international conferences (Neural Information Processing Systems (NIPS) Conference, Int. Conference Current Advances and Trends in Nonparametric Statistics), workshops (Multi-Agent Control summer school, Slovenia, and EPSRC Maynooth Workshop on nonlinear dynamics, Dublin).

5 Research Impact

The publication of control based on nonparametric Bayesian models has exposed engineers to new ways of thinking about their modelling problems, and showed how control can be made more intuitively adaptive and robust by using flexible models such as GPs. The outcome leads to state-dependent degrees of freedom in the controller which is a new concept within the control world. The computational challenges involved in working with the practical data sets encouraged developments on the statistical algorithm side, like the hierarchical GP models, and the use of transformations of GPs.

5.1 Applications beneficiaries

Rehabilitiation patients stand to gain directly from this work, as the collaboration has led to more accurate models than the rehabilitation engineers were using before (the previous work was based on neural networks), and has made a contribution to the design of sensory FES with minimal sensing equipment Kamnik *et al.* (2003). Further improvements from the batch nature of the hierarchical GP approach include potential improvements in data acquisition and experiment design.

The novel approaches to interface design, based on audio feedback from conditional densities, has already been implemented in the Centre for Rehab. Eng. in Glasgow. It is a promising new method of improving patient proprioception and helping to train patients in motor control tasks. The methods also have the potential to benefit the general population as an interaction paradigm for mobile phones. We have filed a patent related to this area, and we are currently negotiating with manufacturers. We have also obtained funding from EPSRC to continue and expand this research (Audioclouds project, GR/R98105). Other aspects are already feeding into follow-on projects (EPSRC grant GR/R79234/01, *Integrated Voluntary Control of Unsupported Paraplegic Standing*, and EPSRC grant GR/R30730/01 *Development of Systems For Tetraplegic Arm Cranking Using Functional Electrical Stimulation: A Pilot Study*).

5.2 Further Research and Dissemination activities

As planned in the original proposal, we presented the research at a range of conferences, from the statistics (e.g. Int. Conference Current Advances and Trends in Nonparametric Statistics, the 5th ICSA international conference in Hong Kong), engineering (IFAC world congress, IFAC System Identification Conference, IEE Signal processing workshop) and computing science (Neural Information Processing Systems, Vancouver and the Gesture Workshop, Padua) areas. Papers have also been submitted to a diverse range of journals (Int. Journal Adaptive Signal Processing, Statistics and Computing , Automatica, Mathematics and Computers in Simulation, IEEE transactions on Rehabilitation Engineering), in order to bring the diverse groups together. Preliminary work was typically published as departmental technical reports to allow early access. As well as the papers already published, papers awaiting results for acceptance include the following.

- R. Murray-Smith, B. Pearlmutter, Transformations of Gaussian process priors, Submitted to NIPS 16.
- J. Williamson, R. Murray-Smith, Dynamics and probabilistic text entry, Submitted to NIPS 16.
- A. Girard, R. Murray-Smith, Learning a Gaussian process prior model with uncertain inputs, Submitted to NIPS 16.
- D. Sbarbaro, R. Murray-Smith, Self-tuning control of non-linear systems using Gaussian process prior models, Submitted to Automatica.
- J.Q. Shi, R. Murray-Smith, and D.M. Titterington, Hierarchical Gaussian process mixtures for regression, Submitted to Statistics and Computing.
- R. Kamnik, J. Q. Shi, R. Murray-Smith, T. Bajd, Nonlinear modelling of FES-supported standing up in paraplegia for selection of feedback sensors, Submitted to IEEE Transactions on Rehabilitation Engineering.

6 Conclusions

We believe that this has been a successful project which has brought together diverse skills, and has overcome the inherent differences in terminology and approach among statisticians, control engineers and computer scientists. The research has already led to a large number of published results, many of which were not forseen at the start of the project. The collaboration expanded during the course of the project to include a number of international colleagues, and will continue in follow-on projects.

References

Girard, A. and R. Murray-Smith (2003). Learning a Gaussian process prior model with uncertain inputs. Technical Report TR-2003-144. University of Glasgow, Scotland, UK.

Girard, A., C. E. Rasmussen and R. Murray-Smith (2002). Multiple-step ahead prediction for nonlinear dynamic systems - a gaussian process treatment with propagation of the uncertainty. Technical Report TR-2002-1194. University of Glasgow, Scotland, UK.

Girard, A., C. E. Rasmussen, J. Quinonero-Candela and R. Murray-Smith (2003). Gaussian process priors with uncertain inputs – application to multiple-step ahead time series forecasting. In: *Advances in Neural Information processing Systems 15* (S. Becker, S. Thrun and K. Obermayer, Eds.).

Gray, G., R. Murray-Smith, K. Thomson and D. J. Murray-Smith (2003). Investigation of Gaussian process modelling applied to laboratory equipment. Technical Report TR-2003-151. University of Glasgow, Scotland, UK.

Holmes, C. C. and B.K. Mallick (1999). Bayesian piecewise linear regression using multivariate linear splines. Technical Note, Imperial College, London.

Kamnik, R., J. Q. Shi, R. Murray-Smith and T. Bajd (2003). Nonlinear modelling of FES-supported standing up in paraplegia for selection of feedback sensors. Technical Report TR-2003-150. University of Glasgow, Scotland, UK.

Kocijan, J., A. Girard, B. Banko and R. Murray-Smith (2003*a*). Dynamic systems identification with Gaussian processes. In: *4th Mathmod conference, Vienna, Int. Association. for Mathematics and Computers in Simulation*.

Kocijan, J., R. Murray-Smith, C.E. Rasmussen and B. Likar (2003*b*). Predictive control with Gaussian process models. In: *IEEE Eurocon 2003: The International Conference on Computer as a Tool*.

Leith, D. and W. Leithead (1999). Analytic framework for blended multiple model systems using linear local models. *International Journal of Control* **72**(7/8), 605–619.

Leith, D. J., W. E. Leithead, E. Solak and R. Murray-Smith (2002). Divide & conquer identification: Using Gaussian process priors to combine derivative and non-derivative observations in a consistent manner. In: *Conference on Decision and Control*.

Murray-Smith, R. and A. Girard (2001). Gaussian Process priors with ARMA noise models. In: *Irish Signals and Systems Conference, Maynooth*. pp. 147–152.

Murray-Smith, R. and B. Pearlmutter (2003). Transformations of Gaussian process priors. Technical Report TR-2003-149. University of Glasgow, Scotland, UK.

Murray-Smith, R. and D. Sbarbaro (2002). Nonlinear adaptive control using non-parametric Gaussian process prior models. In: *15th IFAC World Congress on Automatic Control, Barcelona*.

Murray-Smith, R. and T. A. Johansen (1997). *Multiple Model Approaches to Modelling and Control*. Taylor and Francis, London.

Murray-Smith, R., D. Sbarbaro, C. E. Rasmussen and A. Girard (2003). Adaptive, cautious, predictive control with Gaussian process priors. In: *IFAC International Symposium on System Identification, Rotterdam*.

Murray-Smith, R., R. Shorten and D. Leith (2002). Nonparametric models of nonlinear dynamics. In: *IEE Workshop of Nonlinear and Non-Gaussian signal processing -N2SP, Peebles* (C. Cowan, Ed.).

Murray-Smith, R., T. A. Johansen and R. Shorten (1999). On transient dynamics, off-equilibrium behaviour and identification in blended multiple model structures. In: *European Control Conference, Karlsruhe, 1999*. pp. BA–14.

O'Hagan, A. (1978). On curve fitting and optimal design for regression (with discussion). *Journal of the Royal Statistical Society B* **40**, 1–42.

Sbarbaro, D. and R. Murray-Smith (2003). Self-tuning control of non-linear systems using gaussian process prior models. Technical Report TR-2003-143. University of Glasgow, Scotland, UK.

Shi, J. Q., R. Murray-Smith and D. M. Titterington (2002*a*). Bayesian regression and classification using mixtures of Gaussian processes. Technical Report TR-2002-114. University of Glasgow, Scotland, UK.

Shi, J. Q., R. Murray-Smith and D. M. Titterington (2002*b*). Birth-death MCMC methods for mixtures with an unknown number of components. Technical Report TR-2002-117. University of Glasgow, Scotland, UK.

Shi, J. Q., R. Murray-Smith and D. M. Titterington (2002*c*). Hierarchical Gaussian process mixtures for regression. Technical Report TR-2002-107. University of Glasgow, Scotland, UK.

Shi, J.Q, R. Murray-Smith and D. M. Titterington (2003). Bayesian regression and classification using mixtures of multiple Gaussian processes. *International Journal of Adaptive Control and Signal Processing* **17**(2), 149–161.

Shorten, R., R. Murray-Smith, R. Bjørgan and H. Gollee (1999). On the interpretation of local models in blended multiple model structures. *International Journal of Control* **72**(7/8), 620–628.

Solak, E., R. Murray-Smith, W. E. Leithead, D. J. Leith and C. E. Rasmussen (2003). Derivative observations in Gaussian process models of dynamic systems. In: *Advances in Neural Information processing Systems 15* (S. Becker, S. Thrun and K. Obermayer, Eds.).

Waterhouse, S., D. MacKay and T. Robinson (1996). Bayesian methods for mixtures of experts. In: *Neural Information Processing Systems - 8*. MIT press. Cambridge, MA. pp. 351–357.

Williams, C. K. I. (1998). Prediction with Gaussian processes: From linear regression to linear prediction and beyond. In: *Learning and Inference in Graphical Models* (M. I. Jordan, Ed.). pp. 599–621. Kluwer.

Williamson, J. and R. Murray-Smith (2002). Audio feedback for gesture recognition. Technical Report TR-2002-127. University of Glasgow, Scotland, UK.

Williamson, J. and R. Murray-Smith (2003). Dynamics and probabilistic text entry. Technical Report TR-2003-147. University of Glasgow, Scotland, UK.