Guided Tour of my Publications

Stephen Luttrell
www.stephenluttrell.com

September 22, 2011

Brief Profile

I joined QinetiQ (i.e. its predecessor the Royal Signals and Radar Establishment) as a higher scientific officer in 1982 and began my research into information processing. I was promoted to senior principal scientific officer (individual merit) in 1991 in recognition of the powerful adaptive processing techniques that I had developed. Since then I have being leading a long-term research programme into various aspects of adaptive information processing, focusing particularly on self-organising networks. I have published about 60 refereed papers (in journals, conferences, books), plus *many* unpublished technical reports.

Preamble

By default references are to journal papers, and where there are references to conference papers, technical reports, or other material this is explicitly indicated.

Publications

1981-1982: PhD

Papers [23, 24, 25] and PhD dissertation [127] apply quantum chromodynamics to modelling the quark/gluon structure of hadrons seen in deep inelastic scattering experiments. Various higher order effects that modify naïve perturbative results are studied both phenomenologically and theoretically. The operator algebra techniques used in this PhD research turn out to be very useful for the description of Markov chain Monte Carlo methods, as described in my publications from 2005 onwards.

1984-1989: Bayesian Super-Resolution and Mutual Information (scattered field)

Report [102] describes how linear least squares error reconstruction using an appropriately weighted reconstruction space can be used to introduce prior knowledge to enhance the resolution of coherent images. The material in this report was not published because it was quickly overtaken by the full Bayesian treatment described below.

Paper [22] describes how Bayes' theorem may be used to derive the solution to linear inverse problems under the assumption of Gaussian probability density functions (PDF). For suitable Bayesian priors this leads to super-resolution where details on a scale shorter than the Rayleigh resolution length become visible.

Conference papers [46, 50] describe how a parallel implementation of this super-resolution algorithm can be implemented for analysing synthetic aperture radar (SAR) images. Book chapter [61] and report [101] present an introduction to the use of prior knowledge in the analysis of SAR images.

Papers [20, 21] introduce the principle of mutual information maximisation as a way of optimising the extraction of information from data.

Paper [19] is on clutter and targets in SAR images. My contribution showed how to use a generalisation of the super-resolution techniques originally developed in [22] to analyse images of targets, and showed how to use mutual information (as in [20, 21]) to interpret the information processing in terms of information channels. A brief description of this work is in [15]. A practical implementation of this work is described in [16]. This was patented in [126].

Report [98] describes an analysis of images of point targets in images produced by the Royal Signals and Radar Establishment (RSRE) SAR. The purpose of this was to accurately calibrate the point spread function (PSF) of the RSRE SAR so that its images could be robustly super-resolved. However, this analysis led me to the conclusion that for the RSRE SAR the PSF was itself a function of the data; in other words the system response was *non-linear*. This totally invalidated the super-resolution theory developed thus far, which assumed that the system response was *linear*, which then triggered a shift of my research away from super-resolution. A side effect of this non-linearity was that the higher moments of the image data were biased away from their naïve values, and thus should *not* have been used in their uncorrected form to deduce the statistics of the underlying clutter model.

Book chapter [60] (delayed for several years in publication) reviews the ideas behind the use of information channels, drawing together ideas from mutual information, super-resolution, and Bayesian inference.

1985-1988: Markov Random Fields for Clutter and Texture Modelling

Paper [18] is the first in a long series on Markov random fields (MRF). It describes how to do Monte Carlo sampling of arbitrarily complicated MRFs by bit-flipping operations that could be easily implemented in hardware.

Report [100] is the first description of how to train MRF models using a generalisation of the Boltzmann machine learning algorithm that I called the Gibbs machine. Various generalisations were discussed such as the idea of enhancing an MRF model by using these learning techniques to attach a "brain graft" to the MRF model to patch it up wherever it did not give a good enough approximation to the data.

Paper [17] presents a careful analysis of the use of MRF models for texture modelling, using arguments based on sufficient statistics to understand how the texture information is spread out amongst the various measured statistics. This is conceptually similar to the use of information channels discussed in [19]. Contact with image processing techniques is made via the grey level co-occurrence matrix method and the WISARD n-tuple processing network, both of which are special cases of this MRF analysis.

Paper [14] shows how the maximum entropy method can be used to pick good statistics to measure in textured images (as introduced in [17]), so that an efficient MRF texture model can be built. An efficient algorithm (which does *not* require the fixing of Lagrange multipliers, unlike other approaches) is given for deciding what new statistics to add to the MRF model by comparing *synthetic* textures generated using the current MRF model with the *real* texture to be modelled. As a corollary, a generalisation of the original Boltzmann machine learning algorithm is given for arbitrary MRFs.

Conference papers [51, 54] develop various aspects of the use of MRFs for texture and clutter modelling.

Conference paper [47] compares and contrasts the MRF approach (as in [14]) and a proposed hierarchical "cluster-decomposition" approach to modelling PDFs. The generalised form of the Boltzmann machine learning algorithm is presented in detail. The advantages of the cluster decomposition approach (particularly that it does *not* use Monte Carlo simulations) are explained.

1989-1991: Bayesian Super-Resolution (scatterers themselves)

Paper [13] extends the earlier Bayesian super-resolution method (see [22]) by pointing out the fact that the inverse problem to solve should be the reconstruction of the object *scatterers* that produce the object field, and *not* the reconstruction of the object *field* itself. In paper [8] and report [97] this work is greatly extended by using the expectation-maximisation (EM) method to derive an iterative object reconstruction algorithm from first principles. This allows both super-resolution and auto-focusing to be handled within the same framework. This work was reviewed in [7].

1988-1989: Hierarchical Self-Organising Maps

Paper [12] describes a hierarchical generalisation of the standard Kohonen self-organising map (SOM) network (also reported in conference paper[48]), in which a tree of linked SOMs is trained to encode SAR images (also reported in conference paper [49]). This paper also introduces the idea of the "growing grid" method of training SOMs. In [11] this approach is described in more detail.

1986-1990: Digital Receiver

Paper [10] and report [99] describe the analysis, design and testing of a finite state machine for demodulating signals. My contribution was the analysis. This is essentially an early type of software radio receiver. This was patented in [125].

1989-1992: Distortion Minimising Encoders (single SOM)

Paper [9] presents a novel derivation of SOM networks (also reported in conference papers [44, 45]), based on a generalisation of the standard Linde-Buzo-Gray (LBG) vector quantisation algorithm to account for distortion on the communication channel, which approximates the standard Kohonen SOM algorithm as a special case (the distribution of distortions corresponds to the topographic neighbourhood function). The biggest advantage of the approach used is that it is based on minimisation of an objective function (unlike Kohonen's SOM), which thus allows many properties to be theoretically derived. For example, the density of code vectors in this type of SOM (for 1-dimensional data) is derived in paper [6] and report [96], where it is shown to be the *same* as in a standard vector quantiser (unlike the standard Kohonen SOM which leads to a *different* density which inconveniently depends on the choice of topographic neighbourhood function). A generalisation of this result to data of arbitrary dimensionality was given in report [93].

1988-1994: Lateral Mutual Information Maximising Hierarchical Encoders

Conference paper [43] presents a detailed analysis of the optimisation of the hierarchical "cluster decomposition" introduced in conference paper [47]. The objective function used is the relative entropy between the model PDF and the true PDF, which is equivalent to the *sum* of the (lateral) mutual informations between various nodes in each layer of the network. This is successfully applied to the detection of anomalies in textures images as described in report [95] and in the unpublished arXiv papers [122, 123], and has been patented [124].

1990-1992: Distortion Minimising Encoders (coarse-grain parallel SOMs)

Paper [5] uses the communication channel model of SOMs (see [9]) to analyse networks of connected SOMs (this was also described in conference paper [42] and report [94]). The model used is one in which two communication channels mutually interfere thus causing distortion that is correlated between the channels. This distortion defines the topographic neighbourhood functions that are used in the SOMs (that are the channel encoders) in the first place. The optimisation of the SOMs is influenced by mutual interactions between their outputs, hence the description of this type of training as "self-supervised".

1988-1994: Some Consolidation

Conference paper [41] analyses adaptive n-tuple networks (e.g. WISARD) by using a relative entropy objective function to optimise a PDF model. Both supervised and unsupervised methods of training such networks emerge naturally from this analysis.

Conference paper [40] reviews the Bayesian approach to training and using neural networks. Various models are discussed including the "adaptive cluster expansion" model originally introduced in conference paper [47].

Conference paper [36] presents a rigorous Bayesian analysis of the "adaptive cluster expansion" model.

1992-1994: Partitioned Mixture Distributions

Paper [4] describes a generalisation (the *partitioned* mixture distribution, or PMD) of the standard mixture distribution used to model PDFs (this was also described in conference papers [37, 39] and report [92]). In image processing (the statistical properties of) each correlation area of an image can be modelled with a mixture distribution, and it would be convenient if all of these models could be collected together in a single translation invariant architecture. A nice solution to this problem is to build each mixture distribution using components selected from a common pool of mixture components, where each correlation area has a complete repertoire of components needed for its analysis. Mixture distributions that see overlapping areas of image have a lot of overlap in the components they use. A PMD can be trained using a generalisation of the EM method for training standard mixture distributions. Trained PMDs have many of the properties that are observed in the low-level visual cortex, such as a complete repertoire of processing machinery for each local patch of image.

1990-1994: Folded Markov Chains

Paper [3] gives a Bayesian analysis of the SOM model originally described in [9]. The key to the analysis is the folded Markov chain (FMC) that was described in report [91], in which information is passed along a Markov chain, and then Bayes' theorem is used to pass (virtually) the information back along the chain in the opposite direction until it reaches the start again. This final reconstruction of the input to the chain is compared with the original input, and the transitions in the Markov chain are adjusted to minimise the Euclidean reconstruction distortion (on average). The tight constraints that Bayes' theorem imposes on the various probabilities make it possible to derive a training algorithm that reduces to the algorithm in [9] when the encoder is *deterministic*. This FMC analysis unifies all my earlier work on networks of SOMs, and by allowing *probabilistic* encoders to be handled in the same framework, it sets the scene for future work.

1993-1995: Some Consolidation

Reports [89, 90] describe various ways in which mixture distributions can be used in the analysis of images.

Conference paper [35] and report [87] shows the application of SOM techniques to the analysis of data (e.g. range profiles of radar returns from ships) that lies on a manifold with a circular topology. Prior knowledge of the topology is used to define an appropriate topographic neighbourhood function for use in the SOM.

1994-1997: Unification of PMD and FMC

Book chapter [56] (delayed by several years in publication), reports [83, 86, 88] and unpublished arXiv paper [120] combine the FMC approach (introduced in [3]) with the PMD approach (introduced in [4]) to create an encoder structure that is appropriate for the analysis of images, where the form of the PMD posterior probability allows the encoder to break into local pieces that each encode a local patch of the image. This is then applied to the analysis of images derived from multiple sources. In the case of *pairs* of images the network learns a structure that closely resembles the dominance stripes and orientation maps observed in the visual cortex. The important point is that you can obtain these "biological" results as a consequence of the general properties of encoders – to emphasise this the ancronym VICON (VIsual COrtex Network) was coined in [83, 120].

Conference papers [33, 34] present a detailed analysis of some of the properties of PMDs (introduced in [4]). A first order perturbation analysis reveals the low-order properties of PMDs and allows their behaviour to be related to various neural networks. A dynamical PMD (which has the same relationship to a static PMD that a hidden Markov model has to a mixture distribution) is analysed in detail (this was also presented in report [85]), and is successfully applied to the problem of tracking a weak target in clutter. Some of the results to be presented in [2] below are quoted without proof.

Conference paper [32], report [84] and the unpublished arXiv paper [121] apply the FMC approach (introduced in [3]) to optimising encoders that output several statistically independent (given the input) codes. Various analytic optimum solutions are obtained. However, these solutions are *less* useful than those that are obtained if the optimisation is constrained in various ways, which is the focus of all subsequent papers in this area.

1997: "State of the Nation"

Report [82] summarises all publications produced during the 1994-1997.

1997: Isaac Newton Neural Networks Programme

Report [81] and unpublished arXiv paper [119] describes the work I did at the Neural Networks research programme at the Isaac Newton Institute. I analysed the optimisation an objective function depending on the *joint* PDF of the state of a multi-layer network, and showed that this unified all of my work on the optimisation of hierachical encoders.

1994-1997: Distortion Minimising Encoders (fine-grain parallel)

Paper [2] gives the generalisation of the FMC approach in [3] to the case where multiple codes are output by a *probabilistic* encoder for each of its input vectors. This was also presented earlier in book chapters [58, 59]. Up to this point in my work the encoders have each output a *single* code, and thus each acts as a winner-take-all network. The advantage of allowing the possibility of *multiple* codes is that the encoded information can be split across more than one channel. This is conceptually similar to the use of information channels discussed in [19]. Further, if these codes are chosen stochastically, then by optimising the FMC (minimum Euclidean distortion) the network can decide for itself how many information channels it needs to use. This freedom allows the network to optimise its own architecture. In [2] a simple application of these ideas is analysed in detail, where various (supposedly) optimal choices of information channel are compared with each other, and the conditions under which each type of solution is optimal are established. An important cross-disciplinary result is that the above extension to *multiple* codes allows contact to be made with neural network models based on discrete firing events, where the recent history of which neurons have fired corresponds to the set of codes that are currently active.

Paper [1] applies the multiple sampling ideas in [2] to the hierarchical vector quantisation network [11]. Although the hierarchical network of [11] has its structure manually defined, rather than learnt via optimisation (as *could* be done using [2]), the earlier ideas in [11] automatically fit into the framework of the later theory in [2].

Change in Publication Style

From here on there are currently no further journal papers. However, the research continues to be published in conference papers and technical reports.

1997: Some Consolidation

Report [80] describes various ways of training self-organising encoder networks by modelling their output as the posterior probability over samples (or "firing events") drawn from a stochastic code book. The different training schemes correspond to different conditions under which the samples are drawn.

1998-1999: Distortion Minimising Encoders (analytic optimisation)

Conference paper [31] shows how *analytic* optimisation of a stochastic encoder (an FMC with multiple output codes) leads to optimal encoders being described by *piecewise linear* posterior probabilities. This emerges because the reconstruction is a *linear* superposition of contributions, and the objective function is a *Euclidean* distortion. This piecewise property also holds in networks of linked stochastic encoders. This piecewise linearity will prove to be very useful for deriving analytic solutions to various problems.

Book chapter [57], report [79], and the unpublished arXiv paper [118] use the piecewise linear property of optimal stochastic encoders (see [31]) to derive optimal encoding schemes for circular and toroidal manifolds. The results derived in this paper were obtained analytically by using Mathematica to manipulate the various cumbersome piecewise linear expressions leading to compact results. The circle and the torus are simple curved manifolds that serve as models for the more complicated curved manifolds that arise in signal and image processing. For the toroidal manifold *two* types of optimal encoder emerge depending on the size of the encoder and the number of stochastic codes that it is allowed to output. On the one hand a *joint* encoder is optimal if there are more than a certain minimum number of codes to choose from; this corresponds to an encoder that has so many resources that it can simply partition the torus into localised code cells. On the other hand a *factorial* encoder emerges if the number of codes to choose from is severely reduced *and* the number of codes it is allowed to output is sufficiently large; this corresponds to an encoder that is starved of resources but which is allowed to have several trials at outputting a code, and which therefore chooses to use anisotropic code cells to slice the torus into localised code cells defined by the regions of intersection of *pairs* of anisotropic code cells. The automatic emergence of factorial coding is a crucial property in a network that aims to discover *for itself* what information channels to use when processing data.

1999: Some Consolidation

Report [78] is a user's guide to stochastic encoders, and contains many simple examples of their use. Unpublished arXiv paper [117] is a comprehensive summary of the theory and many numerical simulations of stochastic encoders, and unpublished arXiv paper [116] contains a small subset of these results.

2000-2003: Distortion Minimising Encoders (jammer suppression)

Conference paper [30] applies a stochastic encoder to the problem of separating a signal from a jammer. Because the jammer is much stronger than the signal the optimisation concentrates on encoding the jammer alone. This allows the trained encoder to separate the jammer and signal subspaces (in this case these are *curved* subspaces), which allows the signal to be cleanly detected. This material is also covered in conference papers [28, 29], book chapter [55], report [77], and the unpublished arXiv papers [114, 115].

2001: Distortion Minimising Encoders (classifier fusion)

Conference paper [27] demonstrates some of the key properties of stochastic encoders. One of these is their ability to learn factorial codes in which a code book breaks into separate smaller code books, each of which encodes only one part of the input. Another is the natural way in which supervision can be introduced to steer the code books towards coding their inputs in particular ways. This is all presented in the context of multiple classifier fusion, where the processing is split across multiple information channels and then fused together again at the end.

2001-2004: Internal Reports (part 1)

Report [76] describes how to use a stochastic encoder to suppress sea clutter in coherent images, which allows targets that are masked by the clutter to be revealed. The approach used is essentially the same as for separating a signal from a jammer (as in conference paper [30]).

Report [75] compares and contrasts the use of encoders and the use of PDFs for modelling (and suppressing) sea clutter in images.

Report [68, 74] summarises all of my research results on self-organising networks.

Report [69, 70, 73] describes how to apply self-organising networks to the problem of combining the outputs of multiple classifiers.

Reports [71, 72] describe various interesting and useful properties of the self-organising network known as ACEnet® (Adaptive Cluster Expansion network).

2003: Isaac Newton Neural Networks Programme (published, at last!)

Conference paper [26] introduces a new way of understanding networks of linked stochastic encoders. These ideas were originally described in report [81]. Rather than modelling each encoder as an FMC (see [3] and its generalisation to multiple codes in [2]) and building up the overall network objective function as a *sum* of local contributions, the *whole* network of linked encoders is deemed to have a *single* objective function which compares the joint PDF of the whole network under two different modelling assumptions. For a chain-like network the two PDF models are built from conditional probabilities acting in opposite directions along the chain, and it is similar to Dayan and Hinton's Helmholtz machine *except that* the aim here is to optimise the *joint* network probability rather than only the *marginal* input probability. This approach is then shown to reduce to the FMC approach wherever that had been used earlier. An application of these ideas to training a network on hierarchically correlated data is presented, and the encoders (the bottom-up PDF models) automatically split into smaller encoders in precisely the way that is expected (i.e. process the most uncorrelated pieces separately, then progressively fuse the results until only the most correlated piece is left). This approach is thus capable of making *structural* changes to the network as it is trained, where it creates information channels for doing the lower-level processing, and then progressively fuses these channels to do the higher-level processing.

2003-2004: Distortion Minimising Encoders (multiple correlation scales)

Unpublished arXiv paper [113] describes the use of self-organising stochastic encoders for learning the structure of data manifolds. The main aim of this paper was to demonstrate the automatic separation of correlated components in the data, and their subsequent fusion whilst preserving only their dominant degree(s) of freedom, as was described in [26]. This is a key paper that opens up a very fruitful area of research.

2005-2007: Internal Reports (part 2)

Report [67] comprehensively reviews MRFs and SONs from the perspective of adaptive filters and recurrent networks, and introduces an operator algebra to describe the associated Markov chain Monte Carlo (MCMC) algorithms.

Reports [65, 66] apply the operator algebra methods of [67] to recurrent networks, dealing with both the small-occupancy (i.e. dominated by quantisation effects) and large-occupancy (i.e. quantisation effects negligible) cases.

Report [64] describes a proposed application of recurrent SONs to the problem of language identification.

Report [63] introduces symbolic algebra techniques for manipulating the operator algebra of MCMC algorithms.

Report [62] presents a Mathematica implementation of recurrent SONs.

2006: Discrete Network Dynamics

Unpublished arXiv paper [112] describes the use of operator algebra techniques to reformulate MCMC algorithms, so that the algorithms can be manipulated by purely algebraic methods, which would have applications in adaptive networks. This is a key paper that opens up a very fruitful area of research.

Online Publications

An on-line archive of some of my publications (the Crown Copyright ones) can be found at www.stephenluttrell.com/papers.

References

Journal Papers

[1] Luttrell S P, 1997, A discrete firing event analysis of the adaptive cluster expansion network, Network: Computation in Neural Systems, vol. 7, no.2, pp. 285-290

[2] Luttrell S P, 1997, Self-organisation of multiple winner-take-all neural networks, Connection Science, vol. 9, no.1, pp. 11-30

[3] Luttrell S P, 1994, A Bayesian analysis of self-organising maps, Neural Computation, vol. 6, no.5, pp. 767-794

[4] Luttrell S P, 1994, Partitioned mixture distribution: an adaptive Bayesian network for low-level image processing, IEE Proceedings on Vision, Image,and Signal Processing, vol. 141, no.4, pp. 251-260

[5] Luttrell S P, 1992, Self-supervised adaptive networks, IEE Proceedings-F, vol. 139, no.6, pp. 371-377

[6] Luttrell S P, 1991, Code vector density in topographic mappings: scalar case, IEEE Transactions on Neural Networks, vol. 2, no.4, pp. 427-436

[7] Luttrell S P, 1991, The theory of Bayesian super-resolution of coherent images: a review, International Journal of Remote Sensing, vol. 12, no.2, pp. 303-314

[8] Luttrell S P, 1990, A Bayesian derivation of an iterative autofocus/super-resolution algorithm, Inverse Problems, vol. 6, no.6, pp. 975-996

[9] Luttrell S P, 1990, Derivation of a class of training algorithms, IEEE Transactions on Neural Networks, vol. 1, no.2, pp. 229-232

[10] Luttrell S P and Pritchard J A S, 1989, Rapid acquisition of low signal-to-noise carriers, IEE Proceedings-I, vol. 136, no.1, pp. 100-108

[11] Luttrell S P, 1989, Hierarchical vector quantisation, IEE Proceedings-I, vol. 136, no.6, pp. 405-413

[12] Luttrell S P, 1989, Image compression using a multilayer neural network, Pattern Recognition Letters, vol. 10, no.1, pp. 1-7

[13] Luttrell S P, 1989, The inverse cross section problem for complex data, Inverse Problems, vol. 5, no.1, pp. 35-50

[14] Luttrell S P, 1988, A maximum entropy approach to sampling function design, Inverse Problems, vol. 4, no.3, pp. 829-841

[15] Luttrell S P and Oliver C J, 1988, The role of prior knowledge in coherent image processing, Philosophical Transactions of the Royal Society of London Series A, vol. 324, pp. 397-407

[16] Delves L M, Pryde G C and Luttrell S P, 1988, A super-resolution algorithm for SAR images, Inverse Problems, vol. 4, pp. 681-703

[17] Luttrell S P, 1987, The use of Markov random field models to derive sampling schemes for inverse texture problems, Inverse Problems, vol. 3, no.2, pp. 289-300

[18] Luttrell S P, 1986, An optimisation of the Metropolis algorithm for multibit Markov random fields, Inverse Problems, vol. 2, no.2, pp. L15-L17

[19] Luttrell S P and Oliver C J, 1986, Prior knowledge in synthetic-aperture radar processing, Journal of Physics D: Applied Physics, vol. 19, no.3, pp. 333-356

[20] Luttrell S P, 1985, The use of transinformation in the design of data sampling schemes for inverse problems, Inverse Problems, vol. 1, no.3, pp. 199-218

[21] Luttrell S P, 1985, A new method of sample optimisation, Optica Acta, vol. 32, no.3, pp. 255-257

[22] Luttrell S P, 1985, Prior knowledge and object reconstruction using the best linear estimate technique, Optica Acta, vol. 32, no.6, pp. 703-716

[23] Luttrell S P and Wada S, 1982, The current product expansion of the two quark process at the twist four level, Nuclear Physics, vol. B197, pp. 290 306

[24] Luttrell S P, Wada S and Webber B R, 1981, The Wilson coefficient functions of four quark operators and the four quark process in deep inelastic scattering, Nuclear Physics, vol. B188, pp. 219-232

[25] Luttrell S P and Wada S, 1981, Heavy quark production in electromagnetic deep inelastic scattering and QCD analysis of isosinglet moments, Nuclear Physics, vol. B182, pp. 381-396

Conference Papers

[26] Luttrell S P, 2003, A Markov chain approach to multiple classifier fusion, Proceedings of 4th International Workshop on Multiple Classifier Systems (Guildford, UK), Lecture notes in computer science (Springer-Verlag, ed. Windeatt T and Roli F), vol. 2709, pp. 217-226

[27] Luttrell S P, 2001, A self-organising approach to multiple classifier fusion, Proceedings of 2nd International Workshop on Multiple Classifier Systems (Cambridge, UK), Lecture notes in computer science (Springer-Verlag, ed. Kittler J and Roli F), vol. 2096, pp. 319-328

[28] Luttrell S P, 2000, Using stochastic encoders to discover structure in data, Digest of the 5th IMA International Conference on Mathematics in Signal Processing (Warwick, UK)

[29] Luttrell S P, 2000, Invariant stochastic encoders, Digest of the 5th IMA International Conference on Mathematics in Signal Processing (Warwick, UK)

[30] Luttrell S P, 2000, Self-organised discovery of structure in signal manifolds, Proceedings of 2nd DERA Workshop on Sensor Array Signal Processing (DERA Technical Report, DERA/S&E/SPI/TR000349), pp. 88-93

[31] Luttrell S P, 1999, An adaptive network for encoding data using piecewise linear functions, Proceedings of 9th International Conference on Artificial Neural Networks (Edinburgh, UK), IEE Conference Publication, pp. 198-203

[32] Luttrell S P, 1997, Optimal response functions in a network of discretely firing neurons, Proceedings of 5th International Conference on Artificial Neural Networks (Cambridge, UK), IEE Conference Publication, pp. 216-220

[33] Luttrell S P, 1997, Partitioned mixture distributions: a first order perturbation analysis, Proceedings of 5th International Conference on Artificial Neural Networks (Cambridge, UK), IEE Conference Publication, pp. 280-284

[34] Luttrell S P, 1997, Partitioned mixture distributions: the dynamical case, Proceedings of 5th International Conference on Artificial Neural Networks (Cambridge, UK), IEE Conference Publication, pp. 59-63

[35] Luttrell S P, 1995, Using self-organising maps to classify radar range profiles, Proceedings of 4th International Conference on Artificial Neural Networks (Cambridge, UK), IEE Conference Publication, pp. 335-350

[36] Luttrell S P, 1994, The cluster expansion: a hierarchical density model, Proceedings of 14th International Workshop on Maximum Entropy and Bayesian Methods (Cambridge, UK), (Kluwer, ed. Skilling J and Sibisi S), pp. 269-278

[37] Luttrell S P, 1994, The partitioned mixture distribution: multiple overlapping density models, Proceedings of 14th International Workshop on Maximum Entropy and Bayesian Methods (Cambridge, UK), (Kluwer, ed. Skilling J and Sibisi S), pp. 279-286

[38] Luttrell S P, 1993, An adaptive Bayesian network for texture modelling, Proceedings of IEE Seminar on Texture Analysis in Radar and Sonar (London, UK), pp. 6/1-6/10

[39] Luttrell S P, 1993, An adaptive Bayesian network for low-level image processing, Proceedings of 3rd International Conference on Artificial Neural Networks (Brighton, UK), IEE Conference Publication, pp. 61-65

[40] Luttrell S P, 1992, Adaptive Bayesian networks, Proceedings of SPIE Conference on Adaptive and Learning Systems (Orlando, USA), (ed. Sadjadi F A), vol. 1706, pp. 140-151

[41] Luttrell S P, 1992, Gibbs distribution theory of adaptive n-tuple networks, Proceedings of 2nd International Conference on Artifical Neural Networks (Brighton, UK), (North-Holland, ed. Aleksander I and Taylor J), pp. 313-316

[42] Luttrell S P, 1991, Self-supervised training of hierarchical vector quantisers, Proceedings of 2nd International Conference on Artifical Neural Networks (Bournemouth, UK), IEE Conference Publication, pp. 5-9

[43] Luttrell S P, 1991, A hierarchical network for clutter and texture modelling, Proceedings of SPIE Conference on Adaptive Signal Processing (San Diego, USA), (ed. Haykin S), vol. 1565, pp. 518-528

[44] Luttrell S P, 1989, Hierarchical self-organising networks, Proceedings of 1st International Conference on Artifical Neural Networks (London, UK), IEE Conference Publication, pp. 2-6

[45] Luttrell S P, 1989, Self-organisation: a derivation from first principles of a class of learning algorithms, Proceedings of 3rd International Joint Conference on Neural Networks (Washington DC, USA), pp. II/495-II/498

[46] Pryde G C, Delves L M and Luttrell S P, 1989, The performance of a parallel super-resolution algorithm for synthetic aperture radar images, Proceedings of Conference on Mathematics in Signal Processing (Warwick, UK), pp. 739-747

[47] Luttrell S P, 1988, The use of Bayesian and entropic methods in neural network theory, Proceedings of 8th International Workshop on Maximum Entropy and Bayesian Methods (Cambridge, UK), (Kluwer, ed. Skilling J), pp. 363-370

[48] Luttrell S P, 1988, Self-organising multilayer topographic mappings, Proceedings of 2nd International Conference on Neural Networks (San Diego, USA), pp. I/93-I/100

[49] Luttrell S P, 1988, Image compression using a neural network, Proceedings of International Geoscience and Remote Sensing Symposium on Remote Sensing: Moving Towards the 21st Century (Edinburgh, UK), pp. 1231-1238

[50] Pryde G C, Delves L M and Luttrell S P, 1988, A comparative study of the AMT DAP and of transputer array architectures for the super-resolution of synthetic aperture radar images, Proceedings of Conference on Parallel Processing (Manchester, UK), pp. 340-350

[51] Luttrell S P, 1987, Markov random fields: a strategy for clutter modelling, Proceedings of the AGARD Conference on Scattering and Propagation in Random Media (Rome, Italy), AGARD Conference Preprint, pp. 7.1-7.8

[52] Luttrell S P, 1987, The use of Markov random field models in sampling scheme design, Proceedings of SPIE Conference on Inverse Problems in Optics (The Hague, The Netherlands), (ed. Pike R), vol. 808, pp. 182-188

[53] Luttrell S P, 1987, Designing Markov random field structures for clutter modelling, Proceedings of International Conference on Radar (London, UK), IEE Conference Publication, pp. 222-226

[54] Luttrell S P, 1987, Markov random field image models of targets and clutter, Proceedings of the 1st US/UK Workshop on Parallel Processing (Shrivenham, UK), pp. 786-793

Book Chapters

[55] Luttrell S P, 2002, Using stochastic vector quantisers to characterise signal and noise subspaces, in Mathematics in Signal Processing (Clarendon Press, ed. McWhirter J G and Proudler I K), pp. 193-204

[56] Luttrell S P, 2000, The emergence of dominance stripes and orientation maps in a network of firing neurons, in Information Theory and the Brain (CUP, ed. Baddeley R, Hancock P and Földiák P), pp. 101-121

[57] Luttrell S P, 1999, Self-organised modular neural networks for encoding data, in Combining Artificial Neural Networks: Ensemble and Modular Multi-Net Systems (Perspectives in Neural Computing, Springer-Verlag, ed. Sharkey A J C), pp. 235-263

[58] Luttrell S P, 1997, A theory of self-organising neural networks, in Mathematics of Neural Networks: Models, Algorithms and Applications (Operations Research/Computer Science Interfaces, Kluwer, ed. Ellacott S W, Mason J C and Anderson I J), pp. 240-244

[59] Luttrell S P, 1995, Designing analysable networks, in Handbook of Neural Computation (IOP and OUP, ed. Fiesler E and Beale R), pp. B5.3:1-B5.3:8

[60] Luttrell S P, 1989, Inference Theory, in Mathematics in Remote Sensing (The Institute of Mathematics and its Applications Conference Series, Clarendon Press, ed. Brooks S R), pp. 205-222

[61] Luttrell S P and Oliver C J, 1985, Prior knowledge in synthetic aperture radar (SAR) processing, in RSRE Research Review (RSRE, ed. Grant A J, Jakeman E, and Clarke J), pp. 73-77

QinetiQ, DERA, DRA, RSRE Reports

[62] Luttrell S P, 2007, Recurrent self-organising networks: a software implementation, QinetiQ technical report, QinetiQ/SPS/CSIP/TR0703548

[63] Luttrell S P, 2006, Application of symbolic algebra techniques to Bayesian information processing, QinetiQ technical report, QinetiQ/D&TS/SS/CR0602193

[64] Luttrell S P, 2006, Proposed use of recurrent self-organising networks for language identification, QinetiQ technical report, QinetiQ/D&TS/SS/TWP057427

[65] Luttrell S P, 2006, Recurrent self-organising networks: small occupancy case, QinetiQ technical report, QinetiQ/SPS/CSIP/TR057352

[66] Luttrell S P, 2006, Recurrent self-organising networks: large occupancy case, QinetiQ technical report, QinetiQ/SPS/CSIP/TR057273

[67] Luttrell S P, 2005, Review of recurrent networks, QinetiQ technical report, QinetiQ/SPS/CSIP/TR050877

[68] Luttrell S P, 2004, Self-organising information processing: research position update, QinetiQ technical report, QinetiQ/S&E/APC/CR041282

[69] Luttrell S P, 2004, A self-organising approach to optimally combining classifiers: part 3, QinetiQ technical report, QinetiQ/S&E/APC/CR040126

[70] Luttrell S P, 2003, A self-organising approach to optimally combining classifiers: part 2, QinetiQ technical report, QinetiQ/S&E/APC/CR022003

[71] Luttrell S P, 2002, Multi-layer self-organising networks using ACEnet: part 2, QinetiQ technical report, QinetiQ/S&E/APC/CR021835

[72] Luttrell S P, 2002, Multi-layer self-organising networks using ACEnet, QinetiQ technical report, QinetiQ/S&E/APC/CR021071

[73] Luttrell S P, 2002, A self-organising approach to optimally combining classifiers, QinetiQ technical report, QinetiQ/S&E/SIP/CR011805

[74] Luttrell S P, 2001, Self-organising information processing: research position, QinetiQ technical report, QinetiQ/S&E/SIP/CR011926

[75] Luttrell S P, 2001, A comparison of self-organising network and probability modelling approaches to sea clutter suppression, DERA technical report, DERA/S&E/SPI/TR010477

[76] Luttrell S P, 2001, A self-organising network for sea clutter suppression, DERA technical report, DERA/S&E/SPI/TR000513

[77] Luttrell S P, 2000, Invariance discovery by vector quantisation of noisy data, DERA technical report, DERA/S&P/SPI/TR000070

[78] Luttrell S P, 1999, A user's guide to stochastic encoder/decoders, DERA technical report, DERA/S&P/SPI/TR990290

[79] Luttrell S P, 1998, Encoding data from a curved manifold, DERA technical report, DERA/S&P/SPI/TR98116

[80] Luttrell S P, 1998, A Bayesian analysis of complementary approaches to training self-organising stochastic networks, DERA technical report, DERA/CIS/CIS5/TR97165

[81] Luttrell S P, 1997, A unified theory of density models and auto-encoders, DERA technical report, DERA/CIS/CIS5/TR97303

[82] Luttrell S P and Webber C J S, 1997, Spatio-temporal inference theory: final report (1 April 1994-31 March 1997), DERA technical report, DRA/CIS(SE1)/651/FUN/STIT/RP/4

[83] Luttrell S P, 1996, A self-organising visual cortex network (VICON) for processing data from two imaging sensors, DRA technical report, DRA/CIS(SE1)/651/FUN/STIT/RP/3

[84] Luttrell S P, 1996, Optimal posterior probabilities for self-organised neural networks, DRA technical report, DRA/CIS(SE1)/651/FUN/STIT/RP/1

[85] Luttrell S P, 1995, Dynamical partitioned mixture distributions: an introduction, DRA technical report, DRA/CIS(SE1)/651/11/RP/4

[86] Luttrell S P, 1995, A componential self-organising neural network, DRA technical report, DRA/CIS(SE1)/651/11/RP/3

[87] Luttrell S P, 1995, Range profile classification using topographic mappings, DRA technical report, DRA/CIS(SE1)/651/11/RP/2

[88] Luttrell S P, 1995, A self-organising network for processing data from multiple sensors, DRA technical report, DRA/CIS(SE1)/651/11/RP/1

[89] Luttrell S P, 1994, The use of mixture distributions to model the density of speckled data, DRA technical report, DRA/CIS/CBC3/TR94002

[90] Luttrell S P, 1993, The application of mixture distributions to the detection of linear anomalies in cluttered images: a preliminary survey, DRA technical report, 4787

[91] Luttrell S P, 1993, The Markov chain theory of vector quantisers, DRA technical report, 4742

[92] Luttrell S P, 1992, Partitioned mixture distributions: an introduction, DRA technical report, 4671

[93] Luttrell S P, 1992, Code vector density in topographic mappings, DRA technical report, 4669

[94] Luttrell S P, 1991, Self-supervision in multilayer adaptive networks, DRA technical report, 4467

[95] Luttrell S P, 1990, A trainable texture anomaly detector using the Adaptive Cluster Expansion (ACE) method, RSRE technical report, 4437

[96] Luttrell S P, 1990, Asymptotic code vector density in topographic vector quantisers, RSRE technical report, 4392

[97] Luttrell S P, 1989, A Bayesian derivation of an iterative autofocus/super-resolution algorithm, RSRE technical report, 4337

[98] Luttrell S P, 1987, The complex point spread function of the RSRE SAR, RSRE technical report, 4079

[99] Luttrell S P and Pritchard J A S, 1986, A compact digital communications system part 1: rapid carrier acquisition, RSRE technical report, 3925

[100] Luttrell S P, 1985, The implications of Boltzmann-type machines for SAR data processing: a preliminary survey, RSRE technical report, 3815

[101] Luttrell S P, 1985, A super-resolution model for synthetic aperture radar, RSRE technical report, 3785

[102] Luttrell S P and Oliver C J, 1984, Resolution enhancement in coherent images by the introduction of prior knowledge, RSRE technical report, 3697

RSRE Research Notes

[103] Luttrell S P, 1992, The detection of linear anomalies in images, DRA technical report, CSE1/232

[104] Luttrell S P, 1990, A maximum entropy interpretation of random access memory network computations, RSRE technical report, SP4/137

[105] Luttrell S P, 1990, Error measures in adaptive networks, RSRE technical report, SP4/111

[106] Luttrell S P, 1990, Vector quantisation of K-distributed data, RSRE technical report, SP4/110

[107] Luttrell S P, 1990, Bayesian inference on a tree, RSRE technical report, SP4/109

[108] Luttrell S P, 1989, An EM approach to iterative Bayesian super-resolution with applications to K-distributed data, RSRE technical report, SP4/102

[109] Luttrell S P, 1989, The Gibbs Machine applied to hidden Markov model problems. Part 1: Basic theory, RSRE technical report, SP4/99

[110] Luttrell S P, 1987, The relationship between super-resolution and phase imaging of SAR data, RSRE technical report, BS1/41

[111] Luttrell S P, 1987, A proposal for a new method of calibrating radar sensitivity, RSRE technical report, BS1/37

[112] Luttrell S P, 2005, Discrete network dynamics. Part 1: operator theory, arXiv:cs.NE/0511027 arXiv

[113] Luttrell S P, 2004, Using self-organising mappings to learn the structure of data manifolds, arXiv:cs.NE/0406017

[114] Luttrell S P, 2000, Using stochastic encoders to discover structure in data, arXiv:cs.NE/0408049

[115] Luttrell S P, 2000, Invariant stochastic encoders, arXiv:cs.NE/0408050

[116] Luttrell S P, 2000, Stochastic vector quantisers, arXiv:1012.3705

[117] Luttrell S P, 1999, Self-organising stochastic encoders, arXiv:1012.4126

[118] Luttrell S P, 1998, Self-organised factorial encoding of a toroidal manifold, arXiv:cs.LG/0410036

[119] Luttrell S P, 1998, Modelling the probability density of Markov sources, arXiv:cs.NE/0607019

[120] Luttrell S P, 1996, The development of dominance stripes and orientation maps in a self-organising visual cortex network (VICON), arXiv:1012.3724

[121] Luttrell S P, 1995, A self-organising neural network for processing data from multiple sensors, arXiv:1012.4173

[122] Luttrell S P, 1991, Adaptive cluster expansion (ACE): a hierarchical Bayesian network, arXiv:cs.NE/0410020

[123] Luttrell S P, 1991, Adaptive Cluster Expansion (ACE): a multilayer network for estimating probability density functions, arXiv:1012.3656

[124] Luttrell S P, 1994, Anomaly detector, UK Patent GB 2253697 B Patents

[125] Luttrell S P and Pritchard J A S, 1990, Signal transforming device, UK Patent GB 2190221 B

[126] Luttrell S P and Oliver C J, 1987, Super resolution imaging system, UK Patent GB 2173663 B

[127] Luttrell S P, 1982, Power law corrections to hadronic structure functions, Cambridge University PhD dissertation PhD