Book Reviews

Faithful Representations and Topographic Maps by Marc M. van Hulle (New York: Wiley, 2000, 258 pp, hard cover ISBN 0-471-34 507-5) *Reviewed by Wlodzislaw Duch and Antoine Naud.*

This is probably the most significant book on topographic representations and formation of topographic maps written so far. Kohonen's self organizing map (SOM) is treated with due respect here but it covers much more than that. Review of many other topographic map models, their properties and neurobiological roots is followed by a comprehensive treatment of the equiprobabilistic topographic maps introduced by the author. They are based on a solid information-theoretic principle requiring all neurons to contribute to the final representation with the same probability. Such equiprobabilistic maps, leading to a more faithful representation of the input data, are interesting not only from the theoretical point of view. Van Hulle has shown that these maps have many desirable properties that other self-organizing maps do not posses. A number of learning rules for equiprobabilistic map formation has been introduced. For example, the maximum entropy learning rule (MER) leads to the self-organization by a local weight adaptation without the need for neighborhood functions commonly used in SOM maps. Neurobiological plausibility of such algorithms is also discussed in the book.

In the first chapter, the reader is briefly introduced to topographic maps resulting from the self-organizing processes in the sensory cortices of the brain. A series of topographic map models (purely theoretical or biologically inspired) is presented, starting from Willshaw, von der Malsburg, and Amari models and finishing with the genrative topographic map model of Bishop. This "panorama" of various approaches is quite broad but not exhaustive (for example, the neuroscale model of Lowe is not mentioned). In the second chapter, many properties of the SOM model of Kohonen are presented: topographic ordering, convergence, neighborhood functions, relation to vector quantization, clustering, density estimation, and extensions of the SOM model. Some applications to statistical data analysis, density estimation and regression are presented in the third chapter. Neurons of the SOM map are not active with equal probabilities due to the nonlinear relationship between the weight density at convergence and the input data density. They oversample low probability region and undersample high probability regions of the input space. Some unsupervised competitive learning algorithms are not able to effectively use all available neurons, leaving neurons that are never active.

In Chapter 4, the equiprobabilistic topographic map model is introduced from an information-theoretic point of view. It should preserve a maximum amount of information available about the input data distribution. Approximately equiprobabilistic maps may also be obtained using heuristic modifications of the SOM learning rule or the learning rule of an unsupervised competitive learning algorithm. For example, the algorithms may modify the metric used to compare input vectors with weights of neurons aiming at equal distortion produced by each neuron. Several such distortion-based algorithms are discussed, including some constructive algorithms. Such algorithms are rather improbable from the biological perspective. The maximum information preservation principle (infomax) introduced by Linsker (1988) allows to create equiprobable topographic maps, although the learning rule that results from this principle is rather complex. Learning rules developed by van Hulle for equiprobabilistic topographic map formation are based on maximization of the information-theoretic entropy of the map's outputs. Several variants of lattice-based and kernel-based learning rules are discussed. The first type of maps assumes that the lattice dimensionality is the same as the input space while the second type reduces the dimensionality.

Applications of the kernel-based topographic map learning for nonparametric density estimation and density-based clustering are given in Chapter 5. Some remarks on the use of this approach for Bayesian classification are also made, although no results are given. An interesting application of the kernel MER model to the blind source separation problem is presented. This is followed by applications to speech and image encoding in which specific features (for example Gabor filters) are discovered thanks to the adaptation of kernel subspaces instead of kernel functions. Results are compared with the adaptive subspace SOM and applied to the compression of speech and music signals. The most ambitious application of kernel-based MER approach presented in the last chapter concerns an unsupervised identification of musical instruments and the notes they play using spectrograms (512 components). This is not an easy task but the results are quite impressive.

The book is clearly written and can be recommended to anyone interested in the unsupervised learning, topographical maps and the process of self-organization, both from theoretical and biological perspective. Its structure reflects well the origins of the presented model and its connections with alternative models. Clear presentation of algorithms in pseudocode should be of great help to the readers looking for practical information without the need to go through detailed theoretic considerations. The performance of different algorithms is compared in details and their computational details well presented. On the other hand there are many issues related to self-organizing maps that are not included in this book. The use of topographic maps for visualization of multidimensional data is not covered and several constructive SOM algorithms and interpolation methods useful for regression are omitted. This should not be taken as a criticism of the book but simply as a reminder that its main focus is on formation of topographic maps and unsupervised learning.

Statistics and Neural Networks: Advances at the Interface by J. W. Kay and D.M. Titterington, Eds. Oxford, U.K.: Oxford Univ. Press, 1999, 278 pp., 40 GBP, ISBN 0–19–852422–6. *Reviewed by Wlodzislaw Duch.*

Gone are the old days of ad hoc neural methods, when biological inspiration was sufficient to justify the model. Applying neural methods to classification or approximation problems we are on the grounds claimed by statisticians. For statisticians, feedforward neural models are either nonlinear classifiers performing discriminating/clusterization, or regression models performing approximation. In recent years application of statistical methods has brought deeper insight into methods (boosting, bagging, stacking, Monte Carlo, and Gibbs sampling), regularization, criteria for optimal model choice, mixture models, and many other concepts.

Statistics and Neural Networks has eight chapters written by the top experts in statistics and neural networks. These chapters have between 20–60 pages and although they contain original results they are written in a clear way, introducing their subject in detail. Except for typical neural subjects, such as the radial basis function networks and statistics (by D. Lowe), and typical statistical subjects, such as discriminant

The reviewers are with the Department of Computer Methods, Nicholas Copernicus University, Torun, Poland (e-mail: duch@ieee.org).

Publisher Item Identifier S 1045-9227(00)10092-X.

models (by T. Hasti *et al.*)., they include such topics as information theory, density networks, latent variable models, and visualization of multidimensional data (C. Bishop). I can sincerely recommend this book to every neural researcher–there is a lot to learn here.

Review of ReVision: Reviewing Vision in the Web http://www.ime.usp.br/~cesar/revision/. Reviewed by H. John Caulfield, Fisk University, Nashville, TN (e-mail: hjc@dubois.fisk.edu).

This is a comprehensive and organized set of websites and commentary that make a good place for both newcomers and experienced visual perception workers to begin the search for needed information on visual perception. I sometimes make the same remark about "The Joy of Visual Perception: A Web book" (http://www.yorku.ca/eve/). "Joy" is a self-contained web book with journal and book references and only a couple of dozen links to the web. *Revision* is a set of many organized and rated web sites. Thus one is a book written on the web and the other is a guided tour of the web from a vision viewpoint.

The entry page to ReVision has wonderful graphics of eyes but leaves something to be desired in terms of both helpfulness (You are not told what you get when you click on any eye) and speed (the cute graphics—color of irises in the well-drawn graphics is continuously changing—is slow loading). In fact, speed is a problem throughout (the site is in Brazil), but this should not detract from the overall value of what is present at this site.

Here are a few topics and the number of links for each to give a feel for how comprehensive the survey is:

3D 38, Active models 5, Active vision 5, Applications 12, Artificial Intelligence 30, Attention 9, Biological vision 22

That list is a small but representative sample of the topics and coverage of each.

I believe this to be a significant contribution to vision research. The compilers seem anxious to improve it and keep it updated with user support. This can be a useful tool for the vision community.

Review of Electronic Journal: Neural Computing Surveys by A. Jagota, UCSC and ICSI-Berkeley, CA.

Neural Computing Surveys (NCS) (http://www.icisi.berkeley.edu/~jagota/NCS) is a recently introduced electronic journal that publishes survey articles on all aspects of neural computation, from theory to practice, from engineering to neuroscience, from core to interdisciplinary aspects. NCS articles are available in full-text form free on the world wide web. Volume 1 was also published in print form by Lawrence Earlbaum Associates.

Volumes 1 and 2 were published in 1998 and 1999, respectively, and Volume 3 is being published incrementally in the year 2000.

Volume 1 began with a survey by M. Anthony on a PAC model and its applications to the analysis of learning in neural networks. Next, M. Misra wrote about parallel environments for implementing neural networks. This was followed by D. A. Medler who provided a brief history of connectionism. Finally, S. Kaski, J. Kangas, and T. Kohonen rounded out the volume with an extensive bibiliography of papers on self-organizing maps, from 1981 through 1997, with more than 3000 references, a road map, and multiple indexes. All these surveys were reviewed by experts in their fields, and all required at least one round of revision. Each had more than 100 references.

Volume 2 began with an innovation, the notion of a collectice article. An open call was made for contributions on the topic of "connectionist symbol processing: dead or alive?" Received contributions were informally reviewed. A handful were rejected, but most were put together, upon revision, into an article of the same name, edited by A. Jagota, T. Plate, L. Shastri, and R. Sun. Some positive feedback has been received on the contents and on the form of this article. Next, T. B. Ludermir *et al.* published a survey on weightless neural models. This was followed by a survey by K. MCGarry *et al.* on hybrid neural systems. Fourth, A. Hyvarinen published a nice survey on the hot topic of independent component analysis. This was followed by a survey by Y. Bengio on the even hotter topic of the use of hidden Markov models in the analysis of sequential data. Finally, W. Duch *et al.* rounded out the volume with a survey on neural transfer functions.

Neural Networks Papers on the Internet by W. Duch.

It is no secret that the new generation of students and young researchers tries to avoid going to the library by searching for everything on the Internet. For them information that cannot be found in the Internet does not exist. Last year I received an email from an American high school student asking "who was this Copernicus guy? I have to write a paper and can't find anything."

For the new generation Internet becomes not just the source of information, but the only source of information. Young researchers first check the Internet pages of journals to find relevant papers and only then—if they are lucky and have a good library subscribing to the particular journal—go to the library. Unfortunately good libraries are hard to find even in the United States and Western Europe, not to mention Central Europe and developing countries. Quite a few people (myself included) moved from engineering, physics, or mathematics to neural computing. The fact that they could collect many useful papers downloading them from Internet archives greatly facilitated this movement. Yet many "serious" scientists pay little attention to this new media, spending a significant amount of time on preparation of their articles and conference presentations and devoting virtually no time to maintain archives with their papers and web pages devoted to their projects. There are various prizes for the best papers but no prizes for the most useful web sites, hence little motivation to spend any time on developing and maintaing the sites.

Some conferences publish thousands of pages in thick volumes—who can read it all? It is faster to find the relevant papers on the Internet and print only those one wants. Neural-network societies (INNS, ENNS, APNNA, IEEE NNC), journals and large conferences still continue their business as usual: the web addresses of participants are almost never published, frequently even email addresses are not published, no information about useful web sites is ever mentioned, conference homepages rarely keep a table on contents of their proceedings. There are some signs of change, at least in respect to IEEE conferences: a great new service, called IEEE Xplore (it has been mentioned in the editorial article of the IEEE TNN in January 2000) allows everyone to access tables of contents of IEEE transactions, journals, magazines, conference proceedings, and standards. IEEE members also have full access to articles and other materials in the PDF format. These publications should be cross-linked using hyperlinks. This is a welcome extension to the IEEE bibiliographies online and to the Opera services. The IEEE Xplore may be found at http://ieeexplore.ieee.org.

Although many experts maintain their local archives without a well-organized, central repository of cataloged papers they are not only hard to find, but also they do not have strong motivation to place their papers in a local archive. Without the support of the neural-network community many useful initiatives to collect papers run out of steam and die. For example, the *Neuroprose* archive of papers, which started in 1989 and accumulated more than 600 papers and 53 Ph.D. dissertations, has never been properly organized. The papers were automatically dropped to a huge catalog that finally grew to a size that prevented finding anything interesting unless one knew the name of the file. The readme file was modified for the last time in 1994 and the archive itself seems to have died in 1998. *Neuroprose* was based on the ftp protocol which is out of date and frequently does not work properly due to various firewalls set up for security reasons. Uploading the files to a big catalog is certainly not a good long-term solution. There is no reason why the archives should not be accessible by the http protocol through web

pages. The London and Southeast Center for High Performance Computing (SEL-HPC) has opened an archive storing papers related to parallel and functional programming, vision and image processing, computational mathematics, neural networks, human computer interaction, and other subjects. The archive accumulated more than 700 papers and stored also the links to the home pages of authors that kept their papers there. Unfortunately, in 1998 the SEL-HPC center was shut down and although the archive is still there it does not accept new papers, given the"forbidden access" message. There was no sign on the archive pages that this is the case, but this should have been changed by now.

The internet is full of old files that have not been removed and that show up every time you do a search. Ask for "neural archives" and you will get either a missing link or a 1993 list of archives which have not existed in years. A few people are responsible enough to remove the outdated files and to write the date of last modification on their page. Below, I have made a short review of the current situation—an in-depth review of some of these and of other useful projects would be welcomed.

Los Alamos e-Print Archive (http://arxiv.org/) has continuously served the physics and mathematics community since August 1991. The number of hosts connecting to this archive reaches 9000 per day as of April 2000, excluding numerous mirror sites, and the number of new papers stored here exceeds 2500 per month. The current estimate of the total number of papers in this archive is about 130 000. It contains several subarchives of interest to the neural-network community.

- The Computing Research Repository (CoRR) opened in September 1998 and is sponsored by ACM. Papers in CoRR are classified in two ways: subject area (each subject has a moderator) and by using the ACM classification scheme which covers all of computer science. Subjects relevant to computational intelligence include artificial intelligence, computation and language, computer vision and pattern recognition, human–computer interaction, learning, multiagent systems, neural and evolutionary computation, and robotics. The address is http://arxiv.org/archive/cs.intro.html.
- Nonlinear Sciences repository includes adaptation and self-organizing systems, cellular automata, and chaotic dynamics.
- Physics repository includes disordered systems and neural networks, data analysis, statistics, probability, and Bayesian analysis.
- Other interesting repositories in this archive, like cellular/molecular neuroscience, developmental and behavioral/systems neuroscience, are not too popular yet, but are ready to receive more papers.

The e-Print archive has many copies around the world and is a great service to the physics and mathematics community. A similar initiative has been started in cognitive sciences. The CogPrints archive (http://cogprints.soton.ac.uk/) stores papers in psychology, neuroscience, linguistics, artificial intelligence, robotics, vision, learning, speech, neural network, philosophy of mind and language, behavioral ecology, sociobiology, behavior genetics, evolutionary theory, psychiatry, neurology, human genetics, brain imaging, anthrpology, and other social and mathematical sciences pertinent to the study of cognition. CogPrints archive received the Psychological Science Award for "contribution to psychology on the internet" from PsychologicalScience.net.

Although no complete repository designed specifically for the neural-network community seems to exist several initiatives are worth mentioning. They aim at indexing the Web resources, searching for papers placed in other people's archives. Storing the links to papers at different sites has some disadvantages: links are sometimes changes or papers removed by system administrators, and it is hard to index them in a useful way. On the other hand, it is much easier to index the Web than to create a central repository.

The Princeton NEC Research Institute team (Lee Giles, Steve Lawrence, and Kurt Bollacker) created the CiteSeer, now called a research Index (http://www.researchindex.com), a system for automatically creating digital libraries, with emphasis on citation matching, indexing, and ranking. The links to the postscript or pdf papers are shown as a result of search and cached versions of articles are available directly from the system. The "correct" option will retrieve the first page of the paper and show it in the browser. Since the last and first names may not be correctly ordered the authors should check their papers manually and correct the database, especially if their names contains letters with accents—the system does not deal correctly with French, Spanish, or Polish names! On demand citations are shown in context in which they appeared. Research Index seems to be quite popular since the system reports high load even in the middle of the night (Princeton time). The "Computer Science Directory" shows a list of papers in different subject areas, ordered according to the number of citations, authority of their authors, and their tutorial value.

Research Index is certainly the best bibilographical system at the moment. It does not give references to authors' home pages (it would be rather hard to make it in an automatic way since the authors do not give their web pages in publications) but one can frequently guess them starting from the link to the paper. The HP-search service (Trier University, Germany, http://hpsearch.uni-trier.de/hp/) is the best chance to find personal home pages of computer scientists. In April 2000 there were more than 42 000 entries stored there.

The Collection of Computer Science Bibilographies has more than 1 000 000 references, amounting to 660 Mb of BibTex entries. It is a meta-service, composed of about 1200 specialized bibilographies and updated monthly from their original locations. About 90 000 references URLs to online versions of papers. There are more than 2000 links to other sites carrying bibliographic information, including the large Computer Science Bibliography at Trier University, that stores the bibliographical information from major conferences, books, and journals, but contains relatively little information related to neural networks. The address is http://liinwww.ira.uka.bibliography/index.html.

The ML papers search engine (http://gubbio.cs.berkeley.edu/mlpapers/) created in 1997 to index the machine learning papers was probably the first system implementing automatic extraction of titles, authors, abstracts, and links to postscript papers (this seems to be the only format indexed). Almost 1300 papers with the keyword "neural" were found in April 2000. Adding new papers requires only giving a URL to ftp or http archive.

An interesting approach to the problem of automatic indexing has been taken by the Just Research Company, creators of Cora (http://www.cora.justresearch.com/), a special-purpose search engine covering more than 50 000 research papers found in about 100 computer science departments. Computer science has been divided into ten categories, with "artificial intelligence" subjects involving data mining and machine learning subjects, and in the "machine learning" section neural networks are found among case-based learning, genetic algorithms, probabilistic methods, reinforcement learning, rule learning, and theory. There are 75 computer science categories in this database. Cora allows searching for keywords found in papers that are stored in the postscript format. Results are used to provide automatically generated BibTex entries (sometime with errors) and include titles, authors, abstracts of the paper, main page address where the paper has been found, alternative address where it is stored, a list of references extracted from the paper, and backward references (papers referring to the current paper).

The process of adding new files is very simple, requiring just a submission of the URL of the archive. The 50 highest ranking papers

are automatically displayed in each category, with rank determined by analysis of citations allowing for an automatic identification of papers as survey articles, seminal articles, papers with the largest number of references, and papers written by authoritative authors. Reinforcement learning and probabilistic techniques have been used for automatic classification. The Cora project is led by Andrew McCallum and although it is still in the research phase it is quite useful.

The New Zealand Digital Library (NZDL) project (http://www.nzdl.org/) has created publicly available search engines for domains from computer science technical reports to music videos. The emphasis of this project is on the creation of full-text searchable digital libraries; it does not use machine learning technology to automate the creation of search engines. The NZDL collection provides a huge index to more than 45 000 computer science reports—whole papers were analyzed (over 30 Gb of postscript files, 1 300 000 pages), so any fragment of text may be found. The collection also contains almost 30 000 figures extracted from the reports.

The Networked Computer Science Technical Reference Library (NCSTRL) aims at creation of "premier international online library of computer science technical reports." It provides software allowing institutions (over 160 at the moment, including leading Unites States universities), to create digital libraries which are then integrated into the NCSTRL services. A simpler solution is also possible, with ftp archive servers, requiring contribution of bibliographical information to the central NCSTRL index (http://cs-tr.cs.cornell.edu/). Since the indexing is not automatic the project will probably stay more popular with the libraries (it is run by Cornell University Digital Library Research group) than with computer scientists.

Experimental mthods of searching for the relevant information based on the self-organized mappings exists (for example, WebSOM, http://websom.hut.fi/sebsom/) but so far they are not useful for large-scale searches.

Finally, The IEEE Bibliographies Online (for IEEE members only) provides bibliographical information in IEEE-sponsored conferences and journals, but no links to the actual paper (this is available to subscribers via the Opera Services). IEEE bibliographies are at http://www.biblio.ieee.org/scripts/biblio_home.html.

Links to all services mentioned here (and to many others) are stored at http://www.phys.uni.torun.pl/~duch/neural.html#biblio.