Miguel Á. CARREIRA-PERPIÑÁN

Professor, Dept. of Computer Science and Engineering, UC Merced

RESEARCH INTERESTS

My basic research interests are in machine learning. In recent years, I have been working on topics in the intersection of optimisation and machine learning, in particular in learning algorithms for deep neural nets and other "nested" systems, for decision trees and tree-based models, and for nonlinear embeddings; and algorithms for optimally compressing deep neural nets. Other topics of interest are dimensionality reduction/manifold learning, clustering, denoising and other unsupervised learning problems, and mean-shift algorithms. I often get inspiration from problems in speech processing (e.g. articulatory inversion and model adaptation), computer vision (e.g. segmentation, articulated pose tracking, image registration), sensor networks, robotics (e.g. inverse kinematics) and other application areas. In the past, I also worked on computational neuroscience, specifically on dimension reduction models of the maps of the visual cortex.

PERSONAL DATA

Name: Miguel Á. Carreira-Perpiñán
Address: Dept. of Computer Science and Engineering
School of Engineering
University of California, Merced
5200 N. Lake Road, Merced, CA 95343, USA
E-mail: mcarreira-perpinan@ucmerced.edu
WWW: http://faculty.ucmerced.edu/mcarreira-perpinan
ORCID: 0000–0003–3297–9375
GitHub: https://github.com/UCMerced-ML
Google Scholar: http://scholar.google.com/citations?hl=en&user=SYdYhxgAAAAJ
DBLP: https://dblp.org/pid/23/5257
LinkedIn: https://www.linkedin.com/in/miguel-a-carreira-perpinan

EDUCATION

Sep. 1995 to Apr. 2000: ..... PhD in Computer Science.
University of Sheffield, UK.
PhD thesis (2001) on dimensionality reduction, inverse problems and sequential data reconstruction with latent variable models. Advisor: Prof. Steve Renals.

1986 to 1991: ..... BSc, MSc in Computer Science.
Technical University of Madrid, Spain (6-year degree: *Licenciado en informática*).
MSc thesis (1995) on neural networks for ear and face identification. Advisor: Dr. Ángel Sánchez.

ACADEMIC APPOINTMENTS

Since Jul. 2015: ..... Professor.
Dept. of Computer Science and Engineering, School of Engineering, University of California, Merced.

Jul. 2010 to Jun. 2015: ..... Associate professor with tenure.
Electrical Engineering & Computer Science, School of Engineering, University of California, Merced.

Jul. 2007 to Jun. 2010: ..... Assistant professor.
Electrical Engineering & Computer Science, School of Engineering, University of California, Merced.

Oct. 2004 to Jun. 2007: ..... Assistant professor.
Dept. Computer Science & Electrical Engineering, OGI School of Science & Eng., Oregon Health & Science University.

Jan. 2003 to Aug. 2004: ..... Postdoctoral fellow.
Dept. Computer Science, University of Toronto, Canada.
Spectral clustering, manifold learning, image labeling, Markov-chain Monte Carlo methods for random fields. Advisors: Dr. Richard S. Zemel and Prof. Geoffrey E. Hinton.

May. 2000 to Dec. 2002: ..... Postdoctoral fellow.
Dept. Neuroscience, Georgetown University, Washington, DC.
Modeling visual cortical map development with generalized elastic nets. Advisor: Dr. Geoffrey J. Goodhill.

VISITING ACADEMIC APPOINTMENTS

Jan. 15, 2018 to Jun. 30, 2018: ..... Visiting professor (sabbatical).
Dept. Statistics, University of Washington.

Aug. 25, 2013 to Dec. 25, 2013: ..... Visiting professor (sabbatical).
Robotics Institute, Carnegie-Mellon University.

Aug. 8, 2010 to Aug. 23, 2010: ..... Visiting professor.
Centre for Speech Technology Research, School of Informatics, University of Edinburgh.

Aug. 18, 2006 to Sep. 3, 2006: ..... Visiting professor.
Centre for Speech Technology Research, School of Informatics, University of Edinburgh.

PROFESSIONAL EXPERIENCE

Jan. 1993 to Dec. 1994: ..... Member of technical staff.
European Space Operations Centre, European Space Agency (ESOC/ESA), Darmstadt, Germany.
Real-time simulation of satellite thermal subsystems. Advisor: Mr. Juan E. Miró.

Mar. 1992 to Dec. 1992: ..... Analyst programmer.
LAB 2000, Madrid.

Jan. 1992 to Sep. 1992: ..... Technical officer (rank: lieutenant).
Ministry of Defence, Madrid.

Oct. 1990 to May 1991: ..... Intern.
Telefónica, Madrid.

Jan. 1990 to Oct. 1990: ..... Intern.
IBM, Madrid.

Jan. 1985 to Jan. 1986: ..... Instructor.
Instituto Superior de Estudios (ISE), Madrid.

PUBLICATIONS LIST

Preprints, software and data are available from my web page or from Google Scholar. Supervised students appear underlined.

▷ **Conference proceedings (refereed)**

M. Klemen, M. Á. Carreira-Perpiñán and P. López-García (2023): Solving recurrence relations using machine learning, with application to cost analysis. *39th Int. Conf. Logic Programming (ICLP 2023)*, to appear.

M. Á. Carreira-Perpiñán, M. Gabidolla and A. Zharmagambetov (2023): Towards better decision forests: Forest Alternating Optimization. *IEEE Computer Society Conf. Computer Vision and Pattern Recognition (CVPR 2023)*, to appear. Acceptance rate: 25.8% (2359/9155).

M. Á. Carreira-Perpiñán and S. S. Hada (2023): Very fast, approximate counterfactual explanations for decision forests. *37th AAAI Conf. Artificial Intelligence (AAAI 2023)*, pp. 6935–6943. Acceptance rate: 19.6% (1721/8777).

A. Zharmagambetov and M. Á. Carreira-Perpiñán (2022): Semi-supervised learning with decision trees: graph Laplacian tree alternating optimization. *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*, pp. 2392–2405. Acceptance rate: 25.6% (???/10411). ▶

M. Gabidolla and M. Á. Carreira-Perpiñán (2022): Optimal interpretable clustering using oblique decision trees. *ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining (KDD 2022)*, pp. 400–410. Acceptance rate: 15.0% (254/1695).

M. Gabidolla, A. Zharmagambetov and M. Á. Carreira-Perpiñán (2022): Improved multiclass AdaBoost using sparse oblique decision trees. *Int. Joint Conf. Neural Networks (IJCNN 2022)*, 8 pages. Acceptance rate: 55.4% (1467/2646).

S. S. Hada and M. Á. Carreira-Perpiñán (2022): Sparse oblique decision trees: a tool to interpret natural language processing datasets. *Int. Joint Conf. Neural Networks (IJCNN 2022)*, 8 pages. Acceptance rate: 55.4% (1467/2646).

M. Gabidolla and M. Á. Carreira-Perpiñán (2022): Pushing the envelope of gradient boosting forests via globally-optimized oblique trees. *IEEE Computer Society Conf. Computer Vision and Pattern Recognition (CVPR 2022)*, pp. 285–294. Acceptance rate: 25.3% (2067/8161).

Y. Idelbayev and M. Á. Carreira-Perpiñán (2022): Exploring the effect of ℓ_0/ℓ_2 regularization in neural network pruning using the LC toolkit. *IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP 2022)*, pp. 3373–3377. Acceptance rate: 45% (1721/3967).

S. S. Hada and M. Á. Carreira-Perpiñán (2022): Interpretable image classification using sparse oblique decision trees. *IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP 2022)*, pp. 2759–2763. Acceptance rate: 45% (1721/3967).

A. Zharmagambetov and M. Á. Carreira-Perpiñán (2022): Learning interpretable, tree-based projection mappings for nonlinear embeddings. *25th Int. Conf. Artificial Intelligence and Statistics (AISTATS 2022)*, pp. 9550–9570. Acceptance rate: 29.2% (492/1685). ▶

A. Zharmagambetov, M. Gabidolla and M. Á. Carreira-Perpiñán (2021): Softmax Tree: an accurate, fast classifier when the number of classes is large. *Conf. Empirical Methods in Natural Language Processing (EMNLP 2021)*, pp. 10730–10745. Acceptance rate: 22.6% (840/3717).

Y. Idelbayev and M. Á. Carreira-Perpiñán (2021): LC: A flexible, extensible open-source toolkit for model compression. *30th ACM Int. Conf. Information & Knowledge Management (CIKM 2021)*, resource paper, pp. 4504–4514. Acceptance rate: 32.5% (26/80).

Y. Idelbayev and M. Á. Carreira-Perpiñán (2021): More general and effective model compression via an additive combination of compressions. *European Conf. Machine Learning and Principles and Practice of Knowledge Discovery (ECML/PKDD 2021)*, pp. 233–248. Acceptance rate: 21.5% (147/685).

A. Zharmagambetov, M. Gabidolla and M. Á. Carreira-Perpiñán (2021): Improved multiclass AdaBoost for image classification: the role of tree optimization. *IEEE Int. Conf. Image Processing (ICIP 2021)*, pp. 424–428. Acceptance rate: 45.9% (791/1722).

A. Zharmagambetov and M. Á. Carreira-Perpiñán (2021): A simple, effective way to improve neural net classification: ensembling unit activations with a sparse oblique decision tree. *IEEE Int. Conf. Image Processing (ICIP 2021)*, pp. 369–373. Acceptance rate: 45.9% (791/1722).

Y. Idelbayev and M. Á. Carreira-Perpiñán (2021): Beyond flops in low-rank compression of neural networks: optimizing device-specific inference runtime. *IEEE Int. Conf. Image Processing (ICIP 2021)*, pp. 2843–2847. Acceptance rate: 45.9% (791/1722).

S. S. Hada and M. Á. Carreira-Perpiñán (2021): Sampling the "inverse set" of a neuron. *IEEE Int. Conf. Image Processing (ICIP 2021)*, pp. 3712–3716. Acceptance rate: 45.9% (791/1722).

S. S. Hada, M. Á. Carreira-Perpiñán and A. Zharmagambetov (2021): Understanding and manipulating neural net features using sparse oblique classification trees. *IEEE Int. Conf. Image Processing (ICIP 2021)*, pp. 3707–3711. Acceptance rate: 45.9% (791/1722).

A. Zharmagambetov, S. S. Hada, M. Gabidolla and M. Á. Carreira-Perpiñán (2021): Non-greedy algorithms for decision tree optimization: an experimental comparison. *Int. Joint Conf. Neural Networks (IJCNN 2021)*, 8 pages. Acceptance rate: 59.3% (1183/2032).

A. Zharmagambetov, M. Gabidolla and M. Á. Carreira-Perpiñán (2021): Improved boosted regression forests through non-greedy tree optimization. *Int. Joint Conf. Neural Networks (IJCNN 2021)*, 8 pages. Acceptance rate: 59.3% (1183/2032).

Y. Idelbayev and M. Á. Carreira-Perpiñán (2021): An empirical comparison of quantization, pruning and low-rank neural network compression using the LC toolkit. *Int. Joint Conf. Neural Networks (IJCNN 2021)*, 8 pages. Acceptance rate: 59.3% (1183/2032).

Y. Idelbayev, P. Molchanov, M. Shen, H. Yin, M. Á. Carreira-Perpiñán and Jose M. Alvarez (2021): Optimal quantization using scaled codebook. *IEEE Computer Society Conf. Computer Vision and Pattern Recognition (CVPR 2021)*, pp. 12090–12099. Acceptance rate: 23.6% (1661/7093).

Y. Idelbayev and M. Á. Carreira-Perpiñán (2021): Optimal selection of matrix shape and decomposition scheme for neural network compression. *IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP 2021)*, pp. 3250–3254. Acceptance rate: 48% (1734/3610).

A. Zharmagambetov and M. Á. Carreira-Perpiñán (2021): Learning a tree of neural nets. *IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP 2021)*, pp. 3140–3144. Acceptance rate: 48% (1734/3610).

Y. Idelbayev and M. Á. Carreira-Perpiñán (2021): Neural network compression via additive combination of reshaped, low-rank matrices. *Data Compression Conference (DCC 2021)*, pp. 243–252.

M. Á. Carreira-Perpiñán and S. S. Hada (2021): Counterfactual explanations for oblique decision trees: exact, efficient algorithms.

35th AAAI Conf. Artificial Intelligence (AAAI 2021), pp. 6903–6911. Acceptance rate: 21.4% (1692/7911).

S. S. Hada and M. Á. Carreira-Perpiñán (2021): Style transfer by rigid alignment in neural net feature space. *IEEE Winter Conf. Applications of Computer Vision (WACV 2021)*, pp. 2576–2585.

M. Á. Carreira-Perpiñán and A. Zharmagambetov (2020): Ensembles of bagged TAO trees consistently improve over Random Forests, AdaBoost and Gradient Boosting. *ACM-IMS Foundations of Data Science Conf. (FODS 2020)*, pp. 35–46. Acceptance rate: 29.3% (17/58).

A. Zharmagambetov and M. Á. Carreira-Perpiñán (2020): Smaller, more accurate regression forests using tree alternating optimization. *Int. Conf. Machine Learning (ICML 2020)*, pp. 11398–11408. Acceptance rate: 21.8% (1088/4990). ▶

Y. Idelbayev and M. Á. Carreira-Perpiñán (2020): Low-rank compression of neural nets: learning the rank of each layer. *IEEE Computer Society Conf. Computer Vision and Pattern Recognition (CVPR 2020)*, pp. 8046–8056. Acceptance rate: 22% (1470/6656).

E. Eban, Y. Movshovitz-Attias, H. Wu, A. Poon, M. Sandler, Y. Idelbayev and M. Á. Carreira-Perpiñán (2020): Structured multi-hashing for model compression. *IEEE Computer Society Conf. Computer Vision and Pattern Recognition (CVPR 2020)*, pp. 11900–11909. Acceptance rate: 22% (1470/6656).

M. Á. Carreira-Perpiñán and M. Alizadeh (2019): ParMAC: distributed optimisation of nested functions, with application to learning binary autoencoders. *Proc. Machine Learning and Systems (MLSys 2019)*, pp. 276–288. Acceptance rate: 16.9% (32/189). ▶

M. Á. Carreira-Perpiñán and P. Tavallali (2018): Alternating optimization of decision trees, with application to learning sparse oblique trees. *Advances in Neural Information Processing Systems 31 (NeurIPS 2018)*, pp. 1217–1227. Acceptance rate: 20.8% (1011/4856).

M. Á. Carreira-Perpiñán and Y. Idelbayev (2018): "Learning-compression" algorithms for neural net pruning. *IEEE Computer Society Conf. Computer Vision and Pattern Recognition (CVPR 2018)*, pp. 8532–8541. Acceptance rate: 29% (979/3300+).

D. A. Winkler, M. Á. Carreira-Perpiñán and A. E. Cerpa (2018): Plug-and-play irrigation control at scale. *17th ACM/IEEE Int. Conf. Information Processing in Sensor Networks (IPSN 2018)*, pp. 1–12. Acceptance rate: 26.8% (22/82).

R. Raziperchikolaei and M. Á. Carreira-Perpiñán (2017): Learning supervised binary hashing: optimization vs diversity. *IEEE Int. Conf. Image Processing (ICIP 2017)*, pp. 3695–3699. Acceptance rate: 45.3% (915/2021).

R. Raziperchikolaei and M. Á. Carreira-Perpiñán (2017): Learning circulant support vector machines for fast image search. *IEEE Int. Conf. Image Processing (ICIP 2017)*, pp. 385–389. Acceptance rate: 45.3% (915/2021).

M. Vladymyrov and M. Á. Carreira-Perpiñán (2017): Fast, accurate spectral clustering using locally linear landmarks. *Int. Joint Conf. Neural Networks (IJCNN 2017)*, pp. 3870–3879. Acceptance rate: 66.6% (621/933).

R. Raziperchikolaei and M. Á. Carreira-Perpiñán (2016): Learning supervised binary hashing: optimization vs diversity. *IEEE Int. Conf. Data Mining (ICDM 2016)*, short paper, pp. 1173–1178. Acceptance rate: 19.7% (178/904).

M. Á. Carreira-Perpiñán and R. Raziperchikolaei (2016): An ensemble diversity approach to supervised binary hashing. *Advances in Neural Information Processing Systems 29 (NIPS 2016)*, pp. 757–765. Acceptance rate: 22.7% (568/2500).

R. Raziperchikolaei and M. Á. Carreira-Perpiñán (2016): Optimizing affinity-based binary hashing using auxiliary coordinates. *Advances in Neural Information Processing Systems 29 (NIPS 2016)*, pp. 640–648. Acceptance rate: 22.7% (568/2500).

M. Vladymyrov and M. Á. Carreira-Perpiñán (2016): The Variational Nyström method for large-scale spectral problems. *Int. Conf. Machine Learning (ICML 2016)*, pp. 211–220. Acceptance rate: 24.3% (322/1327).

D. A. Winkler, R. Wang, F. Blanchette, M. Á. Carreira-Perpiñán and A. E. Cerpa (2016): MAGIC: Model-Based Actuation for Ground Irrigation Control. *15th ACM/IEEE Int. Conf. Information Processing in Sensor Networks (IPSN 2016)*, article No. 9. Acceptance rate: 19.5% (23/118).

M. Á. Carreira-Perpiñán and M. Vladymyrov (2015): A fast, universal algorithm to learn parametric nonlinear embeddings. *Advances in Neural Information Processing Systems 28 (NIPS 2015)*, pp. 253–261. Acceptance rate: 21.9% (403/1838).

M. Á. Carreira-Perpiñán and R. Raziperchikolaei (2015): Hashing with binary autoencoders. *IEEE Computer Society Conf. Computer Vision and Pattern Recognition (CVPR 2015)*, pp. 557–566. Acceptance rate: 28.4% (602/2123).

M. Á. Carreira-Perpiñán and W. Wang (2014): LASS: a simple assignment model with Laplacian smoothing. *28th AAAI Conf. Artificial Intelligence (AAAI 2014)*, pp. 1715–1721. Acceptance rate: 28% (398/1406).

W. Wang and M. Á. Carreira-Perpiñán (2014): The role of dimensionality reduction in classification. *28th AAAI Conf. Artificial Intelligence (AAAI 2014)*, pp. 2128–2134. Acceptance rate: 28% (398/1406).

M. Á. Carreira-Perpiñán and W. Wang (2014): Distributed optimization of deeply nested systems. *17th Int. Conf. Artificial Intelligence and Statistics (AISTATS 2014)*, pp. 10–19. **Notable paper award**. Acceptance rate: 36% (121/335). ▶

M. Vladymyrov and M. Á. Carreira-Perpiñán (2014): Linear-time training of nonlinear low-dimensional embeddings. *17th Int. Conf. Artificial Intelligence and Statistics (AISTATS 2014)*, pp. 968–977. Acceptance rate: 36% (121/335).

M. Vladymyrov and M. Á. Carreira-Perpiñán (2013): Locally Linear Landmarks for large-scale manifold learning. *European Conf. Machine Learning and Principles and Practice of Knowledge Discovery (ECML/PKDD 2013)*, pp. 256–271. Acceptance rate: 25% (111/443).

M. Vladymyrov and M. Á. Carreira-Perpiñán (2013): Entropic affinities: properties and efficient numerical computation. *Int. Conf. Machine Learning (ICML 2013)*, pp. 477–485. Acceptance rate: 23.5% (283/1204). ▶

A. Kamthe, M. Á. Carreira-Perpiñán and A. E. Cerpa (2013): Quick construction of data-driven models of the short-term behavior of wireless links. *32nd Annual IEEE Int. Conf. Computer Communications (INFOCOM 2013) Mini-Conference*, pp. 160–164. Acceptance rate: $<$ 25% (???/1613).

M. Vladymyrov and M. Á. Carreira-Perpiñán (2012): Partial-Hessian strategies for fast learning of nonlinear embeddings. *Int. Conf. Machine Learning (ICML 2012)*, pp. 345–352. Acceptance rate: 27.2% (242/890). ▶

W. Wang and M. Á. Carreira-Perpiñán (2012): Nonlinear low-dimensional regression using auxiliary coordinates. *15th Int. Conf. Artificial Intelligence and Statistics (AISTATS 2012)*, pp. 1295–1304. Acceptance rate: 39% (157/402).

M. Farhadloo and M. Á. Carreira-Perpiñán (2012): Learning and adaptation of a tongue shape model with missing data. *IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP 2012)*, pp. 3981–3984. Acceptance rate: 49% (1350/2700+).

M. Farhadloo and M. Á. Carreira-Perpiñán (2012): Regularising an adaptation algorithm for tongue shape models. *IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP 2012)*, pp. 4481–4484. Acceptance rate: 49% (1350/2700+).

W. Wang, M. Á. Carreira-Perpiñán and Z. Lu (2011): A denoising view of matrix completion. *Advances in Neural Information Processing Systems 20 (NIPS 2011)*, pp. 334–342. Acceptance rate: 21.8% (305/1400).

M. Á. Carreira-Perpiñán and Z. Lu (2011): Manifold learning and missing data recovery through unsupervised regression. *IEEE Int. Conf. Data Mining (ICDM 2011)*, short paper, pp. 1014–1019. Acceptance rate: 18.8% (148/786).

A. Kamthe, V. Erickson, M. Á. Carreira-Perpiñán and A. E. Cerpa (2011): Enabling building energy auditing using adapted occupancy models. *3rd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings (Buildsys 2011)*, 33–38. Acceptance rate: 34.4% (10/29).

M. Á. Carreira-Perpiñán (2011): Machine learning models of the tongue shape during speech. *9th Int. Seminar on Speech Production (ISSP 2011)*, pp. 103–110.

A. Kamthe, M. Á. Carreira-Perpiñán and A. E. Cerpa (2011): Adaptation of a mixture of multivariate Bernoulli distributions. *22nd Int. Joint Conf. Artificial Intelligence (IJCAI 2011)*, pp. 1336–1341. Acceptance rate: 30.2% (400/1325).

V. Erickson, M. Á. Carreira-Perpiñán and A. E. Cerpa (2011): OBSERVE: Occupancy-driven system for efficient reduction of HVAC energy. *10th ACM/IEEE Int. Conf. Information Processing in Sensor Networks (IPSN 2011)*, pp. 258–269. Acceptance rate: 20.7% (18/87).

C. Qin, M. Á. Carreira-Perpiñán and M. Farhadloo (2010): Adaptation of a tongue shape model by local feature transformations. *Interspeech 2010*, pp. 1596–1599. Acceptance rate: 58.2% (771/1324).

C. Qin and M. Á. Carreira-Perpiñán (2010): Estimating missing data sequences in X-ray microbeam recordings. *Interspeech 2010*, pp. 1592–1595. Acceptance rate: 58.2% (771/1324).

C. Qin and M. Á. Carreira-Perpiñán (2010): Articulatory inversion of American English /ɹ/ by conditional density modes. *Interspeech 2010*, pp. 1998–2001. Acceptance rate: 58.2% (771/1324).

L. Xie, M. Á. Carreira-Perpiñán and S. Newsam (2010): Semi-supervised regression with temporal image sequences. *Int. Conf. Image Processing (ICIP 2010)*, pp. 2637–2640. Acceptance rate: $<$ 47.6% (1190/2500+).

M. Á. Carreira-Perpiñán (2010): The Elastic Embedding algorithm for dimensionality reduction. *Int. Conf. Machine Learning (ICML 2010)*, pp. 167–174. Acceptance rate: 25.6% (152/594).

W. Wang and M. Á. Carreira-Perpiñán (2010): Manifold blurring mean shift algorithms for manifold denoising. *IEEE Computer*

Society Conf. Computer Vision and Pattern Recognition (CVPR 2010), pp. 1759–1766. Acceptance rate: 26.8% (383/1717). ▶

M. Á. Carreira-Perpiñán and Z. Lu (2010): Parametric dimensionality reduction by unsupervised regression. *IEEE Computer Society Conf. Computer Vision and Pattern Recognition (CVPR 2010)*, pp. 1895–1902. Acceptance rate: 26.8% (383/1717). ▶

C. Qin and M. Á. Carreira-Perpiñán (2010): Reconstructing the full tongue contour from EMA/X-ray microbeam. *IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP 2010)*, pp. 4190–4193. Acceptance rate: 48.8% (1353/2768).

C. Qin and M. Á. Carreira-Perpiñán (2009): The geometry of the articulatory region that produces a speech sound. Invited paper, *Proc. 43rd Annual Asilomar Conference on Signals, Systems, and Computers*, pp. 1742–1746.

A. Kamthe, M. Á. Carreira-Perpiñán and A. E. Cerpa (2009): M&M: Multilevel Markov model for wireless link simulation in sensor networks. *Proc. of the 7th ACM Conference on Embedded Networked Sensor Systems (SenSys 2009)*, pp. 57–70. Acceptance rate: 17.6% (21/119).

C. Qin and M. Á. Carreira-Perpiñán (2009): Adaptation of a predictive model of tongue shapes. *Interspeech 2009*, pp. 772–775. Acceptance rate: 57.7% (750/1300).

D. W. Massaro, M. Á. Carreira-Perpiñán and D. J. Merrill (2009): Optimizing visual feature perception for an automatic wearable speech supplement in face-to-face communication and classroom situations. *42nd Hawaii Int. Conf. Systems Science (HICSS–42 2009)*, pp. 1–10.

C. Qin, M. Á. Carreira-Perpiñán, K. Richmond, A. Wrench and S. Renals (2008): Predicting tongue shapes from a few landmark locations. *Interspeech 2008*, pp. 2306–2309. Acceptance rate: 59.9% (774/1293).

M. Á. Carreira-Perpiñán and Z. Lu (2008): Dimensionality reduction by unsupervised regression. *IEEE Computer Society Conf. Computer Vision and Pattern Recognition (CVPR 2008)*, 8 pages. Acceptance rate: 32% (508/1593).

M. Á. Carreira-Perpiñán (2008): Generalised blurring mean-shift algorithms for nonparametric clustering. *IEEE Computer Society Conf. Computer Vision and Pattern Recognition (CVPR 2008)*, 8 pages. Acceptance rate: 32% (508/1593).

Z. Lu and M. Á. Carreira-Perpiñán (2008): Constrained spectral clustering through affinity propagation. *IEEE Computer Society Conf. Computer Vision and Pattern Recognition (CVPR 2008)*, 8 pages. Acceptance rate: 32% (508/1593).

C. Qin and M. Á. Carreira-Perpiñán (2008): Trajectory inverse kinematics by conditional density modes. *IEEE Int. Conf. Robotics and Automation (ICRA 2008)*, pp. 1979–1986. Acceptance rate: 43% (641/1476).

C. Qin and M. Á. Carreira-Perpiñán (2008): Trajectory inverse kinematics by nonlinear, nongaussian tracking. *IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP 2008)*, pp. 2057–2060. Acceptance rate: 50% (1352/2729).

U. Özertem, D. Erdogmus and M. Á. Carreira-Perpiñán (2008): Density geodesics for similarity clustering. *IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP 2008)*, pp. 1977–1980. Acceptance rate: 50% (1352/2729).

Z. Lu, M. Á. Carreira-Perpiñán and C. Sminchisescu (2008): People tracking with the Laplacian Eigenmaps Latent Variable Model. *Advances in Neural Information Processing Systems 20 (NIPS 2007)*, pp. 1705–1712. Acceptance rate: 22% (217/975).

C. Qin and M. Á. Carreira-Perpiñán (2007): An empirical investigation of the nonuniqueness in the acoustic-to-articulatory mapping. *Interspeech 2007*, pp. 74–77. **Best student paper award**. Acceptance rate: 59% (748/1268).

C. Qin and M. Á. Carreira-Perpiñán (2007): A comparison of acoustic features for articulatory inversion. *Interspeech 2007*, pp. 2469–2472. Acceptance rate: 59% (748/1268).

A. Myronenko, X. Song and M. Á. Carreira-Perpiñán (2007): Free-form nonrigid image registration using generalized elastic nets. *IEEE Computer Society Conf. Computer Vision and Pattern Recognition (CVPR 2007)*, 8 pages. Acceptance rate: 28% (352/1256).

M. Á. Carreira-Perpiñán and Z. Lu (2007): The Laplacian Eigenmaps Latent Variable Model. *11th Int. Conf. Artificial Intelligence and Statistics (AISTATS 2007)*, pp. 59–66. Acceptance rate: 57% (85/150).

A. Myronenko, X. Song and M. Á. Carreira-Perpiñán (2007): Non-rigid point set registration: Coherent Point Drift. *Advances in Neural Information Processing Systems 19 (NIPS 2006)*, pp. 1009–1016. Acceptance rate: 24% (204/833).

M. Á. Carreira-Perpiñán (2006): Fast nonparametric clustering with Gaussian blurring mean-shift. *Int. Conf. Machine Learning (ICML 2006)*, pp. 153–160. Acceptance rate: $<$ 20% (140/700+).

M. Á. Carreira-Perpiñán (2006): Acceleration strategies for Gaussian mean-shift image segmentation. *IEEE Computer Society*

Conf. Computer Vision and Pattern Recognition (CVPR 2006), pp. 1160–1167. Acceptance rate: 28% (318/1131).

D. Erdogmus, M. Á. Carreira-Perpiñán and U. Özertem (2006): Kernel density estimation, affinity-based clustering, and typical cuts. *IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP 2006)*, vol. 5, pp. 569–572. Acceptance rate: 48% (1465/3045).

M. Á. Carreira-Perpiñán, P. Dayan and G. J. Goodhill (2005): Differential priors for elastic nets. *Int. Conf. Intelligent Data Engineering and Automated Learning (IDEAL 2005)*, pp. 335–342. Lecture Notes in Computer Science vol. 3578, Springer-Verlag. Acceptance rate: 46% (76/167).

M. Á. Carreira-Perpiñán and G. E. Hinton (2005): On contrastive divergence learning. *10th Int. Workshop on Artificial Intelligence and Statistics (AISTATS 2005)*, pp. 33–40. Acceptance rate: 39% (58/150).

M. Á. Carreira-Perpiñán and R. S. Zemel (2005): Proximity graphs for clustering and manifold learning. *Advances in Neural Information Processing Systems 17 (NIPS 2004)*, pp. 225–232. Acceptance rate: 25% (207/822).

X. He, R. S. Zemel and M. Á. Carreira-Perpiñán (2004): Multiscale conditional random fields for image labeling. *IEEE Computer Society Conf. Computer Vision and Pattern Recognition (CVPR 2004)*, pp. 695–702. Acceptance rate: 30% (260/873).

M. Á. Carreira-Perpiñán and C. K. I. Williams (2003): On the number of modes of a Gaussian mixture. *Scale-Space Methods in Computer Vision*, pp. 625–640, Lecture Notes in Computer Science vol. 2695, Springer-Verlag. Acceptance rate: 64% (56/88).

M. Á. Carreira-Perpiñán (2000): Reconstruction of sequential data with probabilistic models and continuity constraints. *Advances in Neural Information Processing Systems 12 (NIPS 1999)*, pp. 414–420. Acceptance rate: 32% (150/467).

M. Á. Carreira-Perpiñán and S. Renals (1999): A latent variable modelling approach to the acoustic-to-articulatory mapping problem. *14th Int. Congress of Phonetic Sciences (ICPhS 1999)*, pp. 2013–2016. Acceptance rate: 87% (703/809).

▷ **Journals (refereed)**

S. S. Hada, M. Á. Carreira-Perpiñán and A. Zharmagambetov: Sparse oblique decision trees: a tool to understand and manipulate neural net features. *Data Mining and Knowledge Discovery*, to appear.

D. A. Winkler, M. Á. Carreira-Perpiñán and A. E. Cerpa (2020): OPTICS: OPTimizing Irrigation Control at Scale. *ACM Trans. Sensor Networks*, 16(3):22.1–22.38.

D. A. Winkler, R. Wang, F. Blanchette, M. Á. Carreira-Perpiñán and A. E. Cerpa (2019): DICTUM: Distributed Irrigation aCtuation with Turf hUmidity Modeling. *ACM Trans. Sensor Networks*, 15(4):41.1–41.33.

V. Erickson, M. Á. Carreira-Perpiñán and A. E. Cerpa (2014): Occupancy modeling and prediction for building energy management. *ACM Trans. Sensor Networks*, 10(3):42.1–42.28.

A. Kamthe, M. Á. Carreira-Perpiñán and A. E. Cerpa (2013): Improving wireless link simulation using multi-level Markov models. *ACM Trans. Sensor Networks* 10(1):17.1–17.28.

S. C. Kazmierczak, T. K. Leen, D. Erdogmus and M. Á. Carreira-Perpiñán (2007): Reduction of multi-dimensional laboratory data to a two-dimensional plot: a novel technique for the identification of laboratory error. *Clinical Chemistry and Laboratory Medicine* 45(6):749–752.

M. Á. Carreira-Perpiñán (2007): Gaussian mean shift is an EM algorithm. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 29(5):767–776.

M. Á. Carreira-Perpiñán, R. J. Lister and G. J. Goodhill (2005): A computational model for the development of multiple maps in primary visual cortex. *Cerebral Cortex* 15(8):1222–1233.

M. Á. Carreira-Perpiñán and G. J. Goodhill (2004): Influence of lateral connections on the structure of cortical maps. *J. Neurophysiology* 92(5):2947–2959.

M. Á. Carreira-Perpiñán and G. J. Goodhill (2002): Are visual cortex maps optimized for coverage?. *Neural Computation* 14(7):1545–1560.

M. Á. Carreira-Perpiñán (2000): Mode-finding for mixtures of Gaussian distributions. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 22(11):1318–1323.

M. Á. Carreira-Perpiñán and S. Renals (2000): Practical identifiability of finite mixtures of multivariate Bernoulli distributions. *Neural Computation* 12(1):141–152.

M. Á. Carreira-Perpiñán and S. Renals (1998): Dimensionality reduction of electropalatographic data using latent variable models. *Speech Communication* 26(4):259–282.

▷ **Book chapters (refereed)**

M. Á. Carreira-Perpiñán (2015): Clustering methods based on kernel density estimators: mean-shift algorithms. Invited chapter in *Handbook of Cluster Analysis* (C. Hennig, M. Meila, F. Murtagh and R. Rocci, eds.), CRC/Chapman and Hall, chapter 18, pp. 383–418.

D. Archambault, K. Bunte, M. Á. Carreira-Perpiñán, D. Ebert, T. Ertl and B. Zupan (2015): Machine learning meets visualization: a roadmap for scalable data analytics. In *Bridging Information Visualization with Machine Learning* (D. A. Keim, T. Munzner, F. Rossi and M. Verleysen, eds.), Dagstuhl Reports 5(3):8–12.

D. W. Massaro, M. Á. Carreira-Perpiñán and D. J. Merrill (2010): An automatic visible speech supplement for deaf individuals' speech comprehension in face-to-face and classroom situations. In *Cued Speech and Cued Language for Deaf and Hard of Hearing Children* (C. LaSasso, J. Leybaert and K. Crain, eds.), Plural Publishing Inc., chapter 22, pp. 503–530.

M. Á. Carreira-Perpiñán and G. J. Goodhill (2002): Development of columnar structures in visual cortex. Invited chapter in *Computational Neuroanatomy: Principles and Methods* (G. A. Ascoli, ed.), Humana Press, chapter 15, pp. 337–357.

G. J. Goodhill and M. Á. Carreira-Perpiñán (2002): Cortical columns. Invited article in the *Encyclopedia of Cognitive Science* (L. Nadel, ed.), Macmillan, vol. 1, pp. 845–851.

▷ **Unrefereed technical reports and manuscripts in progress**

M. Á. Carreira-Perpiñán and S. S. Hada (2023): Very fast, approximate counterfactual explanations for decision forests. Unpublished manuscript, Mar. 5, 2023, arXiv:2303.02883 [cs.LG]. 23 pages.

M. Á. Carreira-Perpiñán and Y. Idelbayev (2021): Model compression as constrained optimization, with application to neural nets. Part V: Combining Compressions. Unpublished manuscript, Jul. 9, 2021, arXiv:2107.04380 [cs.LG]. 29 pages.

S. S. Hada, M. Á. Carreira-Perpiñán and A. Zharmagambetov (2021): Sparse oblique decision trees: a tool to understand and manipulate neural net features. Unpublished manuscript, Apr. 6, 2021, arXiv:2104.02922 [cs.LG]. 30 pages.

M. Á. Carreira-Perpiñán and S. S. Hada (2021): Counterfactual explanations for oblique decision trees: exact, efficient algorithms. Unpublished manuscript, Mar. 1, 2021, arXiv:2103.01096 [cs.LG]. 18 pages.

Y. Idelbayev and M. Á. Carreira-Perpiñán (2020): A flexible, extensible software framework for model compression based on the LC algorithm. Unpublished manuscript, May 15, 2020, arXiv:2005.07786 [cs.LG]. 15 pages.

A. Zharmagambetov, S. S. Hada, M. Á. Carreira-Perpiñán and M. Gabidolla (2020): An experimental comparison of old and new decision tree algorithms. Unpublished manuscript, Mar. 20, 2020, arXiv:1911.03054 [cs.LG]. 12 pages.

E. Eban, Y. Movshovitz-Attias, H. Wu, M. Sandler, A. Poon, Y. Idelbayev and M. Á. Carreira-Perpiñán (2019): Structured multi-hashing for model compression. Unpublished manuscript, Nov. 25, 2019, arXiv:1911.11177 [cs.LG]. 10 pages.

S. S. Hada and M. Á. Carreira-Perpiñán (2019): Style transfer by rigid alignment in neural net feature space. Unpublished manuscript, Sep. 27, 2019, arXiv:1909.13690 [cs.LG]. 15 pages.

S. S. Hada and M. Á. Carreira-Perpiñán (2019): Sampling the "inverse set" of a neuron: an approach to understanding neural nets. Unpublished manuscript, Sep. 27, 2019, arXiv: [cs.LG]. 15 pages.

M. Á. Carreira-Perpiñán and Y. Idelbayev (2017): Model compression as constrained optimization, with application to neural nets. Part II: quantization. Unpublished manuscript, Jul. 13, 2017, arXiv:1707.04319 [cs.LG]. 33 pages.

M. Á. Carreira-Perpiñán (2017): Model compression as constrained optimization, with application to neural nets. Part I: general framework. Unpublished manuscript, Jul. 5, 2017, arXiv:1707.01209 [cs.LG]. 23 pages.

M. Á. Carreira-Perpiñán and M. Alizadeh (2016): ParMAC: distributed optimisation of nested functions, with application to learning binary autoencoders. Unpublished manuscript, May 30, 2016, arXiv:1605.09114 [cs.LG]. 40 pages.

M. Á. Carreira-Perpiñán and R. Raziperchikolaei (2016): An ensemble diversity approach to supervised binary hashing. Unpublished manuscript, Feb. 3, 2016, arXiv:1602.01557 [cs.LG]. 17 pages.

R. Raziperchikolaei and M. Á. Carreira-Perpiñán (2016): Optimizing affinity-based binary hashing using auxiliary coordinates. Unpublished manuscript, Feb. 5, 2016, arXiv:1501.05352 [cs.LG]. 22 pages.
Version 1: Jan. 21, 2015, arXiv:1501.05352 [cs.LG], 18 pages.

M. Á. Carreira-Perpiñán (2015): A review of mean-shift algorithms for clustering. Unpublished manuscript, Mar. 2, 2015, arXiv:1503.00687 [cs.LG]. 28 pages.

M. Á. Carreira-Perpiñán and R. Raziperchikolaei (2015): Hashing with binary autoencoders. Unpublished manuscript, Jan. 5, 2015, arXiv:1501.00756 [cs.LG]. 22 pages.

M. Á. Carreira-Perpiñán (2014): An ADMM algorithm for solving a proximal bound-constrained quadratic program. Unpublished manuscript, Dec. 29, 2014, arXiv:1412.8493 [math.OC]. 5 pages.

W. Wang and M. Á. Carreira-Perpiñán (2014): The Laplacian K-modes algorithm for clustering. Unpublished manuscript, Jun. 15, 2014, arXiv:1406.3895 [cs.LG]. 14 pages.

W. Wang and M. Á. Carreira-Perpiñán (2014): The role of dimensionality reduction in classification. Unpublished manuscript, May 25, 2014, arXiv:1405.6444 [cs.LG]. 15 pages.

M. Á. Carreira-Perpiñán and W. Wang (2014): LASS: a simple assignment model with Laplacian smoothing. Unpublished manuscript, May 22, 2014, arXiv:1405.5960 [cs.LG]. 20 pages.

W. Wang and M. Á. Carreira-Perpiñán (2013): Projection onto the probability simplex: An efficient algorithm with a simple proof, and an application. Unpublished manuscript, Sep. 3, 2013, arXiv:1309.1541 [cs.LG]. 5 pages.

M. Á. Carreira-Perpiñán and W. Wang (2013): The K-modes algorithm for clustering. Unpublished manuscript, Apr. 23, 2013, arXiv:1304.6478 [cs.LG]. 13 pages.

M. Á. Carreira-Perpiñán and W. Wang (2012): Distributed optimization of deeply nested systems. Unpublished manuscript, Dec. 24, 2012, arXiv:1212.5921 [cs.LG]. 21 pages.

M. Á. Carreira-Perpiñán and C. K. I. Williams (2003): An isotropic Gaussian mixture can have more modes than components. Tech. rep. EDI–INF–RR–0185, School of Informatics, University of Edinburgh, UK. 8 pages.

M. Á. Carreira-Perpiñán and C. K. I. Williams (2003): On the number of modes of a Gaussian mixture. Tech. rep. EDI–INF–RR–0159, School of Informatics, University of Edinburgh, UK. 17 pages.

M. Á. Carreira-Perpiñán (2004): Reconstruction of sequential data with density models. Unpublished manuscript, Jan. 27, 2004, arXiv:1109.3248 [cs.LG]. 30 pages.

M. Á. Carreira-Perpiñán and G. J. Goodhill (2003): Generalised elastic nets. Unpublished manuscript, Aug. 14, 2003, arXiv:1108.2840 [q-bio.NC]. 52 pages.

M. Á. Carreira-Perpiñán (1999): Mode-finding for mixtures of Gaussian distributions. Tech. rep. CS–99–03, Dept. of Computer Science, University of Sheffield, UK. 23 pages.

M. Á. Carreira-Perpiñán (1997): Density networks for dimension reduction of continuous data: Analytical solutions. Tech. rep. CS–97–09, Dept. of Computer Science, University of Sheffield, UK. 8 pages.

M. Á. Carreira-Perpiñán (1996): A review of dimension reduction techniques. Tech. rep. CS–96–09, Dept. of Computer Science, University of Sheffield, UK. 69 pages.

▷ **Lightly refereed conference or workshop papers, abstracts and other publications**

M. Klemen, M. Á. Carreira-Perpiñán and P. López-García (2023): Solving recurrence relations using machine learning, with application to cost analysis. *10th Workshop on Horn Clauses for Verification and Synthesis (HCVS 2023)*, Apr. 23, 2023. 14 pages.

S. S. Hada and M. Á. Carreira-Perpiñán (2021): Exploring counterfactual explanations for classification and regression trees. *Int. Workshop and Tutorial on eXplainable Knowledge Discovery in Data Mining* at *ECML 2021*, Sep. 13, 2021. 16 pages.

Y. Idelbayev and M. Á. Carreira-Perpiñán (2021): A flexible, extensible software framework for model compression based on the LC algorithm. *2nd On-Device Intelligence Workshop* at *(MLSys 2021)*, Apr. 9, 2021. 5 pages. ▶

Y. Idelbayev and M. Á. Carreira-Perpiñán (2020): A flexible, extensible software framework for model compression based on the LC algorithm. Extended abstract, *Bay Area Machine Learning Symposium (BayLearn 2020)*, virtual event, Oct. 15, 2020. 3 pages. ▶

M. Gabidolla, A. Zharmagambetov and M. Á. Carreira-Perpiñán (2020): Boosted sparse oblique decision trees. Extended abstract, *Bay Area Machine Learning Symposium (BayLearn 2020)*, virtual event, Oct. 15, 2020. 3 pages. ▶

M. Á. Carreira-Perpiñán and S. S. Hada (2020): Inverse classification with logistic and softmax classifiers: efficient optimization.

Workshop *Beyond first order methods in machine learning systems* at *ICML 2020*, Jul. 17, 2020. 11 pages.

M. Á. Carreira-Perpiñán and A. Zharmagambetov (2018): Fast model compression. Extended abstract, *Bay Area Machine Learning Symposium (BayLearn 2018)*, Facebook, Menlo Park, CA, Oct. 11, 2018. 3 pages.

S. S. Hada and M. Á. Carreira-Perpiñán (2018): Sampling the "inverse set" of a neuron. Extended abstract, *Bay Area Machine Learning Symposium (BayLearn 2018)*, Facebook, Menlo Park, CA, Oct. 11, 2018. 3 pages.

M. Á. Carreira-Perpiñán and Y. Idelbayev (2017): Model compression as constrained optimization, with application to neural nets. Workshop on *Optimization for Machine Learning* at *NIPS 2017*, Dec. 8, 2017. 4 pages.

R. Raziperchikolaei and M. Á. Carreira-Perpiñán (2017): Optimizing circulant support vector machines: the exact solution. Workshop on *Optimization for Machine Learning* at *NIPS 2017*, Dec. 8, 2017. 5 pages.

M. Á. Carreira-Perpiñán and Y. Idelbayev (2017): Optimal neural net compression via constrained optimization. Workshop on *Machine Learning on the Phone and other Consumer Devices* at *NIPS 2017*, Dec. 9, 2017. 3 pages.

M. Á. Carreira-Perpiñán and R. Raziperchikolaei (2017): Learning supervised binary hashing without binary code optimization. Workshop on *Nearest Neighbors for Modern Applications with Massive Data: An Age-old Solution with New Challenges* at *NIPS 2017*, Dec. 8, 2017. 4 pages.

M. Á. Carreira-Perpiñán and Y. Idelbayev (2017): Model compression as constrained optimization, with application to neural nets. Extended abstract, *Bay Area Machine Learning Symposium (BayLearn 2017)*, Apple, Cupertino, CA, Oct. 19, 2017. 3 pages. ▶

M. Á. Carreira-Perpiñán and R. Raziperchikolaei (2017): Learning supervised binary hashing without binary code optimization. Extended abstract, *Bay Area Machine Learning Symposium (BayLearn 2017)*, Apple, Cupertino, CA, Oct. 19, 2017. 3 pages.

R. Raziperchikolaei and M. Á. Carreira-Perpiñán (2016): Optimizing affinity-based binary hashing using auxiliary coordinates. Extended abstract, *Bay Area Machine Learning Symposium (BayLearn 2016)*, LinkedIn, Sunnyvale, CA, Oct. 6, 2016. 2 pages.

R. Raziperchikolaei and M. Á. Carreira-Perpiñán (2016): Optimizing affinity-based binary hashing using auxiliary coordinates. Workshop on *Non-Convex Analysis and Optimization* at *ICML 2016*, Jun. 23, 2016. 5 pages.

R. Raziperchikolaei and M. Á. Carreira-Perpiñán (2016): Optimizing binary autoencoders using auxiliary coordinates, with application to learning binary hashing. Workshop on *Optimization Methods for the Next Generation of Machine Learning* at *ICML 2016*, Jun. 24, 2016. 5 pages.

D. A. Winkler, R. Wang, F. Blanchette, M. Á. Carreira-Perpiñán and A. E. Cerpa (2009): MICO: Model-Based Irrigation Control Optimization. Poster abstract, *Proc. of the 13th ACM Conference on Embedded Networked Sensor Systems (SenSys 2015)*, pp. 409–410.

R. Raziperchikolaei and M. Á. Carreira-Perpiñán (2015): Hashing with binary autoencoders. *2015 INFORMS Workshop on Data Mining and Analytics*, 6 pages.

M. Á. Carreira-Perpiñán and R. Raziperchikolaei (2015): An ensemble diversity approach to binary hashing. Extended abstract, *Bay Area Machine Learning Symposium (BayLearn 2015)*, Menlo Park, CA, Oct. 22, 2015. 2 pages. ▶

M. Á. Carreira-Perpiñán and M. Alizadeh (2015): Distributed optimization of binary autoencoders using auxiliary coordinates. Extended abstract, *Bay Area Machine Learning Symposium (BayLearn 2015)*, Menlo Park, CA, Oct. 22, 2015. 2 pages.

M. Á. Carreira-Perpiñán and R. Raziperchikolaei (2014): Hashing with binary autoencoders. Extended abstract, *Bay Area Machine Learning Symposium (BayLearn 2014)*, Berkeley, CA, Oct. 21, 2014. 2 pages.

M. Á. Carreira-Perpiñán and M. Vladymyrov (2014): A fast, universal algorithm to learn parametric nonlinear embeddings. Extended abstract, *Bay Area Machine Learning Symposium (BayLearn 2014)*, Berkeley, CA, Oct. 21, 2014. 2 pages.

M. Vladymyrov and M. Á. Carreira-Perpiñán (2013): Linear-time training of nonlinear low-dimensional embeddings. Extended abstract, *Bay Area Machine Learning Symposium (BayLearn 2013)*, Facebook, Menlo Park, CA, Aug. 28, 2013. 2 pages.

M. Á. Carreira-Perpiñán and W. Wang (2013): A simple assignment model with Laplacian smoothing. Extended abstract, *Bay Area Machine Learning Symposium (BayLearn 2013)*, Facebook, Menlo Park, CA, Aug. 28, 2013. 2 pages.

M. Vladymyrov and M. Á. Carreira-Perpiñán (2013): Locally Linear Landmarks for large-scale manifold learning. Workshop on *Spectral Learning* at *ICML 2013*, Jun. 20, 2013. 5 pages.

M. Vladymyrov and M. Á. Carreira-Perpiñán (2012): Fast training of graph-based algorithms for nonlinear dimensionality re-

duction. Extended abstract, *Bay Area Machine Learning Symposium (BayLearn 2012)*, Google, Mountain View, CA, Aug. 30, 2012. 2 pages.

M. Á. Carreira-Perpiñán and W. Wang (2012): Fast algorithms for learning deep neural networks. Extended abstract, *Bay Area Machine Learning Symposium (BayLearn 2012)*, Google, Mountain View, CA, Aug. 30, 2012. 2 pages.

A. Kamthe, M. Á. Carreira-Perpiñán and A. E. Cerpa (2009): Wireless link simulations using multi-level Markov models. Poster abstract, *Proc. of the 7th ACM Conference on Embedded Networked Sensor Systems (SenSys 2009)*, pp. 391–392.

D. W. Massaro, M. Á. Carreira-Perpiñán, D. J. Merrill, C. Sterling, S. Bigler, E. Piazza and M. Perlman (2008): iGlasses: an automatic wearable speech supplement in face-to-face communication and classroom situations. Demonstration paper, *Proc. Int. Conf. Multimodal Interfaces (ICMI 2008)*, pp. 197–198.

Z. Lu, M. Á. Carreira-Perpiñán and C. Sminchisescu (2007): People tracking with the Laplacian Eigenmaps Latent Variable Model. Invited extended abstract, *The Learning Workshop 2007*, San Juan, Puerto Rico, Mar. 19–22, 2007. 2 pages.

A. Myronenko, X. Song and M. Á. Carreira-Perpiñán (2006): Non-parametric image registration using generalized elastic nets. *Proc. Int. Workshop on Mathematical Foundations of Computational Anatomy (MFCA 2006)*, pp. 156–163 (satellite workshop at the *9th Int. Conf. Medical Image Computing & Computer Assisted Intervention (MICCAI 2006)*).

M. Á. Carreira-Perpiñán and G. J. Goodhill (2001): The effect of variable elastic topologies on the structure of ocular dominance and orientation maps. *Society for Neuroscience Abstracts* 27:475.21.

M. Á. Carreira-Perpiñán (1999): One-to-many mappings, continuity constraints and latent variable models. *Proc. IEE Colloquium on Applied Statistical Pattern Recognition*, pp. 14/1–14/6, Birmingham, UK, Apr. 20, 1999.

M. Á. Carreira-Perpiñán and S. Renals (1998): Experimental evaluation of latent variable models for dimensionality reduction. *Proc. IEEE Signal Processing Society Workshop on Neural Networks for Signal Processing (NNSP 1998)*, pp. 165–173, Cambridge, UK.

M. Á. Carreira-Perpiñán (1994): The modelling of the thermal subsystem in spacecraft real-time simulators. *Proc. 3rd Workshop on Simulators for European Space Programmes*, pp. 69–78, Noordwijk, The Netherlands.

M. Á. Carreira-Perpiñán (1994): ESA Thermal Analysis Program (ESATAN): An evaluation. Tech. rep. DOPS–SST–TN–0215–SIM, European Space Operations Centre (ESA), Darmstadt, Germany. 54 pages.

▷ Theses

M. Á. Carreira-Perpiñán (2001): *Continuous latent variable models for dimensionality reduction and sequential data reconstruction*. PhD thesis, University of Sheffield, UK.

M. Á. Carreira-Perpiñán (1995): *Compression neural networks and feature extraction: Application to human recognition from ear images*. MSc thesis, Facultad de Informática, Technical University of Madrid.

▷ White papers and working group reports

University of California Presidential Working Group on AI (2021): Responsible Artificial Intelligence: Recommendations to Guide the University of California's Artificial Intelligence Strategy. Final report, University of California, Oct. 2021.

▷ Lecture notes

M. Á. Carreira-Perpiñán (2019): *CSE176 Introduction to machine learning: Lecture notes*. UC Merced, 2015–2021.

M. Á. Carreira-Perpiñán (2020): *EECS260 Optimization: Lecture notes*. UC Merced, 2008–2020.

M. Á. Carreira-Perpiñán (2006): *EE 589/689 Foundations of computer vision: Lecture notes*. OGI/OHSU, Fall quarter 2006.

▷ Other publications

M. Á. Carreira-Perpiñán and M. Gabidolla (2023): videos presented at the CVPR 2023 Art Gallery (Jun. 18–22, 2023, Vancouver, Canada), created using the TAO/FAO algorithms for optimizing decision trees and forests: ▶ *Variations on Warhol's Marilyn*, ▶ *Variations on TAO*, ▶ *Mimicking great paintings by oblique decision forests*.

M. Á. Carreira-Perpiñán (2005): MathWorld entry "Simplex Method," `http://mathworld.wolfram.com/SimplexMethod.html`.

RESEARCH GRANTS

Total funding as PI over $1.79M as of 2022.

▷ **Current**

NSF I-Corps Teams (#2228243): *I-Corps: Tree-based artificial intelligence (AI) models for financial fraud detection*, $50,000 (2022–2023), sole PI.

NSF (IIS #2007147): *The TAO algorithm: principled, efficient optimization of decision trees, forests, tree-based neural nets, and beyond*, $425,000 (2020–2023), sole PI.

▷ **Past, as PI**

UC Merced Academic Senate Faculty Research Grant: *Learning decision trees and deep nets with the TAO algorithm*, $5,000 (2019), sole PI.

NVIDIA GPU Grant Program: Titan X Pascal (Feb. 2017), Titan Xp (May 2018), Titan V (Sep. 2019), sole PI.

Xilinx University Program: FPGA board (Nov. 2017), with Prof. Dong Li (UC Merced).

NSF (BCS #1626505): *MRI: Acquisition of robotic tools for studying brain, behavior and embodied cognition*, $182,806 (2016–2019) (my part: $0). R. Balasubramaniam (PI, UC Merced), S. Carpin, M. J. Spivey (co-PIs, UC Merced), M. Á. Carreira-Perpiñán and others (senior personnel, UC Merced).

UC Merced Academic Senate Faculty Research Grant: *A modern optimization approach to deep neural net compression*, $5,000 (2017), sole PI.

NSF (IIS #1423515): *Algorithms for accelerating optimization in deep learning*, $450,000 (2014–2018), sole PI.

REU supplement to NSF IIS #1423515 award: $8,000 (Summer 2015), sole PI.

Google Faculty Research Award: *Parallelization of deep neural net training using the method of auxiliary coordinates*, $60,237 (2013–2014), sole PI.

UC Merced Graduate Research Council Faculty Research Grant: *Linear-time training of nonlinear low-dimensional embeddings*, $4,710 (2013), sole PI.

NSF (IIS #0711186): *Foreign accent conversion through mapping inversion of the vocal-tract frontal cavity*, $220,000 (2008–2011), PI. Collaborative proposal with R. Gutiérrez-Osuna (Texas A & M University).

NSF CAREER (#0546857, #0754089): *Machine learning approaches for articulatory inversion*, $500,000 (2006–2011), sole PI.

REU supplement to NSF CAREER award (#0754089): $8,000 (Summer 2010), sole PI.

REU supplement to NSF CAREER award (#0754089): $8,000 (Summer 2009), sole PI.

▷ **Past, as co-PI or collaborator**

UC Merced Academic Senate Faculty Research Grant: *Tongue shape estimation via magnetic sensing: recording device and AI-based algorithms*, $20,000 (2021), co-PI (my part: $10 000). PI: H. Huang (UC Merced).

NSF (DGE #1633722): *NRT-DESE Intelligent Adaptive Systems: training computational and data-analytics skills for academia and industry*, $2,921,681 (2016–2021) (my part: $0). R. Balasubramaniam (PI, UC Merced), S. Carpin, C. T. Kello, S. S. Sindi (co-PIs, UC Merced), M. Á. Carreira-Perpiñán and others (core participant, UC Merced).

EPSRC EP/I027696/1: *Ultrax: Real-time tongue tracking for speech therapy using ultrasound*, £586,154 (2/2011–7/2014) (my part: $0). S. Renals (PI, U. Edinburgh), J. Cleland (co-PI, Queen Margaret U.), K. Richmond (co-PI, U. Edinburgh), M. Á. Carreira-Perpiñán (UC Merced), external consultant.

Center for Information Technology Research in the Interest of Society (CITRIS) #81: *The Sensing Entity Tracking Initiative (SETI): A CITRIS center for smart energy infrastructure*, $75,000 (2010–2011). A. Cerpa (PI, UC Merced), M. Á. Carreira-Perpiñán (co-PI) and E. Arens (co-PI, UC Berkeley).

Center for Information Technology Research in the Interest of Society (CITRIS) #36: *An automatic wearable speech supplement in face-to-face and classroom situations*, $75,000 (2008–2009). D. W. Massaro (PI, UC Santa Cruz), M. Á. Carreira-Perpiñán (co-PI) and Y. Zhang (co-PI, UC Santa Cruz).

NSF (BCS #0821766): *MRI: Acquisition of robotic hardware for humanoid research in cognitive science and engineering*, $492,500 (2008–2011) (my part: $0). S. Carpin (PI, UC Merced), D. Noelle, T. Matlock, S. Newsam, M. Kallmann (co-PIs, UC Merced), M. Á. Carreira-Perpiñán, S. Oh (senior personnel, UC Merced).

TEACHING GRANTS

UC Merced Zero-Cost Course Materials grant (for *CSE176 Introduction to machine learning*), $1,500 (2021), sole PI.

Intel Faculty Fellowship Award for curriculum development: *Foundations of computer vision* (OGI School of Science & Eng., Oregon Health & Science University.), $24,864 (2006), sole PI.

ENTREPRENEURSHIP

- Summer 2022: participation in the NSF I-Corps Teams program as technical lead (team "TAO Trees", with my PhD students Arman Zharmagambetov and Magzhan Gabidolla as entrepreneurship leads and Amer Kayani as industrial mentor).
- Fall 2021: participation in the CITRIS Foundry incubator (team "TAO Trees", with my PhD students Arman Zharmagambetov, Magzhan Gabidolla and Yerlan Idelbayev).
- Fall 2021: participation in the NSF I-Corps Regional course (UC Berkeley) as technical lead (team "TAO Trees", with my PhD students Arman Zharmagambetov, Magzhan Gabidolla and Yerlan Idelbayev).

AWARDS

- Apr. 2014: notable paper award (AISTATS 2014), "Distributed optimization of deeply nested systems."
- Aug. 2013: Google Faculty Research Award.
- Aug. 2007: best student paper award (Interspeech 2007), "An empirical investigation of the nonuniqueness in the acoustic-to-articulatory mapping."
- Jan. 2006: NSF CAREER award.

Other awards and scholarships

- 2016–2017 UC Merced Senate Award for Excellence in Faculty Mentorship.
- Dec. 1999: travel grant, NIPS 1999.
- Aug. 23–Sep. 17, 1999: scholarship to attend the European Union Advanced Course in *Computational Neuroscience*. International Centre for Theoretical Physics, Trieste, Italy.
- Oct. 1997 to Dec. 1998: bursary, University of Sheffield.
- Aug. 4–15, 1997: scholarship to attend the NATO Advanced Study Institute *Generalization in Neural Networks and Machine Learning*. Isaac Newton Institute, Cambridge, UK.
- Sep. 27–Oct. 7, 1996: scholarship to attend the NATO Advanced Study Institute *Learning in Graphical Models*. E. Majorana Centre, Erice, Italy.
- Oct. 1995 to Sep. 1997: predoctoral (FPU) scholarship, Spanish Ministry of Education and Science.
- Jan. 1993 to Dec. 1994: postgraduate scholarship, European Space Agency/Spanish Centre for Technological & Industrial Development.

Awards to my PhD students

- 2020: Yerlan Idelbayev: UC Merced Outstanding Teaching Award, $1,000.
- 2019: Arman Zharmagambetov: UC Merced Outstanding Teaching Award, $1,000.
- 2017–2021: Arman Zharmagambetov: UC Merced Chancellor's Graduate Fellowship.
- Mar. 2014: Max Vladymyrov: winner, UC Merced Graduate Student Research Poster Competition (School of Engineering).
- Spring 2013: Max Vladymyrov, EECS Graduate Group Fellowship Award, $7,500.
- Spring 2013: Weiran Wang, EECS Graduate Group Fellowship Award, $7,500.
- Oct. 2007 to Mar. 2008: Chao Qin, Marie Curie Fellowship from the European Union (6-month stay at the Centre for Speech Technology Research, University of Edinburgh).

STUDENT SUPERVISION

PhD students (EECS, UC Merced):

- Rasul Kairgeldin, Kuat Gazizov (since Jan. 2022); Magzhan Gabidolla (since Jan. 2020).
- Arman Zharmagambetov (Aug. 2017 – Dec. 2022): *Learning tree-based models with manifold regularization: alternating optimization algorithms*. First appointment: postdoctoral fellow, Fundamental AI Research (FAIR), Meta AI.
- Suryabhan Singh Hada (Aug. 2016 – Jun. 2022): *Some approaches to interpret deep neural networks*. First appointment: AI engineer, LinkedIn.
- Yerlan Idelbayev (Aug. 2016 – Dec. 2021): *Low-rank compression of neural networks: LC algorithms and open-source implementation*. First appointment: applied scientist, Amazon.
- Max Vladymyrov (Aug. 2009 – Dec. 2014): *Large-scale methods for nonlinear manifold learning*. First appointment: research scientist, Yahoo Labs. Currently at Google.
- Weiran Wang (Sep. 2008 – Dec. 2013): *Mean-shift algorithms for manifold denoising, matrix completion and clustering*. First appointment: postdoc, TTI Chicago. Currently at Salesforce Research.
- Ankur Kamthe (Jul. 2007 – May 2012), co-supervised with Alberto Cerpa: *Data-driven modeling of phenomena in wireless sensor networks*. First appointment: software engineer, Aruba Networks.
- Chao Qin (Sep. 2005 – May 2011): *Data-driven approaches to articulatory speech processing*. First appointment: research scientist, `sharethis.com`. Currently at Google.

I also worked closely with Zhengdong Lu (PhD OGI/OHSU 2008).

MS students (CSEE, OGI/OHSU): Bing Liu (2006).

Undergraduate students (CSE, UC Merced): Edward Smith (NSF REU Summer 2009); David King (NSF REU Summer 2010); Jimmy Yih, Mark Crompton (NSF REU Summer 2011); Michael Mcbride, Steven Ng (NSF REU Summer 2015); Marcus Chow (Sep. 2015 – May 2016); Armando Montanez (NSF REU Summer 2017); Michael Wibawa (Spring 2019); Eshant Prakash (Spring 2020).

Visiting students (UC Merced): Javier López-Alberca (May 2013 – Aug. 2013).

TEACHING EXPERIENCE

Professor (EECS, UC Merced), all 4-credit semester courses unless otherwise noted:

- *CSE176 Introduction to machine learning* (Fall 2015–2016, 2019, 2021; Spring 2023): undergraduate course.
- *CSE100 Algorithm design and analysis* (Fall 2018, 2020): undergraduate course.
- *EECS260 Optimization* (Fall 2015–2016, 2018–2020): graduate course.
- *EECS282 Advanced topics in machine learning* (Fall 2017, 2021): graduate course.
- *EECS290 EECS Seminar* (Spring 2020): 1 credit.
- *CSE195 Upper division undergraduate research* (Summer 2015, 4 credits; Spring 2016, 1 credit): undergraduate individualised instruction.

Associate professor (EECS, UC Merced), all 4-credit semester courses unless otherwise noted:

- *CSE100 Algorithm design and analysis* (Spring 2011–2015): undergraduate course.
- *EECS260 Optimization* (Fall 2011–2012, 2014): graduate course.
- *EECS282 Advanced topics in machine learning* (Fall 2010): graduate course.
- *EECS290 EECS Seminar* (Spring 2014): 1 credit.
- *CSE195 Upper division undergraduate research* (Spring 2012): undergraduate individualised instruction; 2 credits.

Assistant professor (EECS, UC Merced), all 4-credit semester courses unless otherwise noted:

- *CSE100 Algorithm design and analysis* (Fall 2009): undergraduate course.
- *CSE140 Computer architecture* (Fall 2007): new undergraduate course.
- *CSE150 Operating systems* (Spring 2008): new undergraduate course.
- *CSE260 Optimization* (Spring 2008): new graduate course.

- *CSE282 Advanced topics in machine learning* (Fall 2008): new graduate course.
- *EECS260 Optimization* (Spring 2010): graduate course.
- *EECS290 EECS Seminar* (Fall 2009): 1 credit.

Assistant professor (Dept. Computer Science & Electrical Engineering, OGI/OHSU), all graduate quarter courses:

- *EE 589/689 Foundations of computer vision* (Fall 2006): new course I created with financial support from Intel; 3 credits.
- *MATH 519/619 Optimization* (Spring 2005, Winter 2006, Spring 2007); 3 credits.

Teaching assistant (Dept. Computer Science, University of Sheffield, UK), all semester courses: *COM129 & COM130 Foundations of computer science I & II* (1996–1999); *COM101 Functional programming with LISP* (1997–1999); *COM201 Data structures and algorithms* (1996–1999); *COM106 Object-oriented programming with Java* (1997–1999); *COM305 Advanced object-oriented design* (1998–1999); *COM240 HCI & computer graphics* (1997–1998); *COM106 Structured programming with Modula/2* (1996–1997).

PROFESSIONAL ACTIVITIES

- Action editor, *J. Machine Learning Research* (Jan. 2017 to present).
- Associate editor, *IEEE Trans. Pattern Analysis and Machine Intelligence* (Sep. 2011 to Feb. 2016).
- Senior area chair for the following conferences: *Neural Information Processing Systems (NeurIPS 2020)*; *AAAI Conf. Artificial Intelligence (AAAI 2020, 2024)*.
- Area chair/senior program committee member for the following conferences: *Neural Information Processing Systems (NIPS 2013–2014, 2016–2017, NeurIPS 2018–2019)*; *Int. Conf. Machine Learning (ICML 2018–2019)*; *Int. Conf. Artificial Intelligence and Statistics (AISTATS 2018–2019)*; *European Conf. Computer Vision (ECCV 2016)*; *SIAM Data Mining Conference (SDM 2012–2013)*.
- Program committee member for the following conferences: *Int. Conf. Machine Learning (ICML 2004, 2011)*; *Int. Conf. Artificial Intelligence and Statistics (AISTATS 2007, 2009, 2011, 2014, 2015)*; *European Conf. Computer Vision (ECCV 2008)*; *CVPR 2007 Workshop on "Component Analysis for Computer Vision"*; *CVPR 2009 "1st Int. Workshop on Visual Scene Understanding (ViSU'09)"*; *IEEE Int. Workshop on Neural Networks for Signal Processing (NNSP 2002–2003)*.
- Reviewer for the following conferences: *Neural Information Processing Systems (NIPS 2003–2010,2012)*; *Int. Conf. Machine Learning (ICML 2004, 2008, 2011, 2013, 2015, 2017)*; *Int. Conf. Artificial Intelligence and Statistics (AISTATS 2007, 2009, 2012)*; *IEEE Computer Society Conf. Computer Vision and Pattern Recognition (CVPR 2003–2004, 2006–2007, 2009–2011)*; *Int. Conf. Computer Vision (ICCV 2007, 2009)*; *European Conf. Computer Vision (ECCV 2008, 2010)*; *SIGGRAPH (2018)*; *IEEE Int. Conf. Robotics and Automation (ICRA 2009)*; *IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS 2009)*; *IEEE Int. Workshop on Neural Networks for Signal Processing (NNSP 2002–2003)*; *Interspeech (2003)*; *Interspeech 2013 Workshop on "Speech Production in Automatic Speech Recognition"*; and *IEEE Int. Conf. Acoustics, Speech and Signal Proc. (ICASSP 2004–2005)*.
- Reviewer for the following journals: *IEEE Trans. Pattern Analysis and Machine Intelligence*; *J. Machine Learning Research*; *Machine Learning*; *SIAM Journal on Mathematics of Data Science*; *IEEE Trans. Neural Networks*; *IEEE Trans. Knowledge and Data Engineering*; *IEEE Trans. Image Processing*; *Int. J. Computer Vision*; *Computer Vision and Image Understanding*; *J. Mathematical Imaging and Vision*; *Speech Communication*; *Computer Speech and Language*; *J. Acoustical Society of America*; *Network: Computation in Neural Systems*; *Neural Systems & Circuits*; *IEEE Signal Processing Letters*; *Pattern Recognition*; *IEEE Trans. Systems, Man, and Cybernetics (Part A)*; *Pattern Recognition Letters*; *Neural Processing Letters*; *Neurocomputing*; *Computational Statistics and Data Analysis*; *Physica A*; *Computer Aided Design*; *Knowledge and Information Systems*; *Knowledge Engineering Review*; *Connection Science*; *Data & Knowledge Engineering*; *EURASIP J. Applied Signal Processing*; *Annals of the Institute of Statistical Mathematics*; *Scientific Reports*; and *Neural Networks*.
- NSF grant proposal reviewing: 2007 (one panel); 2008 (one panel); 2009 (one ad-hoc review); 2012 (two panels); 2013 (one panel); 2015 (one panel, two ad-hoc reviews); 2016 (one ad-hoc review); 2020 (two panels); 2022 (one panel).
- Reviewer for grant proposals from several national agencies in Europe (Austria, Belgium).
- Book reviewer (proposal, chapters, whole book) for the following publishers: *MIT Press*; *Chapman & Hall/CRC Press*.
- Session chair at the following conferences: *IEEE Int. Conf. Data Mining (ICDM 2016)*, *Int. Conf. Artificial Intelligence and Statistics (AISTATS 2012)*, *IEEE Int. Conf. Robotics and Automation (ICRA 2008)*.
- Member of: the *IEEE* (2006–), IEEE senior member; *IEEE Signal Processing Society* (2012, 2016–2017, 2021–2022); the European Laboratory for Learning and Intelligent Systems (ELLIS Society) (2018–); *International Speech Communication*

Association, ISCA (2010–2011); *Society for Neuroscience* (2001–2002); *Association for the Advancement of Science and Technology in Spain* (AACTE) (1999–2003).

UNIVERSITY SERVICE

University of California:

- University of California Presidential Working Group on Artificial Intelligence: member (2020–2021)

UC Merced:

- Academic Senate Committee on Research (COR): member (Fall 2015 – Fall 2020)
- COR representative in the Export Control Work Group (ECWG) (Fall 2018 – Spring 2021)
- COR representative in the Senate Library & Scholarly Communications Committee (LASC) (Fall 2019 – Spring 2020)
- Academic Senate Committee on Committees (CoC): member (2009–2011)
- Faculty Academic Computing Task Force: member (2010)

School of Engineering, UC Merced:

- Faculty chair (2015–2016)
- Executive committee: member-at-large (2011–2012), vice-chair (2014–2015), chair (2015–2016)
- Academic personnel committee: observer (2008–2009), member (2009–2011); member or chair of 22 APC subcommittees for merit, mid-career and promotion cases (2010–2023)
- Curriculum committee: member (2007–2008)
- TA allocation committee: member (Mar. 2007–present)
- Innovate2Grow faculty moderator (2015, 2023)

CSE department and EECS graduate group, UC Merced:

- Faculty mentor for 4 junior faculty members in CSE (from Jan. 2014, May 2015, Aug. 2020, Mar. 2023, respectively)
- EECS admission committee: member (Fall 2018 – Spring 2022)
- EECS TA allocation committee: chair (2015–2016)
- CS faculty recruitment committee as chair: 2008–2009 (1), 2018–2019 (1)
- CS faculty recruitment committees as member: 2007–2008 (2), 2009–2010 (1), 2010–2011 (1), 2012–2013 (1), 2019–2020 (2), 2020–2021 (2), 2021–2022 (3), 2022–2023 (1)
- Academic personnel: chair or member of multiple committees for merit, mid-career and promotion cases (2017–present)

INVITED SEMINARS

- Dec. 14, 2022: Instituto IMDEA Software, Facultad de Informática, Universidad Politécnica de Madrid
- Mar. 3, 2022: Redwood seminar, UC Berkeley ▶
- Jun. 30 to Jul. 6, 2019 (4 lectures): Summer School "Optimization, Big Data and Applications (OBA)", 2nd ed., Veroli, Italy
- Jan. 18, 2019: Instituto IMDEA Software, Facultad de Informática, Universidad Politécnica de Madrid
- Jan. 16, 2019: Signal Theory and Communications Department, Universidad Carlos III de Madrid
- Jan. 9, 2019: Machine Learning Group, Escuela Politécnica Superior, Universidad Autónoma de Madrid
- Jun. 19, 2018: Amazon, Seattle
- Jun. 12, 2018: Allen Institute for Artificial Intelligence (AI2), Seattle
- May 11, 2018: Department of Statistics, University of Washington
- Apr. 5, 2018: Machine Learning and Optimization group, Microsoft Research, Redmond
- Mar. 1, 2018: Geometric Data Analysis Reading Group, University of Washington
- Jan. 12, 2018: Instituto IMDEA Software, Facultad de Informática, Universidad Politécnica de Madrid
- Jan. 11, 2017: Instituto IMDEA Software, Facultad de Informática, Universidad Politécnica de Madrid
- Jul. 12, 2016: machine learning seminar, Google

- Jul. 11, 2016: speech group seminar, Google
- Jul. 1, 2016: invited keynote, CVPR 2016 Third Workshop *DeepVision: Deep learning in computer vision* ▶
- Feb. 24, 2016: Redwood seminar, UC Berkeley ▶
- Nov. 30, 2015: Mini Workshop on Optimization, UC Merced
- Nov. 1, 2015: INFORMS 2015 Annual Meeting, session on "Distributed and Parallel Optimization"
- May 2, 2014: Electrical Engineering & Computer Science, UC Merced
- Jan. 31, 2014: Electrical Engineering & Computer Science, UC Merced
- Jan. 10, 2014: Perception for Computers and Robots Group, Facultad de Informática, Universidad Politécnica de Madrid
- Jan. 8, 2014: Machine Learning Group, Escuela Politécnica Superior, Universidad Autónoma de Madrid
- Nov. 6, 2013: Statistics Reading Group, CMU
- Oct. 28, 2013: Machine Learning Department, CMU ▶
- Oct. 16, 2013: guest lecture on course *10805 Special Topics: Deep Learning*, Machine Learning Department, CMU
- Sep. 16, 2013: Robotics Institute, CMU
- Sep. 9, 2011: 1st Southern California Workshop on Machine Learning (SoCaML 2011)
- Feb. 28, 2011: Electrical Engineering, UC Santa Cruz
- Oct. 2, 2009: Applied Mathematics, UC Merced
- Apr. 28, 2008: CITRIS Health Technologies Tele-Seminar series (UC Santa Cruz, UC Merced)
- Mar. 12, 2008: Psychology Department, UC Santa Cruz
- Feb. 22, 2008: Electrical Engineering & Computer Science, UC Merced
- Mar. 15, 2007: School of Engineering, UC Merced
- Feb. 27, 2007: School of Computing and Information Sciences, Florida International University
- Dec. 8, 2006: NIPS 2006 workshop on *Continuous Attractor Neural Networks*
- Aug. 21, 2006: School of Informatics, University of Edinburgh, UK
- Jul. 5, 2006: Toyota Technological Institute at Chicago (TTI–C)
- Jun. 30, 2006: Dept. of Computer Science, Carnegie-Mellon University
- Jun. 23, 2006: Siemens Medical Solutions, US (Malvern, PA)
- Dec. 5, 2005: workshop on *Contrastive Divergence Learning*, Canadian Institute for Advanced Research, Vancouver
- Nov. 1, 2005: Neurological Sciences Institute (NSI), OHSU
- Feb. 15, 2005: Center for Spoken Language Understanding (CSLU), OGI/OHSU
- Jan. 19, 2005: Dept. of Computer Science & Electrical Engineering, OGI/OHSU
- Aug. 19, 2004: Dept. of Computer Science, University of Toronto
- Jun. 6, 2004: IRIS Machine Learning Workshop, Ottawa, Canada
- Apr. 22, 2004: School of Informatics, University of Edinburgh, UK
- Mar. 24, 2004: Dept. of Mathematical and Computer Sciences, Colorado School of Mines
- Mar. 18, 2004: Dept. of Computer Science & Engineering, OGI, Oregon Health & Science University
- Feb. 26, 2004: Dept. of Computer Science, UC Davis
- Dec. 14, 2002: NIPS 2002 workshop on *Negative Results and Open Problems*
- Nov. 26, 2002: Dept. of Computer Science, University of Toronto
- Apr. 24, 2002: Gatsby Computational Neuroscience Unit, University College London
- May 25, 2001: Program in Neuroscience, University of Maryland, Baltimore
- Mar. 23, 2001: Dept. of Statistics, University of Virginia
- Mar. 13, 2001: Center for Language and Speech Processing (CLSP), Johns Hopkins University ▶
- Mar. 12, 2001: Dept. of Neuroscience, Georgetown University Medical Center
- Nov. 10, 1999: Mathematics Dept., King's College London

- Aug. 5, 1999: Sloan Center for Theoretical Neurobiology, UCSF
- Aug. 3, 1999: International Computer Science Institute (ICSI), Berkeley

July 20, 2023