Geometric Calculus meets Deep Learning

SXD
&
EUMS

UPC·BSC & **Jalisco.gob.mx**

16/07/2019

A

- On Brains and Minds (Readings)
- On Machine Learning (ML)
- On Neural Networks (NNs)
- On Deep Learning (DL)

B

- Complex and quaternionic NNs
- Geometric NNs
- Research opportunities?
- References

On Brains and Minds (Readings)

- *On intelligence: How a new understanding of the brain will lead to the creation of truly intelligent machines* (hawkins-blakeslee-2004 [49])
- *Hawkins on intelligence: fascination and frustration* (perlis-2005 [94])
- *Vision with direction* (bigun-2006 [15])
- *Pattern theory: The Stochastic Analysis of Real-World Signals* (mumford-desolneux-2010 [86])
- *Life 3.0: Being human in the age of artificial intelligence* (tegmark-2017 [109])
- *The AI Spring of 2018* (olhede-wolfe-2018 [89]). The author's discuss the implications as nations race for AI dominance)
- *Superhuman AI for multiplayer poker* (brown-sandholm-2019 [19])

On Machine Learning

Machine learning, a field of computer science, seeks to design machines that learn (*Mathematics of Machine Learning: An introduction*, Plenary lecture at the Rio ICM2018, arora-2018 [7])

In general terms, a rough idea of *machine learning* is to produce algorithms that output a function that

- **F** Gives good approximations of given values y^i for given inputs x^i $(i = 1, \ldots, N)$;
- **G** Has good *generalization capacity*, which means that for any x (of a kind similar to that of the x^i) the value $y' = f(x)$ is a good approximation of the expected value y corresponding to x.

There are many algorithms that have these properties, to some extent:

- *Understanding ML* (shalevshwartz-bendavid-2014 [103]):

  The aim [...] to introduce machine learning, and its algorithmic paradigms, in a principled way.

- arora-2018 [7]:

  Machine learning is related to **artificial intelligence**, *but somewhat distinct because it does not seek to recreate only human-like skills in a machine. Some skills [...] may be easy for a machine, but beyond the cognitive abilities of humans. Conversely, many human skills such as composing good music and proving math theorems seem beyond the reach of current machine learning paradigms.*
  The quest to imbue machines with learning abilities rests upon an emerging body of knowledge that spans computer science, mathematical optimization, statistics, applied math, applied physics etc.

Linear regression: Let w be an unknown vector of weights (one weight for each of the D components of the x vectors) and $f_w(x) = w \cdot x = w_1 x_1 + \cdots + x_D x_D$ (a *weighted sum*) the function to be learned.

A way of fulfilling condition **F** is to pick a w that achieves $\min_w \sum_{i=1}^{N} \left(w \cdot x^i - y^i\right)^2$ (*least squares* method).

Regularized linear regression (improves generalization capacity):

$\min_w \sum_i \left(w \cdot x^i - y^i\right)^2 + \lambda ||w||_2^2$

"where λ is a scalar ... discovered by experimenting with the data".

- *Mathematical foundations of supervised learning* (wolf-2018 [117])
- *Reinforcement learning: An introduction* (sutton-barto-2018 [106])
- *The Hundred-Page Machine Learning Book* (burkov-2019 [27])

On NNs and DL

The inspiration of convolutional layers came from cortical neurons within the visual cortex which only respond to stimuli in a receptive environment (shabbir-anwer-2018 [102])

NN blueprint:

$$\mathcal{N}: \quad \text{Input} \to L_1 \to L_2 \to \cdots \to L_m \to \text{Output}$$

Conventionally, the net is *deep* if $m > 1$.

- Functionally, a layer takes an input x and yields an output x'.
- The map $f : x \mapsto x'$ depends on parameters associated to the layer and whose nature depends on the kind of layer.
- The input to the first layer is the signal to be processed.
- The last layer is the *output layer*, and its output is the result produced by the net on the input signal.

- In general, x, x', and the *layer parameters* are multidimensional arrays whose nature is chosen according to the processing that has to be achieved.

Write $[n_1, n_2, \ldots, n_d]$ to denote the type of a d-dimensional (real) array with axis dimensions $n_1, \ldots, n_d$.

Thus $[n]$ is the type of n-dimensional vectors and $[n_1, n_2]$ the type of matrices with n_1 rows and n_2 columns. Matrices are useful to represent monochrome images, but for RGB images we need arrays of type $[n_1, n_2, 3]$, or $[n_1, n_2, n_3]$ if it is required that the image be represented by n_3 channels, as for example $n_3 = 6$ for *a pair of color stereoscopic images*.

The parameters associated to **convolutional** and **fully connected** layers are represented by an **array of weights**, W, and a bias array, b. In these cases, the expression of f has the form

$$f^{\pi}(x) = g\left(x \star_{\pi} W + b\right) \tag{1}$$

where $\star_{\pi}$ is a pairing specific of the layer and g is an activation function (usually $\mathrm{ReLU}(t) = \max(t - \beta, 0)$), that is applied component-wise to arrays.

For convolutional layers, $\star_{\pi} = \star$ is *array cross-correlation*, while for fully connected layers, $\star_{\pi}$ is *matrix product*, which is denoted by juxtoposition of its factors, xW.

- Given weight arrays and biases W_k and b_k ($k = 1, \dots m$), the net $\mathcal{N}$ computes a function $f = f_{W_1,b_1,\dots,W_m,b_m}$ that is *continuous* and *pice-wise affine*.

- There exist *training algorithms* of $\mathcal{N}$, particularly those of *back-propagation* type, achieving trained weights and biases for which f is 'optimal' in the sense of the conditions **F** and **G**.

- *Convolutional neural networks for images, speech and time series* (lecunn-bengio-1995 [71])
- *ImageNet classification with deep convolutional neural networks* (krizhevsky-sutskever-hinton-2012 [65])
- *Deep learning* (lecunn-bengio-hinton-2015 [72])
- *Neural networks and deep learning* (nielsen-2015 [88])
- *Deep learning in neural networks: an overview* (schmidhuber-2015 [101]; 54 pages of references)
- *Deep learning tutorial* (lisalab-2015 [67], Theano team)
- *Deep learning* (goodfellow-bengio-courville-2016 [47])

- *Understanging deep convolutional networks* (mallat-2016 [78])
- *Deep learning with Python* (brownlee-2017 [20])
- *Deep reinforcement learning for robotic manipulation–the state of the art* (amarjyoti-2017 [3])
- *Mathematics of deep learning* (vidal-bruna-giryes-soatto-2017 [115])
- *Universality of deep convolutional neural networks* (zhou-2019 [121])

Complex and quaternionic NNs

- *Quaternionic neural networks: Fundamental properties and applications* (isokawa-matsui-nishimura-2009 [61])
- *Complex-valued neural networks (second edition)* (hirose-2012 [57])
- *Complex-valued neural networks: Advances and applications* (hirose-2013 [58])
- *A mathematical motivation for complex-valued convolutional networks* (bruna-chintala-lecun-piantino-szlam-tygert-2015 [21])
- *Deep quaternion networks* (gaudet-maida-2018 [45])
- *Quaternion convolutional neural networks for end-to-end automatic speech recognition* (parcollet-et-6-2018 [92])
- *Quaternion convolutional neural networks* (zhu-xu-xu-chen-2018 [122])

Geometric NNs

- *Clifford Wavelets, Singular Integrals, and Hardy Spaces* (mitrea-1994 [82])

- *Geometric computing with Clifford algebras: theoretical foundations and applications in computer vision and robotics* (sommer-2001 [104], editor). Particularly

  Ch. 12: *Introduction to neural computation in Clifford algebras* (S. Buchholz and G. Sommer)

  Ch. 13: *Clifford algebra multilayer perceptrons* (S. Buchholz and G. Sommer)

- *A theory of neural computation with Clifford algebras* (buchholz-2005 [23], PhD thesis)

- *On Clifford neurons and Clifford multi-layer perceptrons* (buchholz-sommer-2008 [25]: "The paper provides a sound theoretical basis to Clifford neural computation")

- *Geometric neurocomputing* (bayro-2018 [10], Ch 13)

  the potential of geometric neural networks [...] for a variety of real applications using multidimensional representations, such as in graphics, augmented reality, machine learning, computer vision, medical image processing, and robotics.

Research opportunities?

- *Pattern theory: the stochastic analysis of real-world signals* (mumford-desolneux-2010)
- *The master algorithm* (domingos-2015 [40])

"Even more astonishing than the breadth of applications of machine learning is that it's the same algorithms doing all of these different things. Outside of machine learning, if you have two different problems to solve, you need to write two different programs. They might use some of the same infrastructure, like the same programming language or the same database system, but a program to, say, play chess is of no use if you want to process credit-card applications. In machine learning, the same algorithm can do both, provided you give it the appropriate data to learn from. In fact, just a few algorithms are responsible for the great majority of machine-learning applications, and we'll take a look at them in the next few chapters".

- *Neural networks and deep learning* (nielsen-2015 [88]): "Is there a simple algorithm for intelligence?"

"science contains many more such examples [besides astronomy]. Consider the myriad chemical substances making up our world, so beautifully explained by Mendeleev's periodic table, which is, in turn, explained by a few simple rules which may be obtained from quantum mechanics".

"My own prejudice is in favour of there being a simple algorithm for intelligence. And the main reason I like the idea, above and beyond the (inconclusive) arguments above, is that it's an optimistic idea. When it comes to research, an unjustified optimism is often more productive than a seemingly better justified pessimism, for an optimist has the courage to set out and try new things. That's the path to discovery, even if what is discovered is perhaps not what was originally hoped. A pessimist may be more "correct" in some narrow sense, but will discover less than the optimist".

- *Group equivariant convolutional networks* (cohen-welling-2016 [33])
- *Dynamic routing between capsules* (sabour-frosst-hinton-2017 [100])
- *Matrix capsules with EM routing* (hinton-sabour-frosst-2018 [56]): work out these remarkable schemes for GNNs.
- *Neural ordinary differential equations* (chen-rubanova-bettencourt-duvenaud-2018 [31]). Can these view of differential equations provide insights of GNNs?
- *Toward an AI Physicist for Unsupervised Learning* (wu-tegmark-2018 [118])
- *Learning algebraic structures: preliminary investigations* (he-kim-2019 [51])

- *Feature extraction using conformal geometric algebra for AdaBoost algorithm based inplane rotated face detection* (pham-doan-hitzer-2019 [96])
- *Geometric Algebra, Gravity and Gravitational Waves* (lasenby-2019 [69])
- *Maxwell's equations are universal for locally conserved quantities* (burns-2019 [28])
- *A 1d up approach to Conformal Geomet- ric Algebra: applications in line fitting and quantum mechanics* (lasenby-2019-1up [68])
- *Using Raising and Lowering Operators from Geometric Algebra for Electroweak Theory in Particle Physics* (mcclellan-2019 [80])

[1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al.

Tensorflow: Large-scale machine learning on heterogeneous distributed systems.

arXiv preprint arXiv:1603.04467, 2016.

[2] Gerardo Altamirano-Gomez and Eduardo Bayro-Corrochano.

Conformal geometric algebra method for detection of geometric primitives.

In *2016 23rd International Conference on Pattern Recognition (ICPR)*, pages 4190–4195. IEEE, 2016.

[3] Smruti Amarjyoti.

Deep reinforcement learning for robotic manipulation–the state of the art.

arXiv: 1701. 08878, 2017.

[4] Michael A Arbib and James J Bonaiuto (editors).

From Neuron to Cognition via Computational Neuroscience.

MIT Press, 2016.

[5] Michael A. Arbib and Allen R. Hancon (editors).

Vision, brain, and cooperative computation.

MIT Press, 1990.

in four parts: Visual neurophysiology, Visual psicophysics, Machine vision and robotics, and Connectionism and cooperative computation.

[6] Paolo Arena, Luigi Fortuna, Giovanni Muscato, and Maria Gabriella Xibilia.

Neural networks in multidimensional domains: fundamentals and new trends in modelling and control, volume 234 of *LNCIS*.

Springer, 1998.

[7] Sanjeev Arora.

Mathematics of Machine Learning: An introduction, 2018.

[8] Will Ballard.

Hands-on deep learning for images with tensorflow.

Packt, 2018.

Build intelligent computer vision applications using TensorFlow and Keras.

[9] Eduardo Bayro-Corrochano.
Geometric computing: for wavelet transforms, robot vision, learning, control and action.
Springer, 2010.

[10] Eduardo Bayro-Corrochano.
Geometric Algebra Applications Vol. I: Computer Vision, Graphics and Neurocomputing.
Springer, 2018.

[11] Eduardo Bayro-Corrochano and Gerik Scheuermann.
Geometric algebra computing: in engineering and computer science.
Springer Science & Business Media, 2010.

[12] Eduardo Bayro-Corrochano, Eduardo Vazquez-Santacruz, Eduardo Moya-Sanchez, and Efrain Castillo-Muñis.

Geometric bioinspired networks for recognition of 2-D and 3-D low-level structures and transformations.

IEEE transactions on neural networks and learning systems, 27(10):2020–2034, 2016.

[13] Mauricio C. L. Belón.

Applications of conformal geometric algebra in mesh deformation, 2013.

XXVI Conference on Graphics, Patterns and Images (IEEE), 39-46 http://sibgrapi.sid.inpe.br/attachment.cgi/sid.inpe.br/sibgrapi/2013/07.04.19.56/doc/PID2845377.pdf.

[14] Fernando Berzal.
Redes neuronales & deep learning.
Author, 2019.

[15] Josef Bigun.
Vision with direction.
Springer, 2006.

[16] Nick Bostrom.
Superintelligence: paths, dangers, strategies.
Oxford University Press, 2014.

[17] Paul Breiding, Sara Kališnik, Bernd Sturmfels, and Madeleine Weinstein.

Learning algebraic varieties from samples.

Revista Matemática Complutense, 31(3):545–593, 2018.

[18] Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst.

Geometric deep learning: going beyond Euclidean data.

IEEE Signal Processing Magazine, 34(4):18–42, 2017.

[19] Noam Brown and Tuomas Sandholm.

Superhuman AI for multiplayer poker.

Science, 10(1126):13 pages, 2019.

[20] Jason Brownlee.

Deep learning with python.

Machine Learning Mastery, 2017.

Develop deep learning models on theano and tensorflow using keras.

[21] Joan Bruna, Soumith Chintala, Yann LeCun, Serkan Piantino, Arthur Szlam, and Mark Tygert.

A mathematical motivation for complex-valued convolutional networks.

arXiv: 1503. 03438, 2015.

[22] Joan Bruna and Stéphane Mallat.

Invariant scattering convolution networks.

IEEE transactions on pattern analysis and machine intelligence, 35(8):1872–1886, 2013.

[23] Sven Buchholz.

A theory of neural computation with Clifford algebras.

PhD thesis, Christian-Albrechts Universität Kiel, 2005.

[24] Sven Buchholz, Eckhard Hitzer, and Kanta Tachibana.

Coordinate independent update formulas for versor Clifford neurons, 2008.

SCIS & ISIS 2008, Japan Society for Fuzzy Theory and Intelligent Informatics, 814-819.

[25] Sven Buchholz and Gerald Sommer.

On Clifford neurons and Clifford multi-layer perceptrons.

Neural Networks, 21(7):925–935, 2008.

[26] Sven Buchholz, Kanta Tachibana, and Eckhard Hitzer.
Optimal learning rates for Clifford neurons, 2007.
International conference on artificial neural networks (Springer), 864-873.

[27] Andriy Burkov.
The Hundred-Page Machine Learning Book.
Andriy Burkov, 2019.

[28] Luke Burns.
Maxwell's equations are universal for locally conserved quantities, 2019.
To appear in AACA.

[29] Wai Lam Chan, Hyeokho Choi, and Richard Baraniuk.
Quaternion wavelets for image analysis and processing.
In *IEEE International Conference on Image Processing*, volume 5, pages 3057–3060, 2004.

[30] Adam S. Charles.
Interpreting deep learning: The machine learning rorschach, 2018.
`https://arxiv.org/pdf/1806.00148.pdf`.

[31] Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud.
Neural ordinary differential equations.
In *Advances in neural information processing systems*, pages 6571–6583, 2018.

[32] Taco Cohen, Mario Geiger, Jonas Köhler, and Max Welling.
Convolutional networks for spherical signals.
arXiv: 1709. 04893, 2017.

[33] Taco Cohen and Max Welling.
Group equivariant convolutional networks.
In *International conference on machine learning*, pages 2990–2999, 2016.

[34] John J. Craig.
Introduction to robotics: Mechanics and control (third edition).
Always learning. Pearson, 2014.

[35] Yuwei Cui, Subutai Ahmad, and Jeff Hawkins.

The HTM spatial pooler–A neocortical algorithm for online sparse distributed coding.

Frontiers in computational neuroscience, 11:111, 2017.

[36] Celia del Buey de Andrés.

Todo es más sencillo con los hipercomplejos.

Gaceta de la Real Sociedad Matematica Española, 22(1):145–157, 2019.

[37] Amit Deshpande, Sandesh Kamath, and KV Subrahmanyam.

Robustness and equivariance of neural networks, 2018.

[38] Sander Dieleman, Jeffrey De Fauw, and Koray Kavukcuoglu.

Exploiting cyclic symmetry in convolutional neural networks.

arXiv: 1602. 02660, 2016.

[39] Samuel Dodge and Lina Karam.

Understanding how image quality affects deep neural networks.

In *Quality of Multimedia Experience (QoMEX), 2016 Eighth International Conference on*, pages 1–6. IEEE, 2016.

[40] Pedro Domingos.

The master algorithm.

Basic Books, 2015.

How the quest for the ultimate learning machine will remake our world.

[41] Carlos Esteves, Christine Allen-Blanchette, Ameesh Makadia, and Kostas Daniilidis.

Learning *so(3)* equivariant representations with spherical CNNS.

In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 52–68, 2018.

[42] Carlos Esteves, Avneesh Sud, Zhengyi Luo, Kostas Daniilidis, and Ameesh Makadia.

Cross-Domain 3D Equivariant Image Embeddings.

arXiv: 1812. 02716 , v2, 2019.

[43] Michael Felsberg.

Low-level image processing with the structure multivector, 2002.

PhD thesis.

[44] H Firdaus and T Ogawa.

Application of Quaternion Neural Network to EMG-Based Estimation of Forearm Motion.

In *Proceedings on the International Conference on Artificial Intelligence (ICAI'17)*, pages 317–318. The Steering Committee of The World Congress in Computer Science, 2017.

[45] Chase J Gaudet and Anthony S Maida.

Deep quaternion networks.

In *2018 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2018.

[46] Aurélien Géron.
Hands-on machine learning with Scikit-Learn and TensorFlow: Concepts, tools, and techniques to build intelligent systems.
O'Reilly, 2017.

[47] Ian Goodfellow, Yoshua Bengio, and Aaron Courville.
Deep learning.
MIT press, 2016.

[48] Gösta H Granlund and Hans Knutsson.
Signal processing for computer vision.
Springer, 1995.

[49] Jeff Hawkins and Sandra Blakeslee.
On intelligence: How a new understanding of the brain will lead to the creation of truly intelligent machines.
Macmillan, 2007.

[50] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun.
Delving deep into rectifiers: Surpassing human-level performance on imagenet classification.
In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.

[51] Yang-Hui He and Minhyong Kim.
Learning algebraic structures: preliminairy investigations, 2019.
arXiv:1905.02263.

[52] Mikael Henaff, Joan Bruna, and Yann LeCun.
Deep convolutional networks on graph-structured data, 2015.
arXiv:1506.05163.

[53] David Hestenes.
Maxwell–Dirac electron theory, 2019.
Preprint 07/2/2019, 20 pages.

[54] David Hestenes and Garrett Sobczyk.
Clifford Algebra to Geometric Calculus.
Reidel, 1984.

[55] Geoffrey E. Hinton, Alex Krizhevsky, and Sida D. Wang.
Transforming auto-encoders, 2011.
International Conference on Artificial Neural Networks (Springer, 2011), 44-51.

[56] Geoffrey E Hinton, Sara Sabour, and Nicholas Frosst.
Matrix capsules with EM routing.
In *ICLR 2018*. Springer, 2018.
15 pages.

[57] Akira Hirose.
Complex-valued neural networks (second edition).
Springer, 2012.
Japanese edition 2004, first English edition 2006.

[58] Akira (editor) Hirose.
Complex-valued neural networks: Advances and applications.
IEEE Computational Intelligence. John Wiley & Sons, 2013.

[59] Eckhard Hitzer and Stephen J Sangwine.
Quaternion and Clifford Fourier transforms and wavelets.
Springer, 2013.

[60] Lester Ingber.
Computational algorithms derived from multiple scales of neocortical processing.
Cognitive Computation, 4(1):38–50, 2012.

[61] Teijiro Isokawa, Nobuyuki Matsui, and Haruhiko Nishimura.
Quaternionic neural networks: Fundamental properties and applications.
In *Complex-valued neural networks: utilizing high-dimensional parameters*, pages 411–439. IGI Global, 2009.

[62] David C. Knill and Whitman Richards (editors).
Perception as Bayesian inference.
Cambridge University Press, 1996.
Paper #1 is *Pattern theory: a unifying perspective*, by D. Mumford.

[63] Masaki Kobayashi.
Rotational invariance of quaternionic Hopfield neural networks.
IEEJ Transactions on Electrical and Electronic Engineering, 11(4):516–520, 2016.

[64] Jens Kober, J Andrew Bagnell, and Jan Peters.
Reinforcement learning in robotics: A survey.
The International Journal of Robotics Research, 32(11):1238–1274, 2013.

[65] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton.
ImageNet classification with deep convolutional neural networks.
In *Proceedings of the Conference Neural Information Processing Systems (NIPS 2011)*, pages 1097–1105, 2012.

[66] Sanjeev R. Kulkarni and Gilbert Harman.
Statistical learning theory: a tutorial.
Wiley Interdisciplinary Reviews: Computational Statistics, 3(6):543–556, 2011.

[67] LISA lab.

Deep learning tutorial, 2015.

Release 0.1. Theano Development Team.

[68] Anthony Lasenby.

A 1*d* up approach to conformal geometric algebra: applications, 2019.

To appear in AACA.

[69] Anthony Lasenby.

Geometric algebra, gravity and gravitational waves.

Advances in Applied Clifford Algebras, 29(4):41 pages, 2019.

[70] Carlile Lavor, Sebastià Xambó-Descamps, and Isiah Zaplana.
A Geometric Algebra Invitation to Space-Time Physics, Robotics and Molecular Geometry.
SBMA/Springerbrief. Springer, 2018.

[71] Y Le Cunn and Y Bengio.
Convolutional neural networks for Images, Speech and Time Series.
AT&T Laboratories, pages 1–14, 1995.

[72] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton.
Deep learning.
nature, 521(7553):436, 2015.

[73] Karel Lenc and Andrea Vedaldi.

Understanding image representations by measuring their equivariance and equivalence.

In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 991–999, 2015.

[74] Ian Lenz.

Deep learning for robotics.

PhD thesis, Cornell University, 2016.

[75] Henry W. Lin, Max abd Tegmark, and David Rolnick.

Why does deep and cheap learning work so well?

Journal of Statistical Physics, 168(6):1223–1247, 2017.

arXiv:1608.08225, v4.

[76] Pertti Lounesto and Esko Latvamaa.
Conformal transformations and Clifford algebras.
Proceedings of the American Mathematical Society, 79(4):533–538, 1980.

[77] Kevin M Lynch and Frank C. Park.
Modern Robotics: Mechanics, Planning, and Control.
Authors, 2017.

[78] Stéphane Mallat.
Understanding deep convolutional networks.
Phil. Trans. R. Soc. A, 374(2065):20150203, 2016.

[79] Diego Marcos, Michele Volpi, Nikos Komodakis, and Devis Tuia.

Rotation equivariant vector field networks.

arXiv: 1612. 09346, 2016.

[80] Gene McClellan.

Using raising and lowering operators from geometric algebra for electroweak theory in particle physics, 2019.

To appear in AACA.

[81] Hrushikesh Mhaskar, Qianli Liao, and Tomaso Poggio.

Learning functions: when is deep better than shallow.

arXiv: 1603. 00988, 2016.

[82] Marius Mitrea.
Clifford wavelets, singular integrals, and Hardy spaces.
Springer, 1994.

[83] Shariq Mobin and Joan Bruna.
Voice conversion using convolutional neural networks.
arXiv: 1610. 08927, 2016.

[84] Jojo Moolayil, Moolayil, and Suresh John.
Learn Keras for Deep Neural Networks.
Springer, 2019.
A fast track approach to modern deep learning with python.

[85] E Ulises Moya-Sánchez and Eduardo Vázquez-Santacruz.

A geometric bio-inspired model for recognition of low-level structures.

In *International Conference on Artificial Neural Networks*, pages 429–436. Springer, 2011.

[86] David Mumford and Agnès Desolneux.

Pattern theory: the stochastic analysis of real-world signals.

A. K. Peters, 2010.

[87] Kevin P. Murphy.

Machine learning: A probabilistic perspective.

MIT Press, 2012.

[88] Michael A Nielsen.

Neural networks and deep learning, volume 25.

Determination press San Francisco, CA, USA, 2015.

[89] Sofia Olhede and Patrick Wolfe.

The ai spring of 2018, 2018.

`http://discovery.ucl.ac.uk/10051919/1/AIspringv5.pdf`, 3 pages. Authors discuss the implications as nations race for AI dominance.

[90] Michael Otte.

A survey of machine learning approaches to robotic path-planning.

International Journal of Robotics Research, 5(1):90–98, 2008.

[91] Michael Otte.

A survey of machine learning approaches to robotic path-planning.

Technical Report PhD Preliminary Exam, University of Colorado at Boulder, 2008.

[92] Titouan Parcollet, Ying Zhang, Mohamed Morchid, Chiheb Trabelsi, Georges Linarès, Renato De Mori, and Yoshua Bengio.

Quaternion convolutional neural networks for end-to-end automatic speech recognition.

arXiv: 1806. 07789, 2018.

[93] Ankit B Patel, Tan Nguyen, and Richard G Baraniuk.

A probabilistic theory of deep learning.

arXiv preprint arXiv:1504.00641, 2015.

[94] Donald Perlis.

Hawkins on intelligence: fascination and frustration.

Artificial Intelligence, 169(2):184–191, 2005.

[95] Minh Tuan Pham and Kanta Tachibana.

A conformal geometric algebra based clustering method and its applications.

Advances in Applied Clifford Algebras, 26(3):1013–1032, 2016.

[96] Tuan M. Pham, Danh C. Doan, and Eckhard Hitzer.

Feature extraction using conformal geometric algebra for AdaBoost algorithm based inplane rotated face detection.

Advances in Applied Clifford Algebras, 29(4):19 pages, 2019.

[97] Tao Qian, Mang I. Vai, and Yuesheng Xu (editors).

Wavelet analysis and applications.

Applied and numerical harmonic analysis. Birkhäuser, 2007.

Includes: *Clifford analysis and the continuous spherical wavelet transform* (P. Cerejeiras, M. Ferreira, U. Kähler).

[98] David Rolnick and Max Tegmark.

The power of deeper networks for expressing natural functions, 2018.

arXiv:1705.05502, v2.

[99] Javier Ruiz-del Solar, Patricio Loncomilla, and Naiomi Soto.

A survey on deep learning methods for robot vision.

arXiv: 1803. 10862, 2018.

[100] Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton.
Dynamic routing between capsules.
In *Advances in neural information processing systems*, pages 3856–3866. Springer, 2017.

[101] Jürgen Schmidhuber.
Deep learning in neural networks: An overview.
Neural networks, 61:85–117, 2015.

[102] Jahanzaib Shabbir and Tarique Anwer.
A survey of deep learning techniques for mobile robot applications.
arXiv: 1803. 07608, 2018.

[103] Shai Shalev-Shwartz and Shai Ben-David.
Understanding machine learning: From theory to algorithms.
Cambridge university press, 2014.

[104] Gerald Sommer (ed.).
Geometric computing with Clifford algebras: theoretical foundations and applications in computer vision and robotics.
Springer, 2001.

[105] GS Staples.
CliffMath: Clifford algebra computations in Mathematica, 2008-2018, 2019.

[106] Richard S Sutton and Andrew G Barto.

Reinforcement learning: An introduction.

MIT press, 2018.

[107] Lei Tai, Jingwei Zhang, Ming Liu, Joschka Boedecker, and Wolfram Burgard.

A survey of deep network solutions for learning control in robotics: From reinforcement to imitation.

arXiv: 1612. 07139, v4, 2018.

[108] Yichuan Tang.

Deep learning using linear support vector machines.

arXiv: 1306. 0239, v4, 2015.

[109] Max Tegmark.
Life 3.0: Being human in the age of artificial intelligence.
Knopf, 2017.

[110] Tijmen Tieleman.
Optimizing neural networks that generate images.
University of Toronto (Canada), 2014.
PhD thesis.

[111] Julian Togelius.
Playing Smart: On Games, Intelligence, and Artificial Intelligence.
MIT Press, 2018.

[112] Jerry R Van Aken.

A high-level model of neocortical feedback based on an event window segmentation algorithm.

arXiv: 1409.6023, 2014.

[113] Ivan Vasiliev, Daniel Slater, Gienmario Scapagna, Peter Roelants, and Valentino Zocca.

Python Deep Learning (second edition).

Packt, 2019.

Exploring deep learning techniques and neural network architectures with PyTorch, Keras, and TensorFlow.

[114] Jayme Vaz and Stephen Mann.

Paravectors and the geometry of 3D euclidean space.

Advances in Applied Clifford Algebras, 28(5):99–141, 2018.

[115] Rene Vidal, Joan Bruna, Raja Giryes, and Stefano Soatto.
Mathematics of deep learning.
arXiv: 1712. 04741, 2017.

[116] Yi-Qing Wang.
An analysis of the viola-jones face detection algorithm.
Image Processing On Line, 4:128–148, 2014.

[117] Michael M Wolf.
Mathematical foundations of supervised learning, 2018.

[118] Tailin Wu and Max Tegmark.
Toward an AI Physicist for Unsupervised Learning, 2018.
arXiv:1810.10525, v2.

[119] S. Xambó-Descamps.

From Leibniz' *characteristica geometrica* to contemporary geometric algebra.

Quaderns d'Història de l'Enginyeria, 16 (Special issue dedicated to commemorate LEIBNIZ (1646-1716)):103–134, 2018.

pdf.

[120] Sebastià Xambó-Descamps.

Real spinorial groups—a short mathematical introduction.

SBMA/Springerbrief. Springer, 2018.

[121] Ding-Xuan Zhou.

Universality of deep convolutional neural networks.

Applied and Computational Harmonic Analysis, 2019.

[122] Xuanyu Zhu, Yi Xu, Hongteng Xu, and Changjian Chen.
Quaternion convolutional neural networks.
In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 631–647, 2018.