

1. Record Nr.	UNINA9910741380503321
Autore	Downing Keith L
Titolo	Gradient expectations : structure, origins, and synthesis of predictive neural networks // Keith L. Downing
Pubbl/distr/stampa	Cambridge, MA : , : The MIT Press, , 2023
ISBN	0-262-37468-4 0-262-37467-6
Edizione	[1st ed.]
Descrizione fisica	1 online resource (280 pages)
Collana	The MIT Press
Disciplina	006.3/2
Soggetti	Neural networks (Computer science)
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Nota di contenuto	Intro -- Title Page -- Copyright Page -- Dedication -- Table of Contents -- Preface -- Acknowledgments -- 1. Introduction -- 1.1. Data from Predictions -- 1.2. Movement and Prediction -- 1.3. Adaptation and Emergence -- 1.3.1. Gradients and Emergence in Neural Networks -- 1.4. Overflowing Expectations -- 2. Conceptual Foundations of Prediction -- 2.1. Compare and Err -- 2.2. Guesses and Goals -- 2.3. Gradients -- 2.3.1. Gradients Rising -- 2.4. Sequences -- 2.5. Abstracting by Averaging -- 2.6. Control and Prediction -- 2.7. Predictive Coding -- 2.8. Tracking Marr's Tiers -- 3. Biological Foundations of Prediction -- 3.1. Gradient-Following Bacteria -- 3.2. Neural Motifs for Gradient Calculation -- 3.3. Birth of a PID Controller -- 3.3.1. Adaptive Control in the Cerebellum -- 3.4. Detectors and Generators -- 3.4.1. The Hippocampus -- 3.4.2. Conceptual Embedding in the Hippocampus -- 3.5. Gradients of Predictions in the Basal Ganglia -- 3.6. Procedural versus Declarative Prediction -- 3.7. Rampant Expectations -- 4. Neural Energy Networks -- 4.1. Energetic Basis of Learning and Prediction -- 4.2. Energy Landscapes and Gradients -- 4.3. The Boltzmann Machine -- 4.4. The Restricted Boltzmann Machine (RBM) -- 4.5. Free Energy -- 4.5.1. Variational Free Energy -- 4.6. The Helmholtz Machine -- 4.7. The Free Energy Principle -- 4.8. Getting a Grip -- 5. Predictive Coding -- 5.1. Information Theory and Perception -- 5.2. Predictive Coding on High -- 5.2.1. Learning Proper Predictions -- 5.3. Predictive Coding for Machine

Learning -- 5.3.1. The Backpropagation Algorithm -- 5.3.2. Backpropagation via Predictive Coding -- 5.4. In Theory -- 6. Emergence of Predictive Networks -- 6.1. Facilitated Variation -- 6.2. Origins of Sensorimotor Activity -- 6.2.1. Origins of Oscillations -- 6.2.2. Activity Regulation in the Brain. 6.2.3. Competition and Cooperation in Brain Development -- 6.2.4. Layers and Modules -- 6.2.5. Running through the Woods on an Icy Evening -- 6.2.6. Oscillations and Learning -- 6.3. A Brief Evolutionary History of the Predictive Brain -- 7. Evolving Artificial Predictive Networks -- 7.1. I'm a Doctor, Not a Connectionist -- 7.2. Evolving Artificial Neural Networks (EANNs) -- 7.2.1. Reconciling EANNs with Deep Learning -- 7.3. Evolving Predictive Coding Networks -- 7.3.1. Preserving Backpropagation in a Local Form -- 7.3.2. Phylogenetic, Ontogenetic, and Epigenetic (POE) -- 7.4. Continuous Time Recurrent Neural Networks (CTRNNs) -- 7.4.1. Evolving Minimally Cognitive Agents -- 7.4.2. Cognitive Robots Using Predictive Coding -- 7.4.3. Toward More Emergent CTRNNs -- 7.5. Predictive POE Networks -- 7.5.1. Simulating Neural Selectionism and Constructivism -- 7.5.2. Predictive Constructivism -- 7.5.3. The D'Arcy Model -- 7.5.4. Neurites to Neurons in D'Arcy -- 7.5.5. Peripherals in D'Arcy -- 7.5.6. Neuromodulators in D'Arcy -- 7.5.7. Predictively Unpredictable -- 7.6. Most Useful and Excellent Designs -- 8. Conclusion -- 8.1. Schrodinger's Frozen Duck -- 8.2. Expectations Great and Small -- 8.3. As Expected -- 8.4. Gradient Expectations -- 8.5. Expecting the Unexpected -- References -- Index.

Sommario/riassunto

An insightful investigation into the mechanisms underlying the predictive functions of neural networks--and their ability to chart a new path for AI. Prediction is a cognitive advantage like few others, inherently linked to our ability to survive and thrive. Our brains are awash in signals that embody prediction. Can we extend this capability more explicitly into synthetic neural networks to improve the function of AI and enhance its place in our world Gradient Expectations is a bold effort by Keith L. Downing to map the origins and anatomy of natural and artificial neural networks to explore how, when designed as predictive modules, their components might serve as the basis for the simulated evolution of advanced neural network systems. Downing delves into the known neural architecture of the mammalian brain to illuminate the structure of predictive networks and determine more precisely how the ability to predict might have evolved from more primitive neural circuits. He then surveys past and present computational neural models that leverage predictive mechanisms with biological plausibility, identifying elements, such as gradients, that natural and artificial networks share. Behind well-founded predictions lie gradients, Downing finds, but of a different scope than those that belong to today's deep learning. Digging into the connections between predictions and gradients, and their manifestation in the brain and neural networks, is one compelling example of how Downing enriches both our understanding of such relationships and their role in strengthening AI tools. Synthesizing critical research in neuroscience, cognitive science, and connectionism, Gradient Expectations offers unique depth and breadth of perspective on predictive neural-network models, including a grasp of predictive neural circuits that enables the integration of computational models of prediction with evolutionary algorithms.
