Realm Retrieval Augmented Language Model Pre Training

Autonomous Horizons
Innovations in Machine and Deep Learning
Business Intelligence and Information Technology
The Psychology of the Language Learner
An Introduction to Machine Learning
Machine Learning Challenges
Advances in Information Retrieval
Foundation Models for Natural Language Processing
Artificial Intelligence in Healthcare
Intelligent Systems in Medicine and Health
Introduction to Information Retrieval
Neural Approaches to Conversational Information Retrieval
First Text Retrieval Conference (TREC-1)
DSLs in Boo
Loan Phonology

Representation Learning for Natural Language Processing
Open-Domain Question Answering
Artificial Intelligence in HCI
Discovering the Brain
Representation Learning for Natural Language Processing
Dynamic Reteaming
Machine Knowledge
Natural Language Processing and Chinese Computing
The Probabilistic Relevance Framework
Computer Vision
SysML Distilled
Talking Nets
Speech & Language Processing
Natural Language Processing and Chinese Computing
Augmented Mind
Personalized Learning
Recent Challenges in Intelligent Information and Database Systems
Pretrained Transformers for Text Ranking
Machine Learning for Text
Chinese Computational Linguistics

The Constitution of Algorithms
Question Answering for the Curated Web
The Road to General Intelligence
The NeurIPS '18 Competition

HEAVEN CAITLYN

Autonomous Horizons Springer Nature
A general-purpose language like C# is designed to handle all programming tasks. By contrast, the structure and syntax of a Domain-Specific Language are designed to match a particular applications area. A DSL is designed for readability and easy programming of repeating problems. Using the innovative Boo language, it's a breeze to create a DSL for your application domain that works on .NET and does not sacrifice performance. DSLs in Boo shows you how to design, extend, and evolve DSLs for .NET by focusing on approaches and patterns. You learn to define an app in terms that match the domain, and to use Boo to build DSLs that generate efficient executables. And you won't deal with the awkward XML-laden syntax many DSLs require. The book concentrates on writing internal (textual) DSLs that allow easy extensibility of the application and framework. And if you don't know Boo, don't worry-you'll learn right here all the

techniques you need. Purchase of the print book comes with an offer of a free PDF, ePub, and Kindle eBook from Manning. Also available is all code from the book.

Innovations in Machine and Deep Learning Academic Press

This book surveys fundamental concepts and practical methods for creating and curating large knowledge bases.

Business Intelligence and Information Technology Independently Published

This three-volume set constitutes the refereed proceedings of the 12th National CCF Conference on Natural Language Processing and Chinese Computing, NLPCC 2023, held in Foshan, China, during October 12–15, 2023. The ____ regular papers included in these proceedings were carefully reviewed and selected from 478 submissions. They were organized in topical sections as follows: dialogue systems; fundamentals of NLP; information extraction and knowledge graph; machine learning for NLP; machine translation and multilinguality; multimodality and explainability; NLP applications and text mining; question answering; large language models; summarization and generation; student workshop; and evaluation workshop.

The Psychology of the Language Learner Springer Nature

This book constitutes the refereed post-proceedings of the First PASCAL Machine Learning Challenges Workshop, MLCW 2005. 25 papers address three challenges: finding an assessment base on the uncertainty of predictions using

classical statistics, Bayesian inference, and statistical learning theory; second, recognizing objects from a number of visual object classes in realistic scenes; third, recognizing textual entailment addresses semantic analysis of language to form a generic framework for applied semantic inference in text understanding.

An Introduction to Machine Learning

Springer Nature

This book constitutes the refereed proceedings of the 2021 International Conference on Business Intelligence and Information Technology (BIIT 2021) held in Harbin, China, during December 18–20, 2021. BIIT 2021 is organized by the School of Computer and Information Engineering, Harbin University of Commerce, and supported by Scientific Research Group in Egypt (SRGE), Egypt. The papers cover current research in electronic commerce technology and application, business intelligence and decision making, digital economy, accounting informatization, intelligent information processing, image processing and multimedia technology, signal detection and processing, communication engineering and technology, information security, automatic control technique, data mining, software development, and design, blockchain technology, big data technology, artificial intelligence technology.

Machine Learning Challenges

Springer Nature

The Probabilistic Relevance Framework (PRF) is a formal framework for

document retrieval, grounded in work done in the 1970-80s, which led to the development of one of the most successful text-retrieval algorithms, BM25. In recent years, research in the PRF has yielded new retrieval models capable of taking into account structure and link-graph information. Again, this has led to one of the most successful web-search and corporate-search algorithms, BM25F. The Probabilistic Relevance Framework: BM25 and Beyond presents the PRF from a conceptual point of view, describing the probabilistic modelling assumptions behind the framework and the different ranking algorithms that result from its application: the binary independence model, relevance feedback models, BM25, BM25F. Besides presenting a full derivation of the PRF ranking algorithms, it provides many insights about document retrieval in general, and points to many open challenges in this area. It also discusses the relation between the PRF and other statistical models for IR, and covers some related topics, such as the use of non-textual features, and parameter optimization for models with free parameters. The Probabilistic Relevance Framework: BM25 and Beyond is self-contained and accessible to anyone with basic knowledge of probability and inference

Advances in Information Retrieval

Springer Nature

Artificial Intelligence (AI) in Healthcare is more than a comprehensive introduction to artificial intelligence as a tool in the generation and analysis of healthcare

data. The book is split into two sections where the first section describes the current healthcare challenges and the rise of AI in this arena. The ten following chapters are written by specialists in each area, covering the whole healthcare ecosystem. First, the AI applications in drug design and drug development are presented followed by its applications in the field of cancer diagnostics, treatment and medical imaging. Subsequently, the application of AI in medical devices and surgery are covered as well as remote patient monitoring. Finally, the book dives into the topics of security, privacy, information sharing, health insurances and legal aspects of AI in healthcare. Highlights different data techniques in healthcare data analysis, including machine learning and data mining Illustrates different applications and challenges across the design, implementation and management of intelligent systems and healthcare data networks Includes applications and case studies across all areas of AI in healthcare data

Foundation Models for Natural Language Processing Routledge

This textbook comprehensively covers the latest state-of-the-art methods and applications of artificial intelligence (AI) in medicine, placing these developments into a historical context. Factors that assist or hinder a particular technique to improve patient care from a cognitive informatics perspective are identified and relevant methods and clinical applications in areas including

translational bioinformatics and precision medicine are discussed. This approach enables the reader to attain an accurate understanding of the strengths and limitations of these emerging technologies and how they relate to the approaches and systems that preceded them. With topics covered including knowledge-based systems, clinical cognition, machine learning and natural language processing, Intelligent Systems in Medicine and Health: The Role of AI details a range of the latest AI tools and technologies within medicine. Suggested additional readings and review questions reinforce the key points covered and ensure readers can further develop their knowledge. This makes it an indispensable resource for all those seeking up-to-date information on the topic of AI in medicine, and one that provides a sound basis for the development of graduate and undergraduate course materials.

Artificial Intelligence in Healthcare Now Pub

This double volume book set constitutes the refereed proceedings of 4th International Conference, AI-HCI 2023, held as part of the 25th International Conference, HCI International 2023, which was held virtually in Copenhagen, Denmark in July 2023. The total of 1578 papers and 396 posters included in the HCII 2023 proceedings was carefully reviewed and selected from 7472 submissions. The first volume focuses on topics related to Human-Centered Artificial Intelligence, explainability, transparency and trustworthiness, ethics

and fairness, as well as AI-supported user experience design. The second volume focuses on topics related to AI for language, text, and speech-related tasks, human-AI collaboration, AI for decision-support and perception analysis, and innovations in AI-enabled systems.

Intelligent Systems in Medicine and Health Springer Nature

This book provides an overview of the recent advances in representation learning theory, algorithms, and applications for natural language processing (NLP), ranging from word embeddings to pre-trained language models. It is divided into four parts. Part I presents the representation learning techniques for multiple language entries, including words, sentences and documents, as well as pre-training techniques. Part II then introduces the related representation techniques to NLP, including graphs, cross-modal entries, and robustness. Part III then introduces the representation techniques for the knowledge that are closely related to NLP, including entity-based world knowledge, sememe-based linguistic knowledge, legal domain knowledge and biomedical domain knowledge. Lastly, Part IV discusses the remaining challenges and future research directions. The theories and algorithms of representation learning presented can also benefit other related domains such as machine learning, social network analysis, semantic Web, information retrieval, data mining and computational biology. This book is

intended for advanced undergraduate and graduate students, post-doctoral fellows, researchers, lecturers, and industrial engineers, as well as anyone interested in representation learning and natural language processing. As compared to the first edition, the second edition (1) provides a more detailed introduction to representation learning in Chapter 1; (2) adds four new chapters to introduce pre-trained language models, robust representation learning, legal knowledge representation learning and biomedical knowledge representation learning; (3) updates recent advances in representation learning in all chapters; and (4) corrects some errors in the first edition. The new contents will be approximately 50%+ compared to the first edition. This is an open access book.

Introduction to Information Retrieval
Springer Nature
Personalized Learning: A Guide for Engaging Students with Technology is designed to help educators make sense of the shifting landscape in modern education. While changes may pose significant challenges, they also offer countless opportunities to engage students in meaningful ways to improve their learning outcomes. Personalized learning is the key to engaging students, as teachers are leading the way toward making learning as relevant, rigorous, and meaningful inside school as outside and what kids do outside school: connecting and sharing online, and engaging in virtual communities of their own Renowned author of the Heck: Where the Bad Kids Go series, Dale

Basye, and award winning educator Peggy Grant, provide a go-to tool available to every teacher today—technology as a way to 'personalize' the education experience for every student, enabling students to learn at their various paces and in the way most appropriate to their learning styles.

Neural Approaches to Conversational Information Retrieval Packt Publishing Ltd

This open access book provides an overview of the recent advances in representation learning theory, algorithms and applications for natural language processing (NLP). It is divided into three parts. Part I presents the representation learning techniques for multiple language entries, including words, phrases, sentences and documents. Part II then introduces the representation techniques for those objects that are closely related to NLP, including entity-based world knowledge, sememe-based linguistic knowledge, networks, and cross-modal entries. Lastly, Part III provides open resource tools for representation learning techniques, and discusses the remaining challenges and future research directions. The theories and algorithms of representation learning presented can also benefit other related domains such as machine learning, social network analysis, semantic Web, information retrieval, data mining and computational biology. This book is intended for advanced undergraduate and graduate students, post-doctoral fellows,

researchers, lecturers, and industrial engineers, as well as anyone interested in representation learning and natural language processing.

First Text Retrieval Conference (TREC-1)

Springer Nature

Humans have always dreamed of automating laborious physical and intellectual tasks, but the latter has proved more elusive than naively suspected. Seven decades of systematic study of Artificial Intelligence have witnessed cycles of hubris and despair. The successful realization of General Intelligence (evidenced by the kind of cross-domain flexibility enjoyed by humans) will spawn an industry worth billions and transform the range of viable automation tasks.The recent notable successes of Machine Learning has lead to conjecture that it might be the appropriate technology for delivering General Intelligence. In this book, we argue that the framework of machine learning is fundamentally at odds with any reasonable notion of intelligence and that essential insights from previous decades of AI research are being forgotten. We claim that a fundamental change in perspective is required, mirroring that which took place in the philosophy of science in the mid 20th century. We propose a framework for General Intelligence, together with a reference architecture that emphasizes the need for anytime bounded rationality and a situated denotational semantics. We given necessary emphasis to compositional reasoning, with the required compositionality being provided

via principled symbolic-numeric inference mechanisms based on universal constructions from category theory.• Details the pragmatic requirements for real-world General Intelligence.• Describes how machine learning fails to meet these requirements.• Provides a philosophical basis for the proposed approach.• Provides mathematical detail for a reference architecture.• Describes a research program intended to address issues of concern in contemporary AI.The book includes an extensive bibliography, with ~400 entries covering the history of AI and many related areas of computer science and mathematics.The target audience is the entire gamut of Artificial Intelligence/Machine Learning researchers and industrial practitioners. There are a mixture of descriptive and rigorous sections, according to the nature of the topic. Undergraduate mathematics is in general sufficient. Familiarity with category theory is advantageous for a complete understanding of the more advanced sections, but these may be skipped by the reader who desires an overall picture of the essential concepts This is an open access book.

DSLs in Boo International Society for Technology in Education

Excitement and momentum in artificial intelligence (AI) and machine learning has been accelerating with a global AI race in progress. After spending a decade on the front lines of the AI revolution, the author discovered the one key ingredient that was missing

from mainstream Al research - humans! This book explains how augmenting humans, combining human intuition and artificial intelligence, will herald an unprecedented era of productivity and financial success. The coming wave of human centered Al will help us solve the biggest problems facing humanity while also protecting us from rogue or weaponized Al systems. It includes a framework for creating hybrid solutions combining Al, machine learning, and human intuition that can, predict the future, improve our social lives, eliminate scarcity, and provide a clear roadmap to abundance and prosperity in the financial, health, and relationship industries that represents trillions in economic opportunity over the next decade.

Loan Phonology National Academies Press

Research results over the past decades have consistently demonstrated that a key reason why many second language learners fail--while some learners do better with less effort--lies in various learner attributes such as personality traits, motivation, or language aptitude. In psychology, these attributes have traditionally been called "individual differences." The scope of individual learner differences is broad--ranging from creativity to learner styles and anxiety--yet there is no current, comprehensive, and unified volume that provides an overview of the considerable amount of research conducted on various language learner differences, until now. Each chapter in this new

volume focuses on a different individual difference variable. Besides a review of the relevant second language literature, Zoltán Dörnyei presents a concise overview of the psychological research involving each topic. A key concern for the author has been to define the various learner factors as measurable constructs and therefore the discussion includes a summary of the most famous tests and questionnaires in each domain. A wide range of readers will benefit from this book--students in linguistics, applied linguistics, modern languages, and psychology programs; second language teachers participating in in-service training courses; and researchers in second language acquisition and psychology.

Representation Learning for Natural Language Processing "O'Reilly Media, Inc."

Held in Gaithersburg, MD, Nov. 4-6, 1992. Evaluates new technologies in information retrieval. Numerous graphs, tables and charts.

Open-Domain Question Answering John Benjamins Publishing

Your team will change whether you like it or not. People will come and go. Your company might double in size or even be acquired. In this practical book, author Heidi Helfand shares techniques for reteaming effectively. Engineering leaders will learn how to catalyze team change to reduce the risk of attrition, learning and career stagnation, and the development of knowledge silos. Based on research into well-known software companies, the patterns in this book

help CTOs and team managers effectively integrate new hires into an existing team, manage a team that has lost members, or deal with unexpected change. You'll learn how to isolate teams for focused innovation, rotate team members for knowledge sharing, break through organizational apathy, and more. You'll explore: Real-world examples that demonstrate why and how organizations reteam Five reteaming patterns: One by One, Grow and Split, Isolation, Merging, and Switching Tactics to help you master dynamic reteaming in your company Stories that demonstrate problems caused by reteaming anti-patterns

Artificial Intelligence in HCI

Cambridge University Press

This open access book provides a comprehensive overview of the state of the art in research and applications of Foundation Models and is intended for readers familiar with basic Natural Language Processing (NLP) concepts. Over the recent years, a revolutionary new paradigm has been developed for training models for NLP. These models are first pre-trained on large collections of text documents to acquire general syntactic knowledge and semantic information. Then, they are fine-tuned for specific tasks, which they can often solve with superhuman accuracy. When the models are large enough, they can be instructed by prompts to solve new tasks without any fine-tuning. Moreover, they can be applied to a wide range of different media and problem domains, ranging from image and video

processing to robot control learning. Because they provide a blueprint for solving many tasks in artificial intelligence, they have been called Foundation Models. After a brief introduction to basic NLP models the main pre-trained language models BERT, GPT and sequence-to-sequence transformer are described, as well as the concepts of self-attention and context-sensitive embedding. Then, different approaches to improving these models are discussed, such as expanding the pre-training criteria, increasing the length of input texts, or including extra knowledge. An overview of the best-performing models for about twenty application areas is then presented, e.g., question answering, translation, story generation, dialog systems, generating images from text, etc. For each application area, the strengths and weaknesses of current models are discussed, and an outlook on further developments is given. In addition, links are provided to freely available program code. A concluding chapter summarizes the economic opportunities, mitigation of risks, and potential developments of AI.

Discovering the Brain MIT Press

Surprising tales from the scientists who first learned how to use computers to understand the workings of the human brain. Since World War II, a group of scientists has been attempting to understand the human nervous system and to build computer systems that emulate the brain's abilities. Many of the early workers in this field of neural networks came from cybernetics; others

came from neuroscience, physics, electrical engineering, mathematics, psychology, even economics. In this collection of interviews, those who helped to shape the field share their childhood memories, their influences, how they became interested in neural networks, and what they see as its future. The subjects tell stories that have been told, referred to, whispered about, and imagined throughout the history of the field. Together, the interviews form a Rashomon-like web of reality. Some of the mythic people responsible for the foundations of modern brain theory and cybernetics, such as Norbert Wiener, Warren McCulloch, and Frank Rosenblatt, appear prominently in the recollections. The interviewees agree about some things and disagree about more. Together, they tell the story of how science is actually done, including the false starts, and the Darwinian struggle for jobs, resources, and reputation. Although some of the interviews contain technical material, there is no actual mathematics in the book. Contributors James A. Anderson, Michael Arbib, Gail Carpenter, Leon Cooper, Jack Cowan, Walter Freeman, Stephen Grossberg, Robert Hecht-Neilsen, Geoffrey Hinton, Teuvo Kohonen, Bart Kosko, Jerome Lettvin, Carver Mead, David Rumelhart, Terry Sejnowski, Paul Werbos, Bernard Widrow

Representation Learning for Natural Language Processing Springer

In recent years, significant progress has been made in achieving artificial intelligence (AI) with an impact on

students, managers, scientists, health personnel, technical roles, investors, teachers, and leaders. This book presents numerous successful applications of AI in various contexts. The innovative implications covered fall under the general field of machine learning (ML), including deep learning, decision-making, forecasting, pattern recognition, information retrieval, and interpretable AI. Decision-makers and entrepreneurs will find numerous successful applications in health care, sustainability, risk management, human activity recognition, logistics, and Industry 4.0. This book is an essential resource for anyone interested in challenges, opportunities, and the latest developments and real-world applications of ML. Whether you are a student, researcher, practitioner, or simply curious about AI, this book provides valuable insights and inspiration for your work and learning.