


Multimodal Knowledge Systems

Construction and Reasoning

This book focuses on advancing the integration of multimodal data (text, images, and structured knowledge) to enable precise knowledge extraction and human-like reasoning. The book's primary objective is to address critical challenges such as modality gaps, semantic misalignment, dataset biases, and static reasoning paradigms. By introducing novel frameworks that unify graph-based learning, hierarchical representation, bias mitigation, and iterative refinement, this book provides systematic solutions to build robust, interpretable, and scalable AI systems. This book addresses gaps caused by incomplete textual semantics, spurious correlations across modalities, and inflexible reasoning pipelines by offering three pivotal contributions. First, the authors offer theoretical innovations in graph alignment techniques, hierarchical learning paradigms, and multi-agent reasoning frameworks. Then, the book goes on to offer practical tools including benchmark datasets, reproducible methodologies, and applications validated on state-of-the-art tasks. Finally, the book offers a broader impact through solutions tailored for low-resource settings, ethical considerations in AI deployment, and integration with emerging technologies like large foundation models. By bridging the divide between theoretical advancements and real-world applicability, the book serves as an essential resource for researchers and practitioners aiming to leverage multimodal data effectively, ethically, and at scale.

This book focuses on advancing the integration of multimodal data (text, images, and structured knowledge) to enable precise knowledge extraction and human-like reasoning. The book's primary objective is to address critical challenges such as modality gaps, semantic misalignment, dataset biases, and static reasoning paradigms. By introducing novel frameworks that unify graph-based learning, hierarchical representation, bias mitigation, and iterative refinement, this book provides systematic solutions to build robust, interpretable, and scalable AI systems. This book addresses gaps caused by incomplete textual semantics, spurious correlations across modalities, and inflexible reasoning pipelines by offering three pivotal contributions. First, the authors offer theoretical innovations in graph alignment techniques, hierarchical learning paradigms, and multi-agent reasoning frameworks. Then, the book goes on to offer practical tools including benchmark datasets, reproducible methodologies, and applications validated on state-of-the-art tasks. Finally, the book offers a broader impact through solutions tailored for low-resource settings, ethical considerations in AI deployment, and integration with emerging technologies like large foundation models. By bridging the divide between theoretical advancements and real-world applicability, the book serves as an essential resource for researchers and practitioners aiming to leverage multimodal data effectively, ethically, and at scale.



160,49 €

149,99 € (zzgl. MwSt.)

Lieferfrist: bis zu 10 Tage

Artikelnummer: 9783032187604

Medium: Buch

ISBN: 978-3-032-18760-4

Verlag: Springer

Erscheinungstermin: 06.06.2026

Sprache(n): Englisch

Auflage: Erscheinungsjahr 2026

Serie: Wireless Networks

Produktform: Gebunden

Gewicht: 539 g

Seiten: 230

Format (B x H): 160 x 241 mm

