Koonce

Verified Deep Learning with Lean4

Formal Backpropagation from MLP to Attention, via MLIR

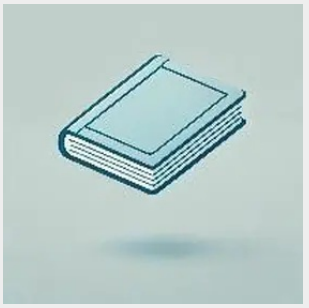

64,19 €
59,99 € (zzgl. MwSt.)

vorbestellbar, Erscheinungstermin ca. November 2027

Artikelnummer: 9798868836473
Medium: Buch
ISBN: 979-8-8688-3647-3
Verlag: APRESS L.P.
Erscheinungstermin: 01.12.2027
Sprache(n): Englisch
Auflage: 1. Auflage 2027
Produktform: Kartoniert
Format (B x H): 178 x 254 mm

This book teaches you how backpropagation actually works. You have called '.backward()' a thousand times. Here you open the hood and check the math, layer by layer.Backpropagation is one operation, the vector-Jacobian product, composed by three rules: the chain rule, additive fan-in, and multiplicative fan-in. Five structural tricks handle the layers whose Jacobians look dense: elementwise activations, convolution, max pooling, the normalizations and softmax, and the weight gradient of a dense layer. Every architecture in this book is a combination of those eight things. So is every architecture in its catalog of forty-six more, from U-Net to Mamba to Stable Diffusion. There is no ninth.Each chapter adds one primitive and climbs one rung: a linear classifier on MNIST, then an MLP, a CNN, BatchNorm on CIFAR-10, ResNet-34, MobileNetV2, EfficientNet, ConvNeXt, and the Vision Transformer. For each layer, you derive the backward pass by hand, state it as a theorem in Lean 4 and check it, read the MLIR it compiles to, and train the network on real data with one command. From ResNet onward, every chapter also trains its ImageNet recipe twice, once through a JAX reference and once through the verified trainer, and scores the two against each other.The gradients are theorems. The compiler and the GPU are trusted, and cross-checked against JAX. Three views of one computation: the math, the proof, and the code that runs. They agree, and this book is about why.What You Will Learn:- Derive the backward pass of every layer in the modern vision stack, from dense layers to self-attention- Read and check Lean 4 theorems that the gradients are correct, with no prior Lean- Follow a network from a fifteen-line spec through MLIR to the GPU, and know what each stage preserves- Reduce an unfamiliar architecture to known primitives and write its specification yourself- Know exactly what the proofs cover and what they leave to the training runWho It Is For:For readers who have trained a neural network and want to know what '.backward()' computed. It assumes derivatives, matrices, and a rough idea of convolution. It assumes no Lean, MLIR, compilers, or formal verification. Skip every Lean snippet and you still get a complete book on backpropagation.

This book teaches you how backpropagation actually works. You have called '.backward()' a thousand times. Here you open the hood and check the math, layer by layer.Backpropagation is one operation, the vector-Jacobian product, composed by three rules: the chain rule, additive fan-in, and multiplicative fan-in. Five structural tricks handle the layers whose Jacobians look dense: elementwise activations, convolution, max pooling, the normalizations and softmax, and the weight gradient of a dense layer. Every architecture in this book is a combination of those eight things. So is every architecture in its catalog of forty-six more, from U-Net to Mamba to Stable Diffusion. There is no ninth.Each chapter adds one primitive and climbs one rung: a linear classifier on MNIST, then an MLP, a CNN, BatchNorm on CIFAR-10, ResNet-34, MobileNetV2, EfficientNet, ConvNeXt, and the Vision Transformer. For each layer, you derive the backward pass by hand, state it as a theorem in Lean 4 and check it, read the MLIR it compiles to, and train the network on real data with one command. From ResNet onward, every chapter also trains its ImageNet recipe twice, once through a JAX reference and once through the verified trainer, and scores the two against each other.The gradients are theorems. The compiler and the GPU are trusted, and cross-checked against JAX. Three views of one computation: the math, the proof, and the code that runs. They agree, and this book is about why.What You Will Learn:- Derive the backward pass of every layer in the modern vision stack, from dense layers to self-attention- Read and check Lean 4 theorems that the gradients are correct, with no prior Lean- Follow a network from a fifteen-line spec through MLIR to the GPU, and know what each stage preserves- Reduce an unfamiliar architecture to known primitives and write its specification yourself- Know exactly what the proofs cover and what they leave to the training runWho It Is For:For readers who have trained a neural network and want to know what '.backward()' computed. It assumes derivatives, matrices, and a rough idea of convolution. It assumes no Lean, MLIR

Kundenservice Verlag Dr. Otto Schmidt KG
Neumannstraße 10, 40235 Düsseldorf | kundenservice@otto-schmidt.de | 0221-93738-997 (Inland)