

1. Record Nr.	UNISA996500065903316
Autore	Avidan Shai
Titolo	Computer Vision - ECCV 2022 : 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXIV
Pubbl/distr/stampa	Cham : , : Springer, , 2022 ©2022
ISBN	3-031-20053-5
Descrizione fisica	1 online resource (803 pages)
Collana	Lecture Notes in Computer Science ; ; v.13684
Altri autori (Persone)	BrostowGabriel CisséMoustapha FarinellaGiovanni Maria HassnerTal
Disciplina	006.37
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Nota di contenuto	Intro -- Foreword -- Preface -- Organization -- Contents - Part XXIV -- Improving Vision Transformers by Revisiting High-Frequency Components -- 1 Introduction -- 2 Related Work -- 3 Revisiting ViT Models from a Frequency Perspective -- 4 The Proposed Method -- 4.1 Adversarial Training with High-Frequency Perturbations -- 4.2 A Case Study Using ViT-B -- 5 Experiments -- 5.1 Experimental Setup -- 5.2 Results on ImageNet Classification -- 5.3 Results on Out-of-distribution Data -- 5.4 Transfer Learning to Downstream Tasks -- 5.5 Ablation Studies -- 5.6 Discussions -- 6 Conclusions and Future Work -- References -- Recurrent Bilinear Optimization for Binary Neural Networks -- 1 Introduction -- 2 Related Work -- 3 Methodology -- 3.1 Preliminaries -- 3.2 Bilinear Model of BNNs -- 3.3 Recurrent Bilinear Optimization -- 3.4 Discussion -- 4 Experiments -- 4.1 Datasets and Implementation Details -- 4.2 Ablation Study -- 4.3 Image Classification -- 4.4 Object Detection -- 4.5 Deployment Efficiency -- 5 Conclusion -- References -- Neural Architecture Search for Spiking Neural Networks -- 1 Introduction -- 2 Related Work -- 2.1 Spiking Neural Networks -- 2.2 Neural Architecture Search -- 3 Preliminaries -- 3.1 Leaky Integrate-and-Fire Neuron -- 3.2 NAS Without Training --

4 Methodology -- 4.1 Linear Regions from LIF Neurons -- 4.2 Sparsity-Aware Hamming Distance -- 4.3 Searching Forward and Backward Connections -- 5 Experiments -- 5.1 Implementation Details -- 5.2 Performance Comparison -- 5.3 Experimental Analysis -- 6 Conclusion -- References -- Where to Focus: Investigating Hierarchical Attention Relationship for Fine-Grained Visual Classification -- 1 Introduction -- 2 Related Work -- 2.1 Fine-Grained Visual Classification -- 2.2 Human Attention in Vision -- 3 Approach -- 3.1 Overview -- 3.2 Region Feature Mining Module.

3.3 Cross-Hierarchical Orthogonal Fusion Module -- 4 Experiments and Analysis -- 4.1 Datasets -- 4.2 Hierarchy Interaction Analysis -- 4.3 Evaluation on Traditional FGVC Setting -- 4.4 Further Analysis -- 5 Conclusions -- References -- DaViT: Dual Attention Vision Transformers -- 1 Introduction -- 2 Related Work -- 3 Methodology -- 3.1 Overview -- 3.2 Spatial Window Attention -- 3.3 Channel Group Attention -- 3.4 Model Instantiation -- 4 Analysis -- 5 Experiments -- 5.1 Image Classification -- 5.2 Object Detection and Instance Segmentation -- 5.3 Semantic Segmentation on ADE20k -- 5.4 Ablation Study -- 6 Conclusion -- References -- Optimal Transport for Label-Efficient Visible-Infrared Person Re-Identification -- 1 Introduction -- 2 Related Work -- 3 Methodology -- 3.1 Problem Formulation and Overview -- 3.2 Discrepancy Elimination Network (DEN) -- 3.3 Optimal-Transport Label Assignment (OTLA) -- 3.4 Prediction Alignment Learning (PAL) -- 3.5 Optimization -- 4 Experiments -- 4.1 Experimental Settings -- 4.2 Implementation Details -- 4.3 Main Results -- 4.4 Ablation Study -- 4.5 Discussion -- 5 Conclusion -- References -- Locality Guidance for Improving Vision Transformers on Tiny Datasets -- 1 Introduction -- 2 Related Work -- 3 Method -- 3.1 The Overall Approach -- 3.2 Guidance Positions -- 3.3 Architecture of the CNN -- 4 Experiments -- 4.1 Main Results -- 4.2 Discussion -- 4.3 Ablation Study -- 5 Conclusion -- References -- Neighborhood Collective Estimation for Noisy Label Identification and Correction -- 1 Introduction -- 2 Related Work -- 2.1 Noise Verification -- 2.2 Label Correction -- 3 The Proposed Method -- 3.1 Neighborhood Collective Noise Verification -- 3.2 Neighborhood Collective Label Correction -- 3.3 Training Objectives -- 4 Experiments -- 4.1 Experimental Setup -- 4.2 Comparisons with the State of the Art -- 4.3 Analysis.

5 Conclusions -- References -- Few-Shot Class-Incremental Learning via Entropy-Regularized Data-Free Replay -- 1 Introduction -- 2 Related Works -- 2.1 Class-Incremental Learning -- 2.2 Few-Shot Class-Incremental Learning -- 2.3 Data-Free Knowledge Distillation -- 3 Preliminaries -- 3.1 Problem Setting -- 3.2 Data-Free Replay -- 4 Methodology -- 4.1 Entropy-Regularized Data-Free Replay -- 4.2 Learning Incrementally with Uncertain Data -- 5 Experiments -- 5.1 Datasets -- 5.2 Implementation Details -- 5.3 Re-implementation of Replay-based Methods -- 5.4 Main Results and Comparison -- 5.5 Analysis -- 6 Conclusion -- References -- Anti-retroactive Interference for Lifelong Learning -- 1 Introduction -- 2 Related Work -- 2.1 Lifelong Learning -- 2.2 Adversarial Training -- 3 Proposed Method -- 3.1 Extracting Intra-Class Features -- 3.2 Generating and Fusing Task-Specific Models -- 4 Experiments and Results -- 4.1 Datasets -- 4.2 Implementation Details -- 4.3 Results and Comparison -- 4.4 Ablation Study -- 5 Conclusion -- References -- Towards Calibrated Hyper-Sphere Representation via Distribution Overlap Coefficient for Long-Tailed Learning -- 1 Introduction -- 2 Related Works -- 3 Methodology -- 3.1 Build vMF Classifier on Hyper-Sphere -- 3.2 Quantify Distribution Overlap Coefficient on Hyper-Sphere -- 3.3 Improve

Representation of Feature and Classifier via o -- 3.4 Calibrate Classifier Weight Beyond Training via o -- 4 Experiments -- 4.1 Long-Tailed Image Classification Task -- 4.2 Long-Tailed Semantic and Instance Segmentation Task -- 4.3 Ablation Study -- 5 Conclusions -- References -- Dynamic Metric Learning with Cross-Level Concept Distillation -- 1 Introduction -- 2 Related Work -- 3 Proposed Approach -- 3.1 Dynamic Metric Learning -- 3.2 Hierarchical Concept Refiner -- 3.3 Cross-Level Concept Distillation -- 3.4 Discussions -- 4 Experiments.

4.1 Datasets -- 4.2 Evaluation Protocol -- 4.3 Implementation Details -- 4.4 Main Results -- 4.5 Experimental Analysis -- 5 Conclusion -- References -- MEnet: A Memory-Based Network with Dual-Branch for Efficient Event Stream Processing -- 1 Introduction -- 2 Related Work -- 2.1 Event-Based Representations -- 2.2 Memory-Based Networks -- 3 Event Camera Model -- 4 Method -- 4.1 Dual-Branch Structure -- 4.2 Double Polarities Calculation Method -- 4.3 Point-Wise Memory Bank -- 4.4 Training and Testing Strategies -- 5 Experiments -- 5.1 Experimental Setup -- 5.2 Ablation Study -- 5.3 Object Recognition -- 5.4 Gesture Recognition -- 6 Conclusion -- References -- Out-of-distribution Detection with Boundary Aware Learning -- 1 Introduction -- 2 Related Work -- 3 Preliminaries -- 4 Boundary Aware Learning -- 4.1 Representation Extraction Module (REM) -- 4.2 Representation Sampling Module (RSM) -- 4.3 Representation Discrimination Module (RDM) -- 5 Experiments -- 5.1 Dataset -- 5.2 Experimental Setup -- 5.3 Ablation Study -- 5.4 Detection Results -- 5.5 Visualization of trivial and hard OOD features -- 6 Conclusion -- References -- Learning Hierarchy Aware Features for Reducing Mistake Severity -- 1 Introduction -- 2 Related Work -- 3 HAF: Proposed Approach -- 3.1 Fine Grained Cross-Entropy (LCE_{fine}) -- 3.2 Soft Hierarchical Consistency (L_{shc}) -- 3.3 Margin Loss (L_m) -- 3.4 Geometric Consistency (L_{gc}) -- 4 Experiments and Results -- 4.1 Experimental Setup -- 4.2 Training Configurations -- 4.3 Results -- 4.4 Coarse Classification Accuracy -- 5 Analysis -- 5.1 Ablation Study -- 5.2 Mistakes Severity Plots -- 5.3 Discussion: Hierarchical Metrics -- 6 Conclusion -- References -- Learning to Detect Every Thing in an Open World -- 1 Introduction -- 2 Related Work -- 3 Learning to Detect Every Thing -- 3.1 Data Augmentation: Background Erasing (BackErase).

3.2 Decoupled Multi-domain Training -- 4 Experiments -- 4.1 Cross-category Generalization -- 4.2 Cross-Dataset Generalization -- 5 Conclusion -- References -- KVT: k-NN Attention for Boosting Vision Transformers -- 1 Introduction -- 2 Related Work -- 2.1 Self-attention -- 2.2 Transformer for Vision -- 3 k-NN Attention -- 3.1 Vanilla Attention -- 3.2 k-NN Attention -- 3.3 Theoretical Analysis on k-NN Attention -- 4 Experiments for Vision Transformers -- 4.1 Experimental Settings -- 4.2 Results on ImageNet -- 4.3 The Impact of Number k -- 4.4 Convergence Speed of k-NN Attention -- 4.5 Other Properties of k-NN Attention -- 4.6 Comparisons with Temperature in Softmax -- 4.7 Visualization -- 4.8 Object Detection and Semantic Segmentation -- 5 Conclusion -- References -- Registration Based Few-Shot Anomaly Detection -- 1 Introduction -- 2 Related Work -- 2.1 Anomaly Detection -- 2.2 Few-Shot Learning -- 2.3 Few-Shot Anomaly Detection -- 3 Problem Setting -- 4 Method -- 4.1 Feature Registration Network -- 4.2 Normal Distribution Estimation -- 4.3 Inference -- 5 Experiments -- 5.1 Experimental Setups -- 5.2 Comparison with State-of-the-Art Methods -- 5.3 Ablation Studies -- 5.4 Visualization Analysis -- 6 Conclusion -- References -- Improving Robustness by Enhancing Weak Subnets -- 1 Introduction -- 2 Related

Work -- 3 EWS: Training by Enhancing Weak Subnets -- 3.1 Subnet Construction and Impact on Overall Performance -- 3.2 Finding Particularly Weak Subnets -- 3.3 EWS: Enhancing Weak Subnets with Knowledge Distillation -- 3.4 Combining EWS with Adversarial Training -- 4 Experiments -- 4.1 Improving Corruption Robustness -- 4.2 Improving Adversarial Robustness -- 5 Ablation and Discussions -- 5.1 Search Strategies and Hyper-Parameters -- 5.2 Vulnerability of Blocks and Layers -- 6 Conclusion -- References.
Learning Invariant Visual Representations for Compositional Zero-Shot Learning.

Sommario/riassunto

The 39-volume set, comprising the LNCS books 13661 until 13699, constitutes the refereed proceedings of the 17th European Conference on Computer Vision, ECCV 2022, held in Tel Aviv, Israel, during October 23-27, 2022. The 1645 papers presented in these proceedings were carefully reviewed and selected from a total of 5804 submissions. The papers deal with topics such as computer vision; machine learning; deep neural networks; reinforcement learning; object recognition; image classification; image processing; object detection; semantic segmentation; human pose estimation; 3d reconstruction; stereo vision; computational photography; neural networks; image coding; image reconstruction; object recognition; motion estimation.
