

1. Record Nr.	UNISA996464420303316
Titolo	Pattern recognition : 42nd DAGM German Conference, DAGM GCPR 2020, Tubingen, Germany, September 28 - October 1, 2020, proceedings / / Zeynep Akata, Andreas Geiger, Torsten Sattler (editors)
Pubbl/distr/stampa	Cham, Switzerland : , : Springer, , [2021] ©2021
ISBN	3-030-71278-8
Descrizione fisica	1 online resource (504 pages)
Collana	Lecture Notes in Computer Science ; ; v.12544
Disciplina	006.4
Soggetti	Pattern recognition systems
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Nota di contenuto	Intro -- Preface -- Organization -- Awards -- Contents -- Characterizing the Role of a Single Coupling Layer in Affine Normalizing Flows -- 1 Introduction -- 2 Related Work -- 3 Single Affine Coupling Layer -- 3.1 Architecture -- 3.2 KL Divergence Minimizer -- 3.3 Tight Bound on Loss -- 3.4 Determining the Optimal Rotation -- 3.5 Independent Outputs -- 4 Layer-Wise Learning -- 5 Conclusion -- References -- Semantic Bottlenecks: Quantifying and Improving Inspectability of Deep Representations -- 1 Introduction -- 2 Related Work -- 3 Semantic Bottlenecks -- 3.1 Supervised Semantic Bottlenecks (SSBs) -- 3.2 Unsupervised Semantic Bottlenecks (USBs) -- 4 Quantification of Layer Output Inspectability -- 4.1 AUIC Metric -- 4.2 Discussion -- 5 Results -- 5.1 Setup -- 5.2 Quantitative Improvements with SBs -- 5.3 Qualitative Improvements with SBs -- 6 Conclusion -- References -- Bias Detection and Prediction of Mapping Errors in Camera Calibration -- 1 Introduction -- 2 Fundamentals -- 3 Related Work -- 4 Detecting Systematic Errors -- 5 The Expected Mapping Error (EME) -- 6 Experimental Evaluation -- 7 Results -- 8 Application in Calibration Guidance -- 9 Conclusion and Future Research -- References -- Learning to Identify Physical Parameters from Video Using Differentiable Physics -- 1 Introduction -- 1.1 Related Work -- 2 Background -- 2.1 Unconstrained and Constrained Dynamics -- 3 Method -- 3.1 Network Architecture -- 3.2 Training

Losses -- 4 Experiments -- 4.1 Simulated Scenarios and Observability Analysis -- 4.2 Results -- 4.3 Qualitative Video Prediction Results -- 4.4 Discussion and Limitations -- 5 Conclusion -- References -- Assignment Flow for Order-Constrained OCT Segmentation -- 1 Introduction -- 2 Assignment Flow -- 3 OCT Data Representation by Covariance Descriptors -- 4 Ordered Layer Segmentation -- 4.1 Ordering Constraint. 4.2 Ordered Assignment Flow -- 5 Experimental Results and Discussion -- 6 Conclusion -- References -- Boosting Generalization in Bio-signal Classification by Learning the Phase-Amplitude Coupling -- 1 Introduction -- 2 Related Work -- 3 Learning to Detect the Phase-Amplitude Coupling -- 4 Experiments -- 4.1 Data Sets -- 4.2 Training Procedures and Models -- 4.3 Evaluation Procedures -- 4.4 Generalization on the Sleep Cassette Data Set -- 4.5 Generalization on the ISRUC-Sleep Data Set -- 4.6 Comparison to the Relative Positioning Task -- 4.7 Results on the Sleep Telemetry and CHB-MIT Data Sets -- 4.8 Impact of the Window Size -- 4.9 Frozen vs Fine-Tuned Encoder -- 4.10 Architecture -- 5 Conclusions -- References -- Long-Tailed Recognition Using Class-Balanced Experts -- 1 Introduction -- 2 Related Work -- 3 Method -- 3.1 Long-Tailed Recognition Using Class-Balanced Experts -- 3.2 Out-of-Distribution Detection for Experts -- 3.3 Fusing Expert Posteriors -- 4 Experiments -- 4.1 Oracle Performance -- 4.2 Effect of Joint Calibration Module -- 4.3 Diverse Ensembles with Experts -- 4.4 Comparison to the State-of-the-Art -- 4.5 Discussion -- 5 Conclusion -- References -- Analyzing the Dependency of ConvNets on Spatial Information -- 1 Introduction -- 2 Related Work -- 3 Methods and Experimental Setup -- 3.1 Approaches to Constrain Information -- 3.2 Experimental Setup -- 4 Results -- 4.1 Spatial and Channel-Wise Shuffle on VGG-16 -- 4.2 Spatial Information at Later Layers is Not Necessary -- 4.3 Patch-Wise Spatial Shuffle -- 4.4 Detection Results on VOC Datasets -- 5 Conclusion -- References -- Learning Monocular 3D Vehicle Detection Without 3D Bounding Box Labels -- 1 Introduction -- 1.1 Related Work -- 1.2 Contribution -- 2 Learning 3D Vehicle Detection Without 3D Bounding Box Labels -- 2.1 Shape Representation -- 2.2 Single-Image Network. 2.3 Loss Functions -- 3 Experiments -- 3.1 Ablation Study -- 3.2 Comparison with Non-Learning-Based Methods -- 4 Conclusion -- References -- Observer Dependent Lossy Image Compression -- 1 Introduction -- 2 Related Work -- 3 Method -- 4 Experiments -- 4.1 Results -- 5 Discussion -- References -- Adversarial Synthesis of Human Pose from Text -- 1 Introduction -- 2 Related Work -- 3 Generating Human Poses from Text -- 3.1 Feature Representation -- 3.2 Architecture -- 4 Dataset and Training -- 5 Experiments -- 6 Conclusion -- References -- Long-Term Anticipation of Activities with Cycle Consistency -- 1 Introduction -- 2 Related Work -- 3 The Anticipation Framework -- 3.1 Sequence-to-Sequence Model -- 3.2 Cycle Consistency -- 3.3 Recognition Module -- 3.4 Loss Function -- 4 Experiments -- 4.1 Ablation Analysis -- 4.2 End-to-End vs. Two-Step Approach -- 4.3 Comparison with the State-of-the-Art -- 5 Conclusion -- References -- Multi-stage Fusion for One-Click Segmentation -- 1 Introduction -- 2 Related Works -- 3 Proposed Method -- 3.1 Overview -- 3.2 Multi-stage Fusion -- 3.3 Transforming User Click -- 3.4 Implementation Details -- 4 Experimental Validation -- 4.1 Datasets -- 4.2 Tap-and-Shoot Segmentation -- 4.3 Interactive Image Segmentation -- 5 User Study -- 6 Conclusion -- References -- Neural Architecture Performance Prediction Using Graph Neural Networks -- 1 Introduction -- 2 Related Work -- 2.1 Neural Architecture Search -- 2.2 Neural Architecture Benchmark Datasets --

2.3 Performance Predictor for Neural Architectures -- 2.4 Graph Neural Networks -- 3 The Graph Encoder -- 3.1 Node-Level Propagation -- 3.2 Graph-Level Aggregation -- 4 Model Details -- 4.1 Message -- 4.2 Update -- 4.3 Aggregation -- 5 The NAS-Bench-101 Dataset -- 6 Experiments -- 6.1 Performance Prediction -- 6.2 Training Behaviour -- 6.3 Comparison to State of the Art.

7 Conclusion -- References -- Discovering Latent Classes for Semi-supervised Semantic Segmentation -- 1 Introduction -- 2 Related Work -- 3 Method -- 3.1 Semantic Branch -- 3.2 Latent Branch -- 3.3 Consistency Loss -- 3.4 Discriminator Network -- 4 Experiments -- 4.1 Implementation Details -- 4.2 Comparison with the State-of-the-Art -- 4.3 Ablation Experiments -- 5 Conclusion -- References --

Riemannian SOS-Polynomial Normalizing Flows -- 1 Introduction -- 2 Preliminaries -- 2.1 Normalizing Flows -- 2.2 Triangular Increasing Maps -- 3 SOS Polynomials and Triangular Increasing Maps -- 4 Riemannian Normalizing Flows -- 4.1 Riemannian Gradient -- 4.2 Exponential Parameterization -- 4.3 Application: Sampling from the Target Measure -- 5 Numerical Experiments -- 5.1 Implementation Details -- 5.2 Riemannian SOS-Polynomial Normalizing Flows -- 5.3 Exponential SOS-Polynomial Normalizing Flows -- 5.4 Comparison Between Riemannian and Exponential SOS Flow -- 6 Conclusion -- References --

Automated Water Segmentation and River Level Detection on Camera Images Using Transfer Learning -- 1 Introduction -- 2 Related Work -- 3 Transfer Learning for Water Segmentation -- 4 Water Segmentation Experiments -- 4.1 Protocol -- 4.2 Datasets -- 4.3 Performance Criteria -- 4.4 Results and Analysis -- 5 River Level Estimation Experiments -- 5.1 Datasets -- 5.2 Protocol -- 5.3 Results and Analysis -- 6 Conclusion -- References --

Does SGD Implicitly Optimize for Smoothness? -- 1 Introduction -- 2 Related Work -- 3 Does SGD Implicitly Optimize for Smoothness? -- 3.1 Measuring the Smoothness of a Function -- 3.2 First-Order Smoothness Measures -- 3.3 Second-Order Smoothness Measures -- 3.4 Smoothness Measures in Related Work -- 4 Experiments -- 4.1 Experimental Setup -- 4.2 Monotonicity -- 4.3 Optimality -- 5 Conclusion -- References.

Looking Outside the Box: The Role of Context in Random Forest Based Semantic Segmentation of PolSAR Images -- 1 Introduction -- 2 Projection-Based Random Forests -- 3 Experiments -- 3.1 Data -- 3.2 Polarimetric Scattering Vectors -- 3.3 Estimation of Polarimetric Sample Covariance Matrices -- 3.4 Polarimetric Sample Covariance Matrices -- 3.5 Summary -- 4 Conclusion and Future Work -- References --

Haar Wavelet Based Block Autoregressive Flows for Trajectories -- 1 Introduction -- 2 Related Work -- 3 Block Autoregressive Modeling of Trajectories -- 3.1 Conditional Normalizing Flows for Sequential Data -- 3.2 Haar Wavelet Based Invertible Transform -- 3.3 Haar Block Autoregressive Framework -- 4 Experiments -- 4.1 Stanford Drone -- 4.2 Intersection Drone -- 5 Conclusion -- References --

Center3D: Center-Based Monocular 3D Object Detection with Joint Depth Understanding -- 1 Introduction and Related Work -- 2 Center3D -- 2.1 CenterNet Baseline -- 2.2 Regressing 3D Center Points -- 2.3 Enriching Depth Information -- 2.4 Reference Area -- 3 Experiments -- 3.1 Implementation Details -- 3.2 Center3D -- 3.3 LID -- 3.4 DepJoint and Reference Area -- 3.5 Comparison to the State of the Art -- 4 Conclusion -- References --

Constellation Codebooks for Reliable Vehicle Localization -- 1 Introduction -- 2 Related Work -- 3 Codebook Generation -- 3.1 Definition of Uniqueness -- 3.2 Codeword Extraction -- 4 Localization Using Constellation Codebooks -- 4.1 Hash Table Generation -- 4.2 Feature Association -- 5 Experimental Results -- 5.1 Dataset -- 5.2 Localization Framework -- 5.3

Constellation Codebooks -- 5.4 Localization Results -- 6 Conclusions
-- References -- Towards Bounding-Box Free Panoptic Segmentation
-- 1 Introduction -- 2 Related Work -- 3 Panoptic Segmentation -- 3.1
Semantic Segmentation -- 3.2 Hough Voting -- 3.3 Watershed
Energies.
3.4 Triplet Loss Network.
