

Twitter Thread by Steven Edwards



Steven Edwards

@stephenwithavee



Time for #PapersThatMakeYouGoHmmm! A weekly summary of new ML papers from arXiv that make me think one or more of:

- 1. That looks useful!**
- 2. That's an interesting approach!**
- 3. A business could be built around this!**
- 4. How did they do that?!**

Integrating prediction in mean-variance portfolio optimization

<https://t.co/QWaryyQzam>

Corporate Social Responsibility and Corporate Governance: A cognitive approach

<https://t.co/7S7MkxrJEH>

The Gender Pay Gap Revisited with Big Data: Do Methodological Choices Matter?

<https://t.co/ykUEv6isZQ>

Sustainable and Resilient Systems for Intergenerational Justice

<https://t.co/1pPKDEh85z>

A Core of E-Commerce Customer Experience based on Conversational Data using Network Text Methodology

<https://t.co/2CgOY7mZH9>

Deep Learning for Market by Order Data

<https://t.co/Nzs544G9ST>

The economic dependency of the Bitcoin security

<https://t.co/UxdJQn9HKg>

On Technical Trading and Social Media Indicators in Cryptocurrencies' Price Classification Through Deep Learning

<https://t.co/f53jC2V5cz>

LATE for History

<https://t.co/iLy6UBfMvc>

Supportive 5G infrastructure policies are essential for universal 6G: Evidence from an open-source techno-economic simulation model using remote sensing

<https://t.co/1WMHwXyhtT>

Integrating Floor Plans into Hedonic Models for Rent Price Appraisal

<https://t.co/01tmCxtBpg>

Personality Traits and the Marriage Market

<https://t.co/JiXL2IJPxd>

Time-varying properties of asymmetric volatility and multifractality in Bitcoin

<https://t.co/rTtGE7PUe8>

REST: Relational Event-driven Stock Trend Forecasting

<https://t.co/FDAoNZgxNd>

How Misuse of Statistics Can Spread Misinformation: A Study of Misrepresentation of COVID-19 Data

<https://t.co/yisQnqgax1>

Aggregate Modeling and Equilibrium Analysis of the Crowdsourcing Market for Autonomous Vehicles

<https://t.co/VVnZYkLenh>

The corruptive force of AI-generated advice

<https://t.co/JxplsShpnJ>

The Golden Age of the Mathematical Finance

<https://t.co/5UWtrMLuFX>

Deep Reinforcement Learning for Portfolio Optimization using Latent Feature State Space (LFSS) Module

<https://t.co/9210MxNh8Y>

Hedging of Financial Derivative Contracts via Monte Carlo Tree Search

<https://t.co/Zrj7YsDivl>

Optimising Long-Term Outcomes using Real-World Fluent Objectives: An Application to Football

<https://t.co/ArOFaGcolK>

SLAKE: A Semantically-Labeled Knowledge-Enhanced Dataset for Medical Visual Question Answering

<https://t.co/587J7vurn6>

Clockwork Variational Autoencoders for Video Prediction

<https://t.co/CnLG5iFFCv>

Gifsplanation via Latent Shift: A Simple Autoencoder Approach to Progressive Exaggeration on Chest X-rays

<https://t.co/fwycwMG0We>

Fake News Detection: a comparison between available Deep Learning techniques in vector space

<https://t.co/jGPmEwwKnU>

ELIXIR: Learning from User Feedback on Explanations to Improve Recommender Models

<https://t.co/Vq6GK26YUm>

L2E: Learning to Exploit Your Opponent

<https://t.co/mj1wN8FwVd>

A Study on the Manifestation of Trust in Speech

<https://t.co/GiE4b80x7R>

AI Can Stop Mass Shootings, and More

<https://t.co/zfHlij5ZNz>

Reinforcement Learning for Datacenter Congestion Control

<https://t.co/3egxwyn3Td>

FIXME: Enhance Software Reliability with Hybrid Approaches in Cloud

<https://t.co/brnJt0ITk6>

Training Microsoft News Recommenders with Pretrained Language Models in the Loop

<https://t.co/OEa0iH5JFo>

Unsupervised Clustering of Time Series Signals using Neuromorphic Energy-Efficient Temporal Neural Networks

<https://t.co/ZhrOxKJx7W>

From Extreme Multi-label to Multi-class: A Hierarchical Approach for Automated ICD-10 Coding Using Phrase-level Attention

<https://t.co/2WwpmBUXiX>

Entity-level Factual Consistency of Abstractive Text Summarization

<https://t.co/7a18lXQw8B>

Understanding and Creating Art with AI: Review and Outlook

<https://t.co/qxd7lLqSJo>

Spacewalker: Rapid UI Design Exploration Using Lightweight Markup Enhancement and Crowd Genetic Programming

<https://t.co/x4jo8hblDX>

Deep Learning Approaches for Forecasting Strawberry Yields and Prices Using Satellite Images and Station-Based Soil Parameters

<https://t.co/quFFINloVd>

Towards AIops in Edge Computing Environments

<https://t.co/Bq1vusl0jJ>

BEDS: Bagging ensemble deep segmentation for nucleus segmentation with testing stage stain augmentation

<https://t.co/j19fLyb9Oz>

Genetically Optimized Prediction of Remaining Useful Life

<https://t.co/J9vxC5PNRI>

An Objective Laboratory Protocol for Evaluating Cognition of Non-Human Systems Against Human Cognition

<https://t.co/g3qsoxcZKt>

Knowledge discovery from emergency ambulance dispatch during COVID-19: A case study of Nagoya City, Japan

<https://t.co/oh7Qm5F7i0>

Towards an AI Coach to Infer Team Mental Model Alignment in Healthcare

<https://t.co/mG86vPY2pl>

Towards the Right Kind of Fairness in AI

<https://t.co/BE20tyAGpf>

Sparsely Factored Neural Machine Translation

<https://t.co/sZDUaCIW1X>

I Want This Product but Different : Multimodal Retrieval with Synthetic Query Expansion

<https://t.co/YtxVyTI6sc>

A Graph Neural Network to Model User Comfort in Robot Navigation

<https://t.co/oq35Ada9gS>

Using exoskeletons to assist medical staff during prone positioning of mechanically ventilated COVID-19 patients: a pilot study

<https://t.co/sTUIYEztwW>

Autonomous Navigation in Dynamic Environments: Deep Learning-Based Approach

<https://t.co/XViAqjhCwJ>

Formation of Social Ties Influences Food Choice: A Campus-Wide Longitudinal Study

<https://t.co/iMMa4QxAtj>

ACTA: A Mobile-Health Solution for Integrated Nudge-Neurofeedback Training for Senior Citizens

<https://t.co/5jV6U6GAhc>

Fully General Online Imitation Learning

<https://t.co/0RHhj8SI8K>

Temporal Memory Attention for Video Semantic Segmentation

<https://t.co/6L5gZP2EOx>

On the Post-hoc Explainability of Deep Echo State Networks for Time Series Forecasting, Image and Video Classification

<https://t.co/OcKtSHWG0U>

Beyond Fully-Connected Layers with Quaternions: Parameterization of Hypercomplex Multiplications with $1/n$ Parameters

<https://t.co/l6AvMftn6F>

Ensemble Transfer Learning of Elastography and B-mode Breast Ultrasound Images

<https://t.co/TpYMFLpTva>

Mode-Assisted Joint Training of Deep Boltzmann Machines

<https://t.co/aEML9if9sE>

Integrating Pre-trained Model into Rule-based Dialogue Management

<https://t.co/0yqJvEZcyH>

Reward Poisoning in Reinforcement Learning: Attacks Against Unknown Learners in Unknown Environments

<https://t.co/V7cvV7lmjK>

TableLab: An Interactive Table Extraction System with Adaptive Deep Learning

<https://t.co/x431QLmZIU>

A Cooperative Memory Network for Personalized Task-oriented Dialogue Systems with Incomplete User Profiles

<https://t.co/p9nFs0U9me>

The Yin-Yang dataset

<https://t.co/L1yG0OWPHX>

Music Harmony Generation, through Deep Learning and Using a Multi-Objective Evolutionary Algorithm

<https://t.co/HZQT4Kqc1D>

ResNet-LDDMM: Advancing the LDDMM Framework Using Deep Residual Networks

<https://t.co/FQXH2WrGRG>

GraphGallery: A Platform for Fast Benchmarking and Easy Development of Graph Neural Networks Based Intelligent Software

<https://t.co/c81dElgiuT>

Engineering Education in the Age of Autonomous Machines

<https://t.co/sJcOhKpDle>

What Do We Want From Explainable Artificial Intelligence (XAI)? -- A Stakeholder Perspective on XAI and a Conceptual Model Guiding Interdisciplinary XAI Research

<https://t.co/MJZgPvBSIz>

Hierarchical VAEs Know What They Don't Know

<https://t.co/6vJN2v1XHd>

An AutoML-based Approach to Multimodal Image Sentiment Analysis

<https://t.co/VWRlp1e8is>

Boosting Deep Transfer Learning for COVID-19 Classification

<https://t.co/A5f6F2jpug>

Steadily Learn to Drive with Virtual Memory

<https://t.co/xm05d92wqs>

On the use of generative deep neural networks to synthesize artificial multichannel EEG signals

<https://t.co/61zKdEpX0Y>

EDITH :ECG biometrics aided by Deep learning for reliable Individual authentication

<https://t.co/jPhElmm5B>

Training Stacked Denoising Autoencoders for Representation Learning

<https://t.co/uKTsoKDKgC>

TransFuse: Fusing Transformers and CNNs for Medical Image Segmentation

<https://t.co/SARCnpNwAg>

Training Larger Networks for Deep Reinforcement Learning

<https://t.co/gzAKBfNcdS>

Anomalous Sound Detection with Machine Learning: A Systematic Review

<https://t.co/yLk3RbxOOm>

How RL Agents Behave When Their Actions Are Modified

<https://t.co/lrc27s8s7S>

Diverse Auto-Curriculum is Critical for Successful Real-World Multiagent Learning Systems

<https://t.co/NRC8Q1Dk4w>

Player-Centered AI for Automatic Game Personalization: Open Problems

<https://t.co/mSfmztneSi>

Jira: a Kurdish Speech Recognition System Designing and Building Speech Corpus and Pronunciation Lexicon

<https://t.co/YLBVXgk79P>

AI Ethics Needs Good Data

<https://t.co/DAtB3CfLzm>

Why Talking about ethics is not enough: a proposal for Fintech's AI ethics

<https://t.co/luZxZoRkyK>

Sequential Recommendation in Online Games with Multiple Sequences, Tasks and User Levels

<https://t.co/CipbumMwLh>

Fast End-to-End Speech Recognition via a Non-Autoregressive Model and Cross-Modal Knowledge Transferring from BERT

<https://t.co/amS3PWqZOS>

CHARET: Character-centered Approach to Emotion Tracking in Stories

<https://t.co/glbPb48kio>

DAC: Deep Autoencoder-based Clustering, a General Deep Learning Framework of Representation Learning

<https://t.co/TKLt9iu2Nz>

Reinforcement Learning for IoT Security: A Comprehensive Survey

<https://t.co/e2F4maHr4a>

Resilient Machine Learning for Networked Cyber Physical Systems: A Survey for Machine Learning Security to Securing Machine Learning for CPS

<https://t.co/32rYU1TZZW>

Domain Adversarial Reinforcement Learning

<https://t.co/JSBY6Vb2M1>

Learning Intents behind Interactions with Knowledge Graph for Recommendation

<https://t.co/OY8vxtBIOY>

Adversarial defense for automatic speaker verification by cascaded self-supervised learning models

https://t.co/gqi8RuIlN_A

PAQ: 65 Million Probably-Asked Questions and What You Can Do With Them

<https://t.co/PXRvg4qvxf>

Interactive Learning from Activity Description

<https://t.co/31YbQLZQvJ>

DeepRA: Predicting Joint Damage From Radiographs Using CNN with Attention

<https://t.co/wG0zctmbTg>

Improving Automated Visual Fault Detection by Combining a Biologically Plausible Model of Visual Attention with Deep Learning

<https://t.co/yyD2nuVTfu>

Multi-class Generative Adversarial Nets for Semi-supervised Image Classification

<https://t.co/ezwaLcp6Ga>

Wasserstein Proximal of GANs

<https://t.co/gu8E7HcGaX>

Universal Adversarial Perturbations for Malware

<https://t.co/hmykJNnJje>

Towards Robust Visual Information Extraction in Real World: New Dataset and Novel Solution

<https://t.co/pt4JfjyLCJ>

A novel method for object detection using deep learning and CAD models

<https://t.co/79JfZ0nLpj>

Intelligent Software Web Agents: A Gap Analysis

<https://t.co/92pGGLfKEp>

VitrAI -- Applying Explainable AI in the Real World

<https://t.co/Nn9YeUsVx0>

Rethinking Eye-blink: Assessing Task Difficulty through Physiological Representation of Spontaneous Blinking

<https://t.co/Qjl0MJ5Vf7>

Exploiting Spline Models for the Training of Fully Connected Layers in Neural Network

<https://t.co/uKVPq84EfT>

Improving Object Detection in Art Images Using Only Style Transfer

<https://t.co/pXbPlx0jQu>

Transformer Language Models with LSTM-based Cross-utterance Information Representation

<https://t.co/RIboGXdizN>

A Decentralized Approach Towards Responsible AI in Social Ecosystems

<https://t.co/x4VivBrqWH>

DEEPPF0: End-To-End Fundamental Frequency Estimation for Music and Speech Signals

<https://t.co/eKGz0hAxXO>

Therapeutics Data Commons: Machine Learning Datasets and Tasks for Therapeutics

<https://t.co/wAHnPiA6wC>

A Machine Learning model of the combination of normalized SD1 and SD2 indexes from 24h-Heart Rate Variability as a predictor of myocardial infarction

<https://t.co/8LEjDbAXNp>

Recurrent Rational Networks

<https://t.co/QsVaicMeu2>

Robust PDF Document Conversion Using Recurrent Neural Networks

<https://t.co/GmRrx5Qiv>

VAE Approximation Error: ELBO and Conditional Independence

<https://t.co/fTzOCBiLWX>

DINO: A Conditional Energy-Based GAN for Domain Translation

<https://t.co/qp0cOgm03Q>

A Mathematical Principle of Deep Learning: Learn the Geodesic Curve in the Wasserstein Space

<https://t.co/VJSUzwRcXF>

Less is More: Pre-training a Strong Siamese Encoder Using a Weak Decoder

<https://t.co/kSqin9QLdn>

Improving Hierarchical Adversarial Robustness of Deep Neural Networks

<https://t.co/mC7Cvg1PcF>

Domain Adaptation for Medical Image Analysis: A Survey

<https://t.co/eJBBMJnI5D>

Regular Expressions for Fast-response COVID-19 Text Classification

<https://t.co/6Eaw9HE42k>

Transfer Learning for Linear Regression: a Statistical Test of Gain

<https://t.co/VjvcLRo4Ir>

DeeperForensics Challenge 2020 on Real-World Face Forgery Detection: Methods and Results

<https://t.co/l9Qrklu3ir>

Enhanced Magnetic Resonance Image Synthesis with Contrast-Aware Generative Adversarial Networks

<https://t.co/PQkL2XhIG8>

NFCNN: Toward a Noise Fusion Convolutional Neural Network for Image Denoising

<https://t.co/fhBNkwq8PX>

A Comprehensive Review of Deep Learning-based Single Image Super-resolution

<https://t.co/fsork7CxTr>

Dynamic Memory based Attention Network for Sequential Recommendation

<https://t.co/o9t1FXypki>

Sparse-Interest Network for Sequential Recommendation

<https://t.co/MS5CI0w4LL>

ADOM: Accelerated Decentralized Optimization Method for Time-Varying Networks

<https://t.co/Vbdrsxve7a>

Deep Neural Networks based Invisible Steganography for Audio-into-Image Algorithm

<https://t.co/gw8Wlpw3ex>

Quiz-Style Question Generation for News Stories

<https://t.co/NUYssrgmSd>

TapNet: The Design, Training, Implementation, and Applications of a Multi-Task Learning CNN for Off-Screen Mobile Input

<https://t.co/rMlymn4dfQ>

Towards Adversarial-Resilient Deep Neural Networks for False Data Injection Attack Detection in Power Grids

<https://t.co/vh8GvCY6ZI>

End-to-end learnable EEG channel selection with deep neural networks

<https://t.co/pKRFwib7LH>

How Faithful is your Synthetic Data? Sample-level Metrics for Evaluating and Auditing Generative Models

<https://t.co/4M9GIITBJG>

POLA: Online Time Series Prediction by Adaptive Learning Rates

<https://t.co/LSTQo0zQlj>

Graph Learning with 1D Convolutions on Random Walks

<https://t.co/bmXrLWx5rk>

Preventing Posterior Collapse Induced by Oversmoothing in Gaussian VAE

<https://t.co/PumNHJh5QZ>

Rethinking Co-design of Neural Architectures and Hardware Accelerators

<https://t.co/pWkCQtXuaB>

Centroid Transformers: Learning to Abstract with Attention

<https://t.co/d5RvbJ4j0E>

Firefly Neural Architecture Descent: a General Approach for Growing Neural Networks

<https://t.co/FzRFIJP7UK>

Re-identification of Individuals in Genomic Datasets Using Public Face Images

<https://t.co/cGC3CLWttM>

Federated Evaluation and Tuning for On-Device Personalization: System Design & Applications

<https://t.co/rxJWi0F0mK>

Globally-Robust Neural Networks

<https://t.co/ypdG6K1yxq>

Scaling Neuroscience Research using Federated Learning

<https://t.co/CYfMhiOFAU>

Complex Momentum for Learning in Games

<https://t.co/t4pROL50V1>

Cross-SEAN: A Cross-Stitch Semi-Supervised Neural Attention Model for COVID-19 Fake News Detection

<https://t.co/vpCq4UiEsm>

DESED-FL and URBAN-FL: Federated Learning Datasets for Sound Event Detection

<https://t.co/YnfdONIDDQ>

Accelerated Simulations of Molecular Systems through Learning of their Effective Dynamics

<https://t.co/wPptFLMYkf>

Variational Autoencoder for Speech Enhancement with a Noise-Aware Encoder

<https://t.co/Z5vy8ttfaT>

CheXternal: Generalization of Deep Learning Models for Chest X-ray Interpretation to Photos of Chest X-rays and External Clinical Settings

<https://t.co/YhMIEtAv34>

Coupled Feature Learning for Multimodal Medical Image Fusion

<https://t.co/odDT1nMCZY>

LambdaNetworks: Modeling Long-Range Interactions Without Attention

<https://t.co/5Nxb3BOo3d>

Evolving GAN Formulations for Higher Quality Image Synthesis

<https://t.co/EatFRyBMbv>

End-to-end lyrics Recognition with Voice to Singing Style Transfer

<https://t.co/xMM0gL8CdX>

Deep Neural Network Based Differential Equation Solver for HIV Enzyme Kinetics

<https://t.co/Zd0n1kYr2e>

Prioritizing Original News on Facebook

<https://t.co/gUWAWGu1TC>

Topological Obstructions to Autoencoding

<https://t.co/weK3WN6AQm>

Interpretable COVID-19 Chest X-Ray Classification via Orthogonality Constraint

<https://t.co/z8C1xWOz27>

Topological Deep Learning: Classification Neural Networks

<https://t.co/BK6OEZgjqe>

GradInit: Learning to Initialize Neural Networks for Stable and Efficient Training

<https://t.co/e2rZM5IFaV>

A Survey of Machine Learning for Computer Architecture and Systems

<https://t.co/AgmBfoG2Ku>

Inverse Reinforcement Learning in the Continuous Setting with Formal Guarantees

<https://t.co/3S57gdjDPe>

Few-Shot Graph Learning for Molecular Property Prediction

<https://t.co/i6BjHPIKRo>

A Koopman Approach to Understanding Sequence Neural Models

<https://t.co/NKXebVYsxL>

Certified Robustness to Programmable Transformations in LSTMs

<https://t.co/rzZrMXfftp>

End-2-End COVID-19 Detection from Breath & Cough Audio

<https://t.co/mDdDbyJk5a>

Boosting Low-Resource Biomedical QA via Entity-Aware Masking Strategies

<https://t.co/mV4GI4VKMz>

Context-Aware Prosody Correction for Text-Based Speech Editing

<https://t.co/5ZexKNxtxJ>

Modeling the Hallucinating Brain: A Generative Adversarial Framework

<https://t.co/JaaqAP8kzp>

Differentiating Surgeon Expertise Solely by Eye Movement Features

<https://t.co/wdkDOooOEz>

End-to-End Automatic Speech Recognition with Deep Mutual Learning

<https://t.co/rhcJnk4LSf>

Flow-Mixup: Classifying Multi-labeled Medical Images with Corrupted Labels

<https://t.co/1mH7eros1q>

Large-Context Conversational Representation Learning: Self-Supervised Learning for Conversational Documents

<https://t.co/akbrYa9RT9>

IronMan: GNN-assisted Design Space Exploration in High-Level Synthesis via Reinforcement Learning

<https://t.co/Yw92O9Diix>

Making the most of your day: online learning for optimal allocation of time

<https://t.co/zaWvxl4lld>

Learning the Noise of Failure: Intelligent System Tests for Robots

<https://t.co/H8XgGCjcEW>

Axial Residual Networks for CycleGAN-based Voice Conversion

<https://t.co/pONkkt3VH1>

Semi Supervised Learning For Few-shot Audio Classification By Episodic Triplet Mining

<https://t.co/PQUEI5yhoL>

Exploring Transformers in Natural Language Generation: GPT, BERT, and XLNet

<https://t.co/0BUklyDcgQ>

Twin Augmented Architectures for Robust Classification of COVID-19 Chest X-Ray Images

<https://t.co/ALgZIH3etJ>

Machine Learning Based Cyber Attacks Targeting on Controlled Information: A Survey

<https://t.co/TdbFrPwsIP>

Hierarchical Transformer-based Large-Context End-to-end ASR with Large-Context Knowledge Distillation

<https://t.co/GGUpwmvC48>

Self-Supervised Features Improve Open-World Learning

<https://t.co/Q2xNlyfJW6>

A Data Quality-Driven View of MLOps

<https://t.co/eQKhvYhUGI>

How to Learn when Data Reacts to Your Model: Performative Gradient Descent

<https://t.co/QN4iO6ki4b>

Learning Student-Friendly Teacher Networks for Knowledge Distillation

<https://t.co/DgKFAF1V1V>

GradPIM: A Practical Processing-in-DRAM Architecture for Gradient Descent

<https://t.co/4NbcthTmEm>

Intermediate Layer Optimization for Inverse Problems using Deep Generative Models

<https://t.co/lnKWDTeh0B>

Intelligent Electric Vehicle Charging Recommendation Based on Multi-Agent Reinforcement Learning

<https://t.co/MrbgPYe3h0>

Guided Interpolation for Adversarial Training

<https://t.co/7k6Hck9ByS>

CAP-GAN: Towards Adversarial Robustness with Cycle-consistent Attentional Purification

<https://t.co/pZaKRrybX8>

Perceptually Constrained Adversarial Attacks

<https://t.co/6fQi3Z3PAG>

Machine Learning Methods for the Design and Operation of Liquid Rocket Engines -- Research Activities at the DLR Institute of Space Propulsion

<https://t.co/g48JL4vPPo>

CATE: Computation-aware Neural Architecture Encoding with Transformers

<https://t.co/4GANBBXbEB>

Doping: A technique for efficient compression of LSTM models using sparse structured additive matrices

<https://t.co/AgB279ov90>

CrossLight: A Cross-Layer Optimized Silicon Photonic Neural Network Accelerator

<https://t.co/KdRZEfjYZZ>

Online Apprenticeship Learning

<https://t.co/5yFlv6jPQc>

Self-Reorganizing and Rejuvenating CNNs for Increasing Model Capacity Utilization

<https://t.co/87a52OkRCn>

Machine Learning for Mechanical Ventilation Control

<https://t.co/fCD6PZsg92>

Towards automatic extraction and validation of on-street parking spaces using park-out events data

<https://t.co/2s0Gww6xOj>

Neural Network Libraries: A Deep Learning Framework Designed from Engineers' Perspectives

<https://t.co/twjVSW3DIA>

How Convolutional Neural Networks Deal with Aliasing

<https://t.co/Xk7kHhbs0l>

Detection and severity classification of COVID-19 in CT images using deep learning

<https://t.co/zOsH8AKwDT>

Cancer Gene Profiling through Unsupervised Discovery

<https://t.co/zjklQLWbB8>

Generating 3D structures from a 2D slice with GAN-based dimensionality expansion

<https://t.co/sO53SsXarN>

Accelerating COVID-19 research with graph mining and transformer-based learning

<https://t.co/Rnf0Ny8x2Y>

UserReg: A Simple but Strong Model for Rating Prediction

<https://t.co/WTWRJe6qFO>

Certiably Robust Variational Autoencoders

<https://t.co/yJILKtrZAd>

Confidence-Aware Learning Assistant

<https://t.co/1HWHeg6l7V>

Vehicle to Vehicle (V2V) Communication Protocol: Components, Benefits, Challenges, Safety and Machine Learning Applications

<https://t.co/365f93sRgS>

Query-by-Example Keyword Spotting system using Multi-head Attention and Softtriple Loss

<https://t.co/T8DvRUxwbS>

Deep Convolutional and Recurrent Networks for Polyphonic Instrument Classification from Monophonic Raw Audio Waveforms

<https://t.co/rsu2GF9Ufo>

Saliency-Aware Class-Agnostic Food Image Segmentation

<https://t.co/6XwYnXeRfQ>

Representing Alzheimer's Disease Progression via Deep Prototype Tree

<https://t.co/l7V2BZnkuF>

Bi-APC: Bidirectional Autoregressive Predictive Coding for Unsupervised Pre-training and Its Application to Children's ASR

<https://t.co/WDiLhG5kre>

VIPPrint: A Large Scale Dataset of Printed and Scanned Images for Synthetic Face Images Detection and Source Linking

<https://t.co/Qg0tFYzql4>

INSTA-YOLO: Real-Time Instance Segmentation

<https://t.co/temd1GTKUg>

DeepGLEAM: a hybrid mechanistic and deep learning model for COVID-19 forecasting

<https://t.co/18A3HP8zfw>

A model for traffic incident prediction using emergency braking data

<https://t.co/DbcuGFZSK7>

MetaGrad: Adaptation using Multiple Learning Rates in Online Learning

<https://t.co/ddE1rzSbR9>

Cockpit: A Practical Debugging Tool for Training Deep Neural Networks

<https://t.co/m38h270ytx>

Neural Architecture Search as Program Transformation Exploration

<https://t.co/Xa8husAHix>

Universal Adversarial Perturbations Through the Lens of Deep Steganography: Towards A Fourier Perspective

<https://t.co/PpGmmXgzRV>

Broad-UNet: Multi-scale feature learning for nowcasting tasks

<https://t.co/8DiCf40AVw>

What does LIME really see in images?

<https://t.co/sGAaD331kd>

Bayesian Uncertainty Estimation of Learned Variational MRI Reconstruction

<https://t.co/wBYJuYuW7F>

Deep Reinforcement Learning for Backup Strategies against Adversaries

<https://t.co/umrehSyR0G>

Enhancing into the codec: Noise Robust Speech Coding with Vector-Quantized Autoencoders

<https://t.co/2OTRj0jL5x>

Bayesian Neural Network Priors Revisited

<https://t.co/8CC4eC64bG>

Hybrid quantum convolutional neural networks model for COVID-19 prediction using chest X-Ray images

<https://t.co/iupCTtwqfo>

Comparison of Machine Learning Classifiers to Predict Patient Survival and Genetics of GBM: Towards a Standardized Model for Clinical Implementation

<https://t.co/2OUySQbMMF>

Mediastinal lymph nodes segmentation using 3D convolutional neural network ensembles and anatomical priors guiding

<https://t.co/G4zmOU6gBt>

A Non-Intrusive Machine Learning Solution for Malware Detection and Data Theft Classification in Smartphones

<https://t.co/36LhXJkfKz>

Content-Aware Speaker Embeddings for Speaker Diarisation

<https://t.co/4dMtS1Z7Ba>

Deep Sound Field Reconstruction in Real Rooms: Introducing the ISOBEL Sound Field Dataset

<https://t.co/AqllgUBtbO>

Guided Variational Autoencoder for Speech Enhancement With a Supervised Classifier

<https://t.co/M5tX4C1uvO>

Annotation Cleaning for the MSR-Video to Text Dataset

<https://t.co/cBQcwrZmkv>

Data Analytics and Machine Learning Methods, Techniques and Tool for Model-Driven Engineering of Smart IoT Services

<https://t.co/dnkaAUxbE6>

Bootstrapping Large-Scale Fine-Grained Contextual Advertising Classifier from Wikipedia

<https://t.co/PFad4l7j7O>

Contrastive Unsupervised Learning for Speech Emotion Recognition

<https://t.co/2n55JXA6mS>

Personalized Visualization Recommendation

<https://t.co/Xo37595w1J>

On Automatic Parsing of Log Records

<https://t.co/sDoZtEPhMs>

Embracing Domain Differences in Fake News: Cross-domain Fake News Detection using Multimodal Data

<https://t.co/1Jd7H6bQd0>

COVID-19 detection from scarce chest x-ray image data using deep learning

<https://t.co/x7r5lJ3dSx>

Deep Gait Recognition: A Survey

<https://t.co/89aFNgWBF8>

StablePose: Learning 6D Object Poses from Geometrically Stable Patches

<https://t.co/EWOtKiBQ2s>

HVAQ: A High-Resolution Vision-Based Air Quality Dataset

<https://t.co/6Flz0Uc0Of>

HandTailor: Towards High-Precision Monocular 3D Hand Recovery

<https://t.co/p2YYdJ813S>

Multi-Agent Reinforcement Learning of 3D Furniture Layout Simulation in Indoor Graphics Scenes

<https://t.co/9G4GUwdVmL>

Mobile Computational Photography: A Tour

<https://t.co/rr2nK3qjZT>

One-shot action recognition towards novel assistive therapies

<https://t.co/QtCRtssIWS>

Conceptual 12M: Pushing Web-Scale Image-Text Pre-Training To Recognize Long-Tail Visual Concepts

<https://t.co/SubgUHWjki>

Biometrics in the Era of COVID-19: Challenges and Opportunities

<https://t.co/PqJRPMD272>

Automatic Face Understanding: Recognizing Families in Photos

<https://t.co/vpyohaD6pc>

Crop mapping from image time series: deep learning with multi-scale label hierarchies

<https://t.co/CsxAa7WuYx>

SPAN: a Simple Predict & Align Network for Handwritten Paragraph Recognition

<https://t.co/G7PKLUR5t1>

Active Face Frontalization using Commodity Unmanned Aerial Vehicles

<https://t.co/7QYhzRKtdH>

Shape-Tailored Deep Neural Networks

<https://t.co/mSP09bcrtw>

A Review of Testing Object-Based Environment Perception for Safe Automated Driving

<https://t.co/Ol29qNObFm>

Galaxy Zoo DECaLS: Detailed Visual Morphology Measurements from Volunteers and Deep Learning for 314,000 Galaxies

<https://t.co/EDJS3hRhYm>

Instance Localization for Self-supervised Detection Pretraining

<https://t.co/ad4mwWQOEV>

Just Noticeable Difference for Machine Perception and Generation of Regularized Adversarial Images with Minimal Perturbation

<https://t.co/hd36ZxnZkl>

Learning to Recognize Actions on Objects in Egocentric Video with Attention Dictionaries

<https://t.co/4VTzNk4gfO>

LEAD: LiDAR Extender for Autonomous Driving

<https://t.co/DyQNHNaq3n>

Feature Pyramid Network with Multi-Head Attention for Se-mantic Segmentation of Fine-Resolution Remotely Sensed Images

<https://t.co/zdRnto5ltU>

MITNet: GAN Enhanced Magnetic Induction Tomography Based on Complex CNN

<https://t.co/QUI3l7HkR8>

PSA-Net: Deep Learning based Physician Style-Aware Segmentation Network for Post-Operative Prostate Cancer Clinical Target Volume

<https://t.co/V2Zkl520om>

Hough2Map -- Iterative Event-based Hough Transform for High-Speed Railway Mapping

<https://t.co/8GooisP8MC>

A Deep-Learning Approach For Direct Whole-Heart Mesh Reconstruction

<https://t.co/8vkdVZuGXH>

OmniDet: Surround View Cameras based Multi-task Visual Perception Network for Autonomous Driving

<https://t.co/NfBdxqEqbv>

Generation for adaption: a Gan-based approach for 3D Domain Adaption inPoint Cloud

<https://t.co/CVcU6xXUVF>

Capturing Detailed Deformations of Moving Human Bodies

<https://t.co/wQg1phtYkq>

3D Fully Convolutional Neural Networks with Intersection Over Union Loss for Crop Mapping from Multi-Temporal Satellite Images

<https://t.co/ILS56ugAOH>

Naturalizing Neuromorphic Vision Event Streams Using GANs

<https://t.co/mmHStWczXA>

Improved Bengali Image Captioning via deep convolutional neural network based encoder-decoder model

<https://t.co/B9URyhTYyS>

ChipNet: Budget-Aware Pruning with Heaviside Continuous Approximations

<https://t.co/mKftaYYAcN>

TransGAN: Two Transformers Can Make One Strong GAN

<https://t.co/7qLEeW33Cq>

Robust Lane Detection via Expanded Self Attention

<https://t.co/DwDsl0KgiK>

Fast, Accurate Barcode Detection in Ultra High-Resolution Images

<https://t.co/ubkOreEylU>

CPP-Net: Context-aware Polygon Proposal Network for Nucleus Segmentation

<https://t.co/T3UDQGdiOf>

Learning Speech-driven 3D Conversational Gestures from Video

<https://t.co/zwpi7gEZeJ>

Plug-and-Play external and internal priors for image restoration

<https://t.co/NFXdwteLwl>

Multi-Texture GAN: Exploring the Multi-Scale Texture Translation for Brain MR Images

<https://t.co/b7Tg9z7YZl>

FastHand: Fast Hand Pose Estimation From A Monocular Camera

<https://t.co/NM0pS7IXeP>

Fusion of convolution neural network, support vector machine and Sobel filter for accurate detection of COVID-19 patients using X-ray images

<https://t.co/swBKHkS3v5>

Collaborative Intelligence: Challenges and Opportunities

<https://t.co/hij4x2wq4g>

Efficient Conditional GAN Transfer with Knowledge Propagation across Classes

<https://t.co/FOU52KXX5f>

End-to-end Audio-visual Speech Recognition with Conformers

<https://t.co/QJLYOfeLMD>

Learning Depth via Leveraging Semantics: Self-supervised Monocular Depth Estimation with Both Implicit and Explicit Semantic Guidance

<https://t.co/7fkOjPyrap>

Outdoor inverse rendering from a single image using multiview self-supervision

<https://t.co/JlPqOVow2U>

K-Hairstyle: A Large-scale Korean hairstyle dataset for virtual hair editing and hairstyle classification

<https://t.co/FChVd0QSPi>

A Generative Model for Hallucinating Diverse Versions of Super Resolution Images

<https://t.co/ol2hso1j7G>

Predicting and Attending to Damaging Collisions for Placing Everyday Objects in Photo-Realistic Simulations

<https://t.co/M4OzXV8Gf9>

Segmentation-Renormalized Deep Feature Modulation for Unpaired Image Harmonization

<https://t.co/xwufYQ9iZR>

Uncertainty-Aware Semi-supervised Method using Large Unlabelled and Limited Labeled COVID-19 Data

<https://t.co/lgnu7a1kWv>

[@threadreaderapp](#) unroll