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?!

How can I choose an explainer? An Application-grounded Evaluation of Post-hoc Explanations

https://t.co/GTAaqHRgi1

Validating Label Consistency in NER Data Annotation

https://t.co/eeYBU4AUvT

A two-stage data association approach for 3D Multi-object Tracking

https://t.co/h3LijchIzC

Neural Networks, Artificial Intelligence and the Computational Brain

https://t.co/2jJHGlHYsn

Mindless Attractor: A False-Positive Resistant Intervention for Drawing Attention Using Auditory Perturbation

https://t.co/QhahrfRN5T

Boost then Convolve: Gradient Boosting Meets Graph Neural Networks

https://t.co/gM4vCIMMzW

Deep Reinforcement Learning with Spatio-temporal Traffic Forecasting for Data-Driven Base Station Sleep Control

https://t.co/xvyYotmqBG

Discussion of Ensemble Learning under the Era of Deep Learning

https://t.co/eqdMw8WNEU

Do we need to go Deep? Knowledge Tracing with Big Data

https://t.co/Z3EtibSYA3

mt5b3: A Framework for Building AutonomousTraders

https://t.co/w6sBscN3uo

SUGAR: Subgraph Neural Network with Reinforcement Pooling and Self-Supervised Mutual Information Mechanism

https://t.co/oO0r8IOshz

Classifying Scientific Publications with BERT -- Is Self-Attention a Feature Selection Method?

https://t.co/7pAUbTGfLb

Collision-Free Flocking with a Dynamic Squad of Fixed-Wing UAVs Using Deep Reinforcement Learning

https://t.co/1R4qJ5M0Eq

Adversarial Attacks for Tabular Data: Application to Fraud Detection and Imbalanced Data

https://t.co/S3fQcgbNcK

UPDeT: Universal Multi-agent Reinforcement Learning via Policy Decoupling with Transformers

https://t.co/8sPKuqAPnQ

DynaComm: Accelerating Distributed CNN Training between Edges and Clouds through Dynamic Communication Scheduling

https://t.co/UcgL7WDUGv

Noise Learning Based Denoising Autoencoder

https://t.co/hPSClsZTTp

Illuminating the Space of Beatable Lode Runner Levels Produced By Various Generative Adversarial Networks

https://t.co/7xawUMSYSW

Spatial Assembly: Generative Architecture With Reinforcement Learning, Self Play and Tree Search

https://t.co/b6PQNPyDef

Creation and Evaluation of a Pre-tertiary Artificial Intelligence (AI) Curriculum

https://t.co/7qA7BomthH

Dissonance Between Human and Machine Understanding

https://t.co/nRBcIDIoSP

A System for Automated Open-Source Threat Intelligence Gathering and Management

https://t.co/zRIE873tMW

Classification of Pedagogical content using conventional machine learning and deep learning model

https://t.co/kFt1Vr11DS

GLocalX -- From Local to Global Explanations of Black Box AI Models

https://t.co/jNEAt3yDei

An Artificial Intelligence based approach to estimating time of arrival and bus occupancy for public transport systems in Africa

https://t.co/oQkFARvo0e

Edge-Featured Graph Attention Network

https://t.co/5jRr0yngHA

Situation and Behavior Understanding by Trope Detection on Films

https://t.co/2tTVFlj7BM

Meta-Reinforcement Learning for Adaptive Motor Control in Changing Robot Dynamics and Environments

https://t.co/wsMBpdo3zG

Disentangled Recurrent Wasserstein Autoencoder

https://t.co/KNKdFN9RII

GIID-Net: Generalizable Image Inpainting Detection via Neural Architecture Search and Attention

https://t.co/eGlkz92WGB

Grounding Language to Entities and Dynamics for Generalization in Reinforcement Learning

https://t.co/8AC8HngYl1

An attention model to analyse the risk of agitation and urinary tract infections in people with dementia

https://t.co/51FwnK9i5v

Faster Convergence in Deep-Predictive-Coding Networks to Learn Deeper Representations

https://t.co/4lp4UxSYMV

Adversarial Interaction Attack: Fooling AI to Misinterpret Human Intentions

https://t.co/abJTLSptss

Understanding in Artificial Intelligence

https://t.co/b5kufxxoL5

A Literature Review of Recent Graph Embedding Techniques for Biomedical Data

https://t.co/6TRfUgvv1v

Artificial Intelligence for Emotion-Semantic Trending and People Emotion Detection During COVID-19 Social Isolation

https://t.co/OYEPMDtf5I

An Empirical Comparison of Deep Learning Models for Knowledge Tracing on Large-Scale Dataset

https://t.co/VtsS1g7pVx

Leveraging AI to optimize website structure discovery during Penetration Testing

https://t.co/tZxbNs99Tu

Is it a great Autonomous FX Trading Strategy or you are just fooling yourself

https://t.co/ted6dt7jBd

Deep Reinforcement Learning for Active High Frequency Trading

https://t.co/C4iR7RNfs2

Studying Catastrophic Forgetting in Neural Ranking Models

https://t.co/GtzumvDLEj

Motor-Imagery-Based Brain Computer Interface using Signal Derivation and Aggregation Functions

https://t.co/TxUEKByuBX

DeepPayload: Black-box Backdoor Attack on Deep Learning Models through Neural Payload Injection

https://t.co/aNjfrPm8Gd

Cooperative and Competitive Biases for Multi-Agent Reinforcement Learning

https://t.co/4lYbSyvDH3

CheXtransfer: Performance and Parameter Efficiency of ImageNet Models for Chest X-Ray Interpretation

https://t.co/8tMpbWAu1b

Stacked LSTM Based Deep Recurrent Neural Network with Kalman Smoothing for Blood Glucose Prediction

https://t.co/KzvOxzqexs

Deep Parametric Continuous Convolutional Neural Networks

https://t.co/m3jGJWSnXr

Coarse Temporal Attention Network (CTA-Net) for Driver's Activity Recognition

https://t.co/B3EG8k37SB

GENIE: A Leaderboard for Human-in-the-Loop Evaluation of Text Generation

https://t.co/rjx4yjnNQj

TrafficSim: Learning to Simulate Realistic Multi-Agent Behaviors

https://t.co/fAEPrJmbYB

AdvSim: Generating Safety-Critical Scenarios for Self-Driving Vehicles

https://t.co/wPPbWpOR36

AdvSim: Generating Safety-Critical Scenarios for Self-Driving Vehicles

https://t.co/wPPbWpOR36

GeoSim: Photorealistic Image Simulation with Geometry-Aware Composition

https://t.co/AcLGlax0fk

SceneGen: Learning to Generate Realistic Traffic Scenes

https://t.co/JNOxqvAeKB

Towards Searching Efficient and Accurate Neural Network Architectures in Binary Classification Problems

https://t.co/Tjh0adUiNv

Slot Machines: Discovering Winning Combinations of Random Weights in Neural Networks

https://t.co/XB0dKTws4J

NNStreamer: Efficient and Agile Development of On-Device AI Systems

https://t.co/s6SOklUTsp

AR-based Modern Healthcare: A Review

https://t.co/pcoZB3J3ka

Attention Based Video Summaries of Live Online Zoom Classes

https://t.co/M4ZStoy1AN

When SIMPLE is better than complex: A case study on deep learning for predicting Bugzilla issue close time

https://t.co/YTIAAWqAcd

On the Verification and Validation of AI Navigation Algorithms

https://t.co/gwBaqOeu4c

Local Navigation and Docking of an Autonomous Robot Mower using Reinforcement Learning and Computer Vision

https://t.co/YHLUQ5SpIa

LIME: Learning Inductive Bias for Primitives of Mathematical Reasoning

https://t.co/rsqGk5fQ8b

Player-AI Interaction: What Neural Network Games Reveal About AI as Play

https://t.co/DIzXf1wuW6

Probabilistic Inference for Learning from Untrusted Sources

https://t.co/PFyTH6jZyp

Teaming up with information agents

https://t.co/fpaMVZlH0k

How AI Developers Overcome Communication Challenges in a Multidisciplinary Team: A Case Study

https://t.co/kStzf3RoTX

Black-box Adversarial Attacks in Autonomous Vehicle Technology

https://t.co/0B2xkavOWt

Motion-Based Handwriting Recognition

https://t.co/lT0ybiyiAl

Affordance-based Reinforcement Learning for Urban Driving

https://t.co/Do6J5eo7j6

Randomized Ensembled Double Q-Learning: Learning Fast Without a Model

https://t.co/St76T1uGLq

Responsible AI Challenges in End-to-end Machine Learning

https://t.co/u3drnpONrR

Mining Knowledge Graphs From Incident Reports

https://t.co/Sm1wZA2gYQ

Descriptive AI Ethics: Collecting and Understanding the Public Opinion

https://t.co/3u3By4VxIz

Hostility Detection and Covid-19 Fake News Detection in Social Media

https://t.co/GlrtTAalKE

Robusta: Robust AutoML for Feature Selection via Reinforcement Learning

https://t.co/ihKgihjDfV

KDLSQ-BERT: A Quantized Bert Combining Knowledge Distillation with Learned Step Size Quantization

https://t.co/rZBlbHStLd

Neural Attention Distillation: Erasing Backdoor Triggers from Deep Neural Networks

https://t.co/oPTKZbiaHd

Interpretable Multi-Head Self-Attention model for Sarcasm Detection in social media

https://t.co/aU5g6hXOaS

Knowledge-Preserving Incremental Social Event Detection via Heterogeneous GNNs

https://t.co/mHFCoQBzCm

ItNet: iterative neural networks with tiny graphs for accurate and efficient anytime prediction

https://t.co/38Ivf3iTys

Adversarial Machine Learning in Text Analysis and Generation

https://t.co/OAaQGy3VD9

Dive into Decision Trees and Forests: A Theoretical Demonstration

https://t.co/YEd3NpQqLc

Stress Testing of Meta-learning Approaches for Few-shot Learning

https://t.co/CP6QDkSUPF

Collaborative Teacher-Student Learning via Multiple Knowledge Transfer

https://t.co/UW74MHAEwV

Analysis of Information Flow Through U-Nets

https://t.co/H7Oay6mTpj

Distilling Interpretable Models into Human-Readable Code

https://t.co/tNph0XpFDT

Invariance, encodings, and generalization: learning identity effects with neural networks

https://t.co/kCJt8bajEP

Can stable and accurate neural networks be computed? -- On the barriers of deep learning and Smale's 18th problem

https://t.co/cHHRGBip9h

Copycat CNN: Are Random Non-Labeled Data Enough to Steal Knowledge from Black-box Models?

https://t.co/8Zx7fVfIIp

Explainable Patterns: Going from Findings to Insights to Support Data Analytics Democratization

https://t.co/kuvkN4ZMZu

MPASNET: Motion Prior-Aware Siamese Network for Unsupervised Deep Crowd Segmentation in Video Scenes

https://t.co/ffTvTzz5rZ

LEAF: A Learnable Frontend for Audio Classification

https://t.co/3WaQ8MIC8A

Customer Price Sensitivities in Competitive Automobile Insurance Markets

https://t.co/3nIj6XJNYQ

Pre-training without Natural Images

https://t.co/m1lBQFT9KJ

Arabic Speech Recognition by End-to-End, Modular Systems and Human

https://t.co/6rWnWJvJiR

Ensemble learning and iterative training (ELIT) machine learning: applications towards uncertainty quantification and automated experiment in atom-resolved microscopy

https://t.co/lumhRoaVhf

Influence Estimation for Generative Adversarial Networks

https://t.co/ypMwRmSJqv

Text Line Segmentation for Challenging Handwritten Document Images Using Fully Convolutional Network

https://t.co/R9SiU4pWEy

TensorBNN: Bayesian Inference for Neural Networks using Tensorflow

https://t.co/5wEXw5ECXL

Bayesian Neural Networks for Fast SUSY Predictions

https://t.co/eEjxi4ybgU

Probabilistic Solar Power Forecasting: Long Short-Term Memory Network vs Simpler Approaches

https://t.co/FdvIcYrTJZ

Rank the Episodes: A Simple Approach for Exploration in Procedurally-Generated Environments

https://t.co/XFZLPMcllL

Deep Learning for Intelligent Demand Response and Smart Grids: A Comprehensive Survey

https://t.co/WfdmTkw69N

Intelligent Icing Detection Model of Wind Turbine Blades Based on SCADA data

https://t.co/1OjpuDYWcY

Machine learning applications for COVID-19: A state-of-the-art review

https://t.co/9DSd3TfB4D

Implicit Bias of Linear RNNs

https://t.co/IfgXSmwaFb

Open-Domain Conversational Search Assistant with Transformers

https://t.co/8iMUvwb0To

Machine learning for rapid discovery of laminar flow channel wall modifications that enhance heat transfer

https://t.co/IQKlMjf5Hv

Variational Autoencoders with a Structural Similarity Loss in Time of Flight MRAs

https://t.co/zS4DRvfhI7

Bridge the Vision Gap from Field to Command: A Deep Learning Network Enhancing Illumination and Details

https://t.co/EA9GZDwoQO

Cross-domain few-shot learning with unlabelled data

https://t.co/DW7JGPcNwS

Classification of COVID-19 X-ray Images Using a Combination of Deep and Handcrafted Features

https://t.co/JswvK6YlQJ

The Devils in the Point Clouds: Studying the Robustness of Point Cloud Convolutions

https://t.co/Z9cI9d1n2b

A Unifying Generative Model for Graph Learning Algorithms: Label Propagation, Graph Convolutions, and Combinations

https://t.co/bWrX68Pxl1

Image Denoising using Attention-Residual Convolutional Neural Networks

https://t.co/5aG06Yf2RF

Interpretable Models for Granger Causality Using Self-explaining Neural Networks

https://t.co/DR34Ed1qGd

Continual Deterioration Prediction for Hospitalized COVID-19 Patients

https://t.co/OsBrVfw7kj

Momentum^2 Teacher: Momentum Teacher with Momentum Statistics for Self-Supervised Learning

https://t.co/bKhILExhy7

PeerGAN: Generative Adversarial Networks with a Competing Peer Discriminator

https://t.co/944IW7qRsm

Collaborative Federated Learning For Healthcare: Multi-Modal COVID-19 Diagnosis at the Edge

https://t.co/jBfVPPcvlY

Optimizing Hyperparameters in CNNs using Bilevel Programming in Time Series Data

https://t.co/Y1hnvcMo2K

Deep Reinforcement Learning Optimizes Graphene Nanopores for Efficient Desalination

https://t.co/r3B4ST2XIX

Handling Non-ignorably Missing Features in Electronic Health Records Data Using Importance-Weighted Autoencoders

https://t.co/MeBl9H7kS5

Does Continual Learning = Catastrophic Forgetting?

https://t.co/jk7tasU0IR

A survey on shape-constraint deep learning for medical image segmentation

https://t.co/OuSpETBxYG

Predicting Pneumonia and Region Detection from X-Ray Images using Deep Neural Network

https://t.co/qpDYzwLw2b

Comparative Evaluation of 3D and 2D Deep Learning Techniques for Semantic Segmentation in CT Scans

https://t.co/Q2nOU79sNv

Deep Learning Models for Calculation of Cardiothoracic Ratio from Chest Radiographs for Assisted Diagnosis of Cardiomegaly

https://t.co/GOe2yLmOWF

Collaboration among Image and Object Level Features for Image Colourisation

https://t.co/bVqv6ZWQTV

Electrocardiogram Classification and Visual Diagnosis of Atrial Fibrillation with DenseECG

https://t.co/EbrLI7pYe9

The Unreasonable Effectiveness of Patches in Deep Convolutional Kernels Methods

https://t.co/DcOl3AbCds

COVID-Net CT-2: Enhanced Deep Neural Networks for Detection of COVID-19 from Chest CT Images Through Bigger, More Diverse Learning

https://t.co/5uPWFcYtzQ

Using Shape to Categorize: Low-Shot Learning with an Explicit Shape Bias

https://t.co/VTDmjNd9QD

Challenges in the application of a mortality prediction model for COVID-19 patients on an Indian cohort

https://t.co/tx5Sc89T4Y

A simple geometric proof for the benefit of depth in ReLU networks

https://t.co/TT25NN3Z3M

Emotional EEG Classification using Connectivity Features and Convolutional Neural Networks

https://t.co/CgZhQGXLhC

Deep Learning for Moving Blockage Prediction using Real Millimeter Wave Measurements

https://t.co/8MlUI1ezGy

Discrete Graph Structure Learning for Forecasting Multiple Time Series

https://t.co/MXCAmJNoaX

Heterogeneous Similarity Graph Neural Network on Electronic Health Records

https://t.co/6WhEGluvBZ

Learning from pandemics: using extraordinary events can improve disease now-casting models

https://t.co/7jWkDBEzv6

Physics-Informed Deep Learning for Traffic State Estimation

https://t.co/43BlFlrB6W

Diverse Complexity Measures for Dataset Curation in Self-driving

https://t.co/CZCCyoONoE

Phases of learning dynamics in artificial neural networks: with or without mislabeled data

https://t.co/5hgjP1yYgN

Multi-objective Search of Robust Neural Architectures against Multiple Types of Adversarial Attacks

https://t.co/aQbjE43vtA

Visual Analytics approach for finding spatiotemporal patterns from COVID19

https://t.co/pSrj6m6zj5

Learning by Watching: Physical Imitation of Manipulation Skills from Human Videos

https://t.co/SMWKivSQuU

Latent Space Analysis of VAE and Intro-VAE applied to 3-dimensional MR Brain Volumes of Multiple Sclerosis, Leukoencephalopathy, and Healthy Patients

https://t.co/XcyMard8Nf

Trilevel Neural Architecture Search for Efficient Single Image Super-Resolution

https://t.co/s5cJeOGptC

MultiBodySync: Multi-Body Segmentation and Motion Estimation via 3D Scan Synchronization

https://t.co/B9E2XzC4zC

Data-driven discovery of multiscale chemical reactions governed by the law of mass action

https://t.co/M2HkwCZBQm

Temporal Clustering of Disorder Events During the COVID-19 Pandemic

https://t.co/5M4wjqpubk

Mispronunciation Detection in Non-native (L2) English with Uncertainty Modeling

https://t.co/zYC4jpnXkH

Comparison of Machine Learning for Sentiment Analysis in Detecting Anxiety Based on Social Media Data

https://t.co/0B1e6SVE2d

Exponential Kernels with Latency in Hawkes Processes: Applications in Finance

https://t.co/S44KgRfJjs

Deciding What to Learn: A Rate-Distortion Approach

https://t.co/s6TPmGNN5M

Artificial Intelligence for IT Operations (AIOPS) Workshop White Paper

https://t.co/ATqL7pEWkP

The Geometry of Deep Generative Image Models and its Applications

https://t.co/EwSU9rEiaw

Comparisons of Graph Neural Networks on Cancer Classification Leveraging a Joint of Phenotypic and Genetic Features

https://t.co/jcJIQk6gwJ

A Neophyte With AutoML: Evaluating the Promises of Automatic Machine Learning Tools

https://t.co/diHKuYRFdO

Empirical Evaluation of Supervision Signals for Style Transfer Models

https://t.co/0lxQ3WXzaN

Needmining: Designing Digital Support to Elicit Needs from Social Media

https://t.co/xRhSg5WL9I

A New Artificial Neuron Proposal with Trainable Simultaneous Local and Global Activation Function

https://t.co/252Oegv871

Video Summarization Using Deep Neural Networks: A Survey

https://t.co/LffqAz9gVb

Nowcasting Gentrification Using Airbnb Data

https://t.co/tySM5cSpy9

How Shift Equivariance Impacts Metric Learning for Instance Segmentation

https://t.co/OPdhtdDmC2

@threadreaderapp unroll