GHOST DAY

APPLIED MACHINE LEARNING CONFERENCE

Book of Abstracts

Poznan University of Technology
8 November 2019

Organized by

Group of
Horribly
Optimistic
STatisticians

About

GHOST Day: AMLC

The aim of GHOST Day: AMLC is to create an outstanding space for exchanging experiences between machine learning practitioners and, most importantly, a friendly environment for updating knowledge in the rapidly changing area of data analysis. To this end, conference speakers include recognized representatives of the scientific community publishing at top-tier global conferences such as NeurIPS or ICML and experts from leading companies building machine learning-based products. During the talks, methods for dealing with practical machine learning problems were presented. Both large companies such as IBM, Intel, PSI, Allegro, Pearson as well as startups such as Logic.AI, Tooploox or StethoMe participated in the event.

Unlike other events regarding machine learning applications, the conference was not focused on any specific programming technology (e.g. R or python), but on the techniques, methods, and algorithms. The conference also created a gateway for interested software developers and other talented peoples who wanted to enter the fascinating world of artificial intelligence. The event gave them the opportunity to meet companies operating in the industry, participate in presentations at an intermediate level and to present their own projects during a poster session. The conference was held on 8th November 2019 in the modern facilities of Poznań University of Technology – one of the two Polish institutions ranked in prestigious Shanghai Ranking 2019 in Computer Science Engineering (on positions 401-500, ex aequo with University of Warsaw) and the first Polish higher education institution offering bachelor degree studies on Artificial Intelligence.

The conference is organized by Group of Horribly Optimistic STatisticians (GHOST), a machine learning research student group operating at Faculty of Computing, Poznań University of Technology. The group gathers talented students and graduates interested in deepening their understanding of machine learning algorithms. GHOST organizes advanced lectures, works on machine learning projects (including internships), participates in data competitions and provides a unique environment for meeting and working with peoples fascinated by the recent progress in artificial intelligence.

Organizing committee

- Mateusz Lango
- Adam Czyżewski
- Zuzanna Kocur
- Magda Miara
- Piotr Miara

Volunteers

- Julia Będziechowska
- Damian Horna
- Hanna Klimczak
- Marcin Konieczny
- Oliwia Masian
- Sławomir Pioroński
- Kamil Pluciński
- Marek Prywer
- Agnieszka Rusin
- Mateusz Skiba
- Jakub Tomczak
- Martyna Wilczyńska

Official Sponsors

Platinum Sponsor

Pearson

Gold Sponsor

OLX GROUP

WYDZIAŁ
INFORMATYKI
POLITECHNIKI POZNAŃSKIEJ

Silver Sponsor

antmicro

F-Secure

Honorary Patronages

Ministry
of Digital Affairs

MINISTRY
OF ENTREPRENEURSHIP
AND TECHNOLOGY

POZnań
Honorary Patronage
of the Mayor of Poznań

Polskie Stowarzyszenie Sztucznej Inteligencji

GHOST DAY
APPLIED MACHINE LEARNING CONFERENCE

AGENDA

9:00 - 10:00
Registration (Lecture and Conference Centre of Poznan University of Technology)

10:00 - 10:15
Conference opening (Aula Magna, CW 4)

10:15 - 11:15
Keynote lecture: "Quo vadis Reinforcement Learning?"
by Wojciech Jaśkowski, PhD (Aula Magna, CW 4)

11:15 - 11:45
Coffee break (Pasaż Bernardyński, CW)

11:45 - 13:45

Session 1: NLP (CW 8)	Session 2: General ML (CW 9)	Session 3: Computer Vision (CW 7)

13:45 - 14:45
Lunch break (CW 0.53)

14:45 -16:45

Session 4: ML in Medicine (CW 8)	Session 5: ML Applications (CW 9)	Session 6: Student Session (CW 7)

16:45 - 17:45
Poster session (Pasaż Bernardyński, CW)

17:45 - 18:00
Closing session (CW 7)

Keynote lecture

Chair: Prof. Jerzy Stefanowski, Vice-president of Polish AI Society

Quo vadis Reinforcement Learning?

by **Wojciech Jaśkowski**, NNAISENSE, Switzerland

Reinforcement Learning (RL) is a fascinating field spanning supervised learning, control theory, optimization, and human-computer interaction with potential applications ranging from game-playing through trading to industrial control. However, there is a significant gap between its potential and reality. Thus, how wide is the RL reality gap? What RL can deliver and what it cannot do (yet) and why? What are the main problems and research directions? In the talk, I will attempt to address those questions balancing between high- and low-level perspectives on RL and its applications.

Wojciech Jaśkowski, PhD works as a Research Scientist in the Intelligent Automation group at NNAISENSE, a Swiss company co-founded by Jürgen Schmidhuber (most noted for his work on LSTMs) which aims at building the first practical general-purpose AI. Previously Wojciech worked as a post-doctoral researcher at the Swiss AI Lab IDSIA and as an Assistant Professor at PUT. Winner of many international competitions, like Google ROADEF/EURO Challenge and Microsoft Imagine Cup. In 2014 his team developed AI for a race car that won Hello World Open 2014, beating the algorithms developed by the teams from tech giants like Facebook and Google. Recently, he was part of the team that won Learning to Run Competition - one of the official challenges of the NeurIPS conference. He is also one of the creators of ViZDoom framework - a popular research platform for Reinforcement Learning.

Talks

Session 1: Natural Language Processing

Chair: Agnieszka Ławrynowicz, PhD, Dr Habil.

Word sense disambiguation and mechanisms improving NLP

by **Mateusz Półtorak**, Pearson

Tematem spotkania będzie ujednoznacznianie słów należących do grupy homonimów (przykład - znaczenie słowa "zamek" w kontekście "zgubiłem klucz do zamku"). Porozmawiamy o tym jak reprezentować podobieństwa między zdaniami oraz jak przetwarzać te podobieństwa w taki sam sposób jak przetwarza się obrazy lub zdjęcia. Po zapoznaniu słuchaczy ze szkieletem architektury zdolnej do powyższej analizy, talk skupi się na prezentacji dwóch bardziej zaawansowanych metod jej modyfikacji. Pierwszą z nich są dwuwymiarowe embeddingi części mowy, wprowadzające do modelu namiastkę mechanizmu atencji. Druga metoda, to implementacja miary soft-cosine, jako część grafu Tensorflow, umożliwiająca maksymalne wykorzystanie wiedzy zgromadzonej w pre-trenowanych embeddingach słów.

Mateusz Półtorak jest hobbystą modelowania i aproksymacji złożonych mechanizmów otaczającego nas świata. W czasie swojej przygody z deep learningiem, najwięcej uwagi poświęcił zagadnieniom rozpoznawania emocji w sygnałach audio, transferowi stylu artystycznego oraz tematyce word sense disambiguation w dziedzinie przetwarzania języka naturalnego. Od ponad roku zajmuje się szeroko pojętym data science w Pearson.

Debugging black-box NLP models

by **Filip Graliński**, Applica.ai

Uczenie maszynowe przyniosło w ostatnich latach w przetwarzaniu języka naturalnego olbrzymie postępy. Modele oparte na uczeniu maszynowym, zwłaszcza sieci neuronowe, rodzą przy tym nowe wyzwania z punktu widzenia inżynierii oprogramowania. W jaki sposób w zrozumiały dla człowieka sposób ewaluować i porównywać wyniki modeli? Jak wykrywać „lokalne" regresje modeli w sytuacji ogólnego polepszenia wyników? W trakcie wys-

tąpienia zaprezentowane zostaną techniki i narzędzia, które, przynajmniej częściowo, pozwolą odpowiedzieć na te pytania.

dr Filip Graliński, Chief Data Scientist w Applica.ai oraz pracownik Wydziału Matematyki i Informatyki UAM, zajmuje się zastosowaniem uczenia maszynowego do problemów przetwarzania tekstu oraz metodami ewaluacji modeli ML

Artificial intelligence - human emotions

by **Konrad Juszczyk**, Sentimenti

Algorytmy rekomendują nam hotele na weekend, dobierają reklamy do czytanego tekstu, a także sortują i filtrują komentarze. Czy to wszystko może czynić sztuczna inteligencja pozbawiona emocji? W SENTIMENTI uważamy, że sztuczna inteligencja potrzebuje ludzkich emocji. Dlatego przebadaliśmy ponad 20.000 Polaków, którzy oznaczyli po osiem emocji dla ponad 30.000 słów i znaczeń. Uzyskaliśmy największą w Polsce emotywną bazę języka, która w połączeniu ze Słowosiecią i uczeniem maszynowym pozwala na mierzenie emocji w każdym tekście – od krótkich wzmianek przez posty i artykuły do baz na dowolny temat. Podczas wystąpienia zdradzimy kilka szczegółów dotyczących naszych badań.

dr Konrad Juszczyk to specjalista językoznawstwa kognitywnego i komputerowego, badacz multimodalnej komunikacji międzyludzkiej, metafor, gestów i emocji. W projekcie odpowiada za badania rynku innowacji związanych z komunikacją.

A long long time ago in an NLP model far far away - are computers able to create stories?

by **Dawid Wiśniewski**, Applica.ai

Ostatnie lata obfitowały w liczne przełomy w dziedzinie przetwarzania języka naturalnego, których efekty widujemy w codziennej pracy (asystenci głosowi, systemy zamiany mowy na tekst, automatyczne tłumaczenie dokumentów między językami). Czy wśród tych przełomów wyłania się zdolność do, wydawałoby się zarezerwowanej dla naszego gatunku, kreatywności? Czy komputery potrafią generować użyteczne teksty, które w przyszłości będą mogły stanowić automatycznie wytworzone artykuły, recenzje lub przemówienia? Jak wyglądają podejścia używane do automatycznego gen-

erowania tekstów?

Dawid Wiśniewski jest doktorantem i asystentem na Wydziale Informatyki Politechniki Poznańskiej, gdzie zajmuje się przetwarzaniem języka naturalnego oraz technologiami semantycznymi. Od 2013 zawodowo związany z przetwarzaniem języka, Deep Learning Specialist w Applica.ai.

Session 2: General Machine Learning

Chair: Marek Wydmuch

Hacking AI

by **Michał Łukaszewsk**, Intel

Sieci neuronowe wydają się być niesamowitymi algorytmami o zaskakujących możliwościach i niezwykłej potędze ujawnianej w starciu z ogromnymi wolumenami danych. A jednak są bardzo wrażliwe i podatne na zakłócenia, oszukanie ich jest niesamowicie proste. Podatności te stanowią duży problem przy zastosowaniach produkcyjnych. Prezentacja przedstawi kilka najbardziej popularnych sztuczek jakimi możemy ogłupić sieć neuronową.

Michał Łukaszewski pracuje na stanowisku Deep Learning Manager w firmie Intel. Po ponad 18 latach jako developer zajął się prowadzeniem zespołu specjalizującego się w benchmarkingu rozwiązań Deep Learningowych. Pasjonat czystego kodu, optymalnych procesów, dobrze zaprojektowanych baz danych.

Scale-invariant online learning

by **Wojciech Kotłowski**, PUT

The talk concerns online learning with linear models, where the algorithm predicts on sequentially revealed instances (feature vectors), and is compared against the best linear function (comparator) in hindsight. Popular algorithms in this framework, such as Online Gradient Descent (OGD), have parameters (learning rates), which ideally should be tuned based on the scales of the features and the optimal comparator, but these quantities only become available at the end of the learning process. We show how to resolve the tuning problem by proposing online algorithms making predictions which are invariant under arbitrary rescaling of the features. The algorithms have no parameters to tune, do not require any prior knowledge on the scale of the instances or the comparator, and achieve regret bounds matching (up to a logarithmic factor) that of OGD with optimally tuned separate learning rates per dimension, while retaining comparable runtime performance.

Wojciech Kotłowski, PhD is an Assistant Professor at Poznań University of Technology. From 2009 to 2012, he was a post-doctoral researcher in

Centrum Wiskunde & Informatica (Amsterdam, Netherlands) in the group of Peter Grünwald. His main research interests are in the theory of machine learning, particularly in the online learning with adversarial data. He is the author of many publications on top-tier machine learning conferences such as NeurIPS, ICML and COLT.

Is practical AutoML more than CASH?

by **Alexandre Quemy**, IBM

For a broader adoption and scalability of machine learning systems, the construction and configuration of machine learning workflow need to gain in automation. In the last few years, several techniques have been developed in this direction, known as AutoML. AutoML focuses on solving the black-box optimization problem named CASH, for Combined Algorithm Selection and Hyperparameter optimization. As its name states, CASH consists in selection the proper algorithm and its set of hyperparameters. However, a rule of thumb regarding the effort in building a machine learning workflows is that 80% of the time is spent on data preparation. Therefore, one might wonder why most effort is spent on automating the least time consuming part. In this talk, we will try to answer this question and to show that indeed, under time constraint, it is often - but not always - more interesting to allocate time to automated data preparation rather than hyperparameter tuning. Last, we will briefly present a robust two-stage optimization process to allocate time between data preparation and CASH.

Alexandre Quemy joined IBM in 2015 where he currently holds a senior engineer position in the Data and AI organization. Prior joining IBM, he worked at Inria research center in France, in the team Machine Learning and Optimization on multi-objective optimization and meta-heuristics. In parallel of his activities at IBM, he is pursuing a Ph.D. at PUT, under the supervision of Pr. Wrembel, in the field of AutoML in unstructured space with a focal point on explainability.

Scaling Machine Learning Systems up to Billions of Predictions per Day

by **Carmine Paolino**, OLX

Whether it's a linear regressor or a system of connected deep learning models, getting your models ready is half the battle. Did you design your machine learning system to survive the onslaught of visitors from your

latest Reddit and Hacker News post? Or the influx of users shopping during Black Friday? Are you ready for a world filled with flakey networks, invalid data, and impatient users? In this talk you'll learn how to design and architect your machine learning systems for the harsh realities it will face. We will show you how we tackled these problems in a real, complex machine learning system at OLX and scaled it to serve up to billions of predictions per day, using software engineering principles while debunking the myth that Python code cannot scale.

Carmine Paolino is a Senior Data Scientist at OLX Group where he works on computer vision, interpretable machine learning, and recommender systems. He has more than 4 years of industry experience as a Data Scientist, and more than 9 as a Software Engineer. Soon after his bachelor thesis he published a book chapter about large-scale distributed graph analysis, and he has recently collaborated on a state-of-the-art saliency detection technique. He's also an avid open source contributor, with patches to MXNet Model Server and Theano amongst others. When he's not working, he loves to take photos, write music, DJ, and learn more about Machine Learning, psychology, and the stock market. Once saved the Coding Horror blog from losing all its content.

Session 3: Computer Vision

Chair: Prof. Krzysztof Krawiec

Learning binary image representations and the path towards 3D

by **Tomasz Trzciński**, Tooploox

Learning a compact and efficient representation of images, words or 3D shapes is critical to many real-life applications, including visual search, 3D reconstruction and object recognition. In this talk, we will present several methods that address this problem, starting from binary image descriptors learned with linear projections to unsupervised training of 3D point cloud representations built with Generative Adversarial Networks to be applied in the context of autonomous cars. We will start with binary local image descriptors trained with linear projections (D-BRIEF, ECCV'12) and boosting (BinBoost, CVPR'13) and then move on to more recent works on learning constellation descriptors with Siamese architecturs (SCoNE, ECCVW'18) and unsupervised learning with GANs (BinGAN, NeurIPS'18). We will conclude this talk with a few samples of our ongoing research on learning compact representations of 3D point clouds using various generative models.

Tomasz Trzciński is an Assistant Professor at Warsaw University of Technology and Chief Scientist and Partner at Tooploox. He received his PhD in computer vision from EPFL, previously worked at Google and Qualcomm and held a Visiting Scholar position at Stanford University. Member of IEEE and Computer Vision Foundation, he frequently serves as a committee member fo major vision and machine learning converences (CVPR, ICCV, ECCV, NeurIPS). Recently co-founded Comixify, VC-backed startup that uses GANs and reinforcement learning to turn your video into comics.

How to find that red dress? - Visual Search in Allegro

by **Piotr Rybak**, Allegro

Jak skutecznie znaleźć na Allegro koszulę widzianą ostatnio w sklepie? Albo te nietypowe filiżanki, w których podają kawę w kawiarni na dole? Można oczywiście próbować opisać je słownie, ale czy nie byłoby prościej gdyby wystarczyło zrobić im zdjęcie? Podczas tej prezentacji przedstawię różne

modele używane do wyszukiwania obrazem, opiszę jak one działają oraz pokażę przykładowe wyniki. Będzie dużo obrazków :)

Piotr Rybak pracuje jako Senior Research Engineer w zespole Machine Learning Research w Allegro i zajmuje się rozwojem modeli do wyszukiwania obrazowego. Absolwent matematyki, pasjonat uczenia maszynowego, miłośnik gier planszowych i wspinaczki.

Session 4: Machine Learning in Medicine

Chair: Jerzy Błaszczyński, PhD

Transfer learning in healthcare industry

by **Michał Kierzynka**, Roche

Transfer learning is a powerful technique to boost the performance of a deep learning model by reusing the features learned on a related task. The use of pretrained models became a de facto standard in many practical applications. However, the healthcare industry often has very specific data sets that are rather dissimilar to the large-scale and publicly available data sets used to pretrain the models. Therefore, are there any benefits of applying transfer learning in healthcare? Come and listen to find out.

Michał Kierzynka works as a Data Scientist at Roche. He received his PhD in computer science from PUT, and previously worked at PSNC. He was serving as a certified instructor at NVIDIA Deep Learning Institute. With his strong background in parallel computing and GPUs (early adopter of CUDA), he loves to work on high performance software tools. He specializes in deep learning and computer vision.

Sound event detection using self-consistency checking neural networks

by **Tomasz Grzywalski**, StethoMe

Sound event detection is a challenging task even for most sophisticated neural networks. Effective merging of local and global context of time series is the key for achieving good performance. In my talk I will present an interesting new approach for utilization of global context, developed by my team at StethoMe. The method introduces new definition of training target for audio events that is inspired by Capsule Networks proposed by Geoffrey Hinton and residual training. Examples will include real-life cases of detecting normal and abnormal lung sounds in stethoscope recordings.

Tomasz Grzywalski: Researcher, developer and tester of cutting-edge machine learning algorithms for sound & image processing. Currently works as a Lead Machine Learning Specialist at StethoMe where he, together with his team, develops algorithms for for automatic detection and classification of lung and heart sounds. In 2018 his team won 2nd place in FEMH

Voice Data Challenge, a competition for diagnosing voice disorders based on vowel sounds.

The machine learning for de novo drug design

by **Rafał A. Bachorz**, PSI

The Recurrent Neural Networks (RNNs) are capable of solving the problems involving some sort of sequence. Actually, the sequence can be considered as the context prototype. This means that not only the observation recorded at certain point carries the content, but also the relation with the neighbors brings significant piece of knowledge. One can mention a substantial number of use cases where the same computational techniques, i.e. RNNs, are employed in order to solve semantically very distant problems. For instance the natural language modelling, speech recognition, machine translation, artificial music generation are only a few examples. Here I would like to introduce a particular perspective of the application of the RNNs, where the chemical information contained in carefully chosen molecular data set is transferred into the properly designed predictive model. The considered chemical information involves the chemical species itself represented as the SMILES code, as well as the supplementary information reflecting the chemical and physical properties. Both components, dubbed as static and dynamic features, are simultaneously treated in the neural network architecture. This is possible thanks to the technique of the static feature incorporation into the RNNs. At the end, the obtained generative predictive model provides very strong capability of molecule generation where the property of new species can be controlled by providing desired properties at the prediction phase. This, on the other hand, can be considered as the scientific tool applicable in the areas of de novo drug design or computational material science.

Rafał A. Bachorz, PhD is Data Scientist at PSI Poland and Head of Molecular Modeling Laboratory, Institute of Medical Biology, Polish Academy of Sciences.

Prediction of phase change in CHAD bipolar disorder

by **Olga Kamińska**, Britenet

Zmiana fazy u pacjentów w afektywnej chorobie dwubiegunowej może powodować nieodwracalne zmiany. Celem projektu była próba przewidzenia zmiany stanu w kierunku depresji/manii u badanych pacjentów. Każdy osoba

inaczej przechodzi tą chorobę dlatego tak ważna jest personalizacja algorytmu. W badaniu zostały zastosowane metody zarówno uczenia nadzorowanego jak i nienadzorowanego na danych pochodzących od prawdziwych pacjentów. Wykorzystano dane pochodzące z nagrań rozmów telefonicznych od pacjentów. Przygotowane rozwiązanie jest pierwszym krokiem do stworzenia kompletnego systemu prognostycznego, który ułatwiłby życie lekarzom oraz pacjentom.

Olga Kamińska pracuje na stanowisku Data Scientist w Grupie Britenet oraz jest doktorantką w Instytucie Badań Systemowych Polskiej Akademii Nauk w Warszawie.

Session 5: Machine Learning Applications

Chair: Julia Będziechowska

How to model students' knowledge using recurrent neural networks?

by **Mikołaj Olszewski & Mateusz Otmianowski**, Pearson

Podczas wystąpienia przedstawimy nasze doświadczenia związane z implementacją metody modelowania wiedzy uczniów za pomocą Rekurencyjnych Sieci Neuronowych (Deep Knowledge Tracing). Opowiemy o wyzwaniach jakie napotkaliśmy w trakcie tego procesu, a które były związane z tworzeniem niskopoziomowego kodu w TensorFlow, wydajnością procesu trenowania sieci, śledzeniem wyników i organizacją pracy zespołu kilku osób pracujących jednocześnie nad tym samym modelem.

Mikołaj Olszewski i Mateusz Otmianowski od kilku lat pracują jako specjaliści Data Science w firmie Pearson. W swojej codziennej pracy badają i aplikują najnowsze odkrycia w zakresie uczenia maszynowego w produktach edukacyjnych.

Machine Learning at the very edge with TensorFlow and beyond

by **Michael Gielda**, Antmicro

Recent advances in deep learning mean that not only desktop and server class machines can now perform fairly complex AI tasks. With open source ML frameworks such as Google's TensorFlow expanding to cover entirely new areas, as well as open, ML-oriented hardware and compiler stacks, we can expect a host of new applications where low-power AI will play a key role. We will look into TensorFlow Lite Micro, TVM, NVDLA and other open source technologies that shape the landscape of "Tiny ML" as it is sometimes called.

Michael Gielda is VP Business Development and co-founder of Antmicro. With a background in both computer science and the humanities, he is an ardent believer in using open source to advance entire industries. Michael is vice-chair of Marketing in the RISC-V Foundation and Chair of Marketing and Outreach in CHIPS Alliance.

OCR for Data Extraction from Receipts

by **Bartosz Ludwiczuk**, intive

Jak stworzyć własnego OCR, dopasowanego do swoich potrzeb, np. do paragonów? Jakie moduły taki OCR powinien posiadać? A co się stanie, gdy klient zrobi zdjęcie kotu lub innemu dokumentowi? Na te i podobne pytania odpowiem w czasie prezentacji dotyczącej opisu pełnego pipelinu systemu typu OCR, stworzonego od zera, którego zadaniem jest wyciąganie informacji z paragonów. Cały proces developmentu będzie opisany od pierwszego prototypu do finalnego modelu używanego na produkcji w celu pokazania iteracyjnego procesu dostarczania produktu. Cały projekt jest silnie związany zarówno z Deep-Learningiem, przetwarzaniem obrazu jak i przetwarzaniem tekstu.

Bartosz Ludwiczuk to Deep Learning Enginnier z 7-letnim doświadczeniem w temacie, który przede wszystkim zajmuje się zagadnieniami związanymi z przetwarzaniem obrazu. W wolnym czasie zgłębia tematy uczenia głębokiego lub jeździ na rowerze.

Data Science in Cybersecurity

by **Zuzanna Kunik**, F-Secure

Czy data science jest w stanie rozwiązać problemy cyber security? W jaki sposób możemy wykorzystać machine learning, aby wesprzeć analityków i konsultantów ds. bezpieczeństwa w podejmowaniu decyzji nt. tysięcy potencjalnych zagrożeń dziennie? Czy polecenia linii komend mogą być traktowane jako język naturalny? Czym jest concept drift i jaki ma wpływ na tworzenie systemów AI w cyber security?

Zuzanna Kunik od ponad 2 lat pracuje jako Data Scientist w F-Secure. Jej codzienność wykracza poza analizę danych czy przygotowywanie modeli - do jej obowiązków należy też implementacja i utrzymywanie systemów machine learningowych na backendzie.

Session 6: Student Session

Chair: Mateusz Langa

How to win Data Science competitions?

by **Paweł Sienkowski**, Fathom

Konkursy z dziedzin Machine Learning, Data Science i Data Mining to świetne pole do nauki. Dobrze zdefiniowana miara, wstępnie przygotowany zbiór danych znacznie obniżają próg wejścia - pierwszy model często można zbudować w mniej niż pół godziny. Ale nie dla samej nauki człowiek buduje modele :) Jak wygrywać konkursy? Jak pławić się w blasku i chwale? Na prezentacji usłyszycie kilka praktycznych rad i trików.

Paweł Sienkowski to Machine Learning Engineer w Fathom, laureat 8 konkursów Data Science. Pracował m. in. w Microsoft, Facebook i Palantir Technologies. W wolnych chwilach buduje autorskie escape roomy.

Making your neural networks pay more attention

by **Dariusz Max Adamski**, GHOST

Learning very long-term dependencies in sequential data is crucial in many ML tasks, such as natural language understanding and time-series analysis. Attention allows us to model dependencies regardless of the distance in data. It is even possible to replace recurrent networks completely, and achieve superior performance with attention only.In the talk, we will introduce attention, explain the popular transformer model, and finally, show how we use attention mechanisms at IMIND to improve the performance and interpretability of our models.

Dariusz Max Adamski is undergraduate student at the Poznań University of Technology, co-founder of IMIND, a startup applying various ML techniques, to help you lower your electric bill and save energy in the process. Max is also an active member of Group of Horribly Optimistic STatisticians (GHOST).

Unsupervised anomaly detection in multivariate time series data

by **Damian Horna**, GHOST

Czwarta Rewolucja Przemysłowa przynosi ze sobą olbrzymią szansę na rozwój firm, poprawę jakości produktów, a także obniżenie kosztów ich wytwarzania. Dane mające charakter wielowymiarowych szeregów czasowych są coraz chętniej gromadzone m.in. w elektrowniach lub zakładach produkcyjnych. W obliczu tak dużej ilości informacji wykrywanie różnego rodzaju anomalii staje się zadaniem bardzo wymagającym i kosztownym. W jaki sposób możemy wykorzystać możliwości uczenia maszynowego aby zautomatyzować ten proces? Podczas prezentacji przedstawię kilka metod wykrywania anomalii w danych sensorycznych i podzielę się moimi spostrzeżeniami dotyczącymi ich praktycznego wykorzystania.

Damian Horna to student ostatniego roku studiów inżynierskich na kierunku Informatyka na Wydziale Informatyki Politechniki Poznańskiej, a także członek Koła Naukowego GHOST. Obecnie pracuje jako Machine Learning Engineer w Poznańskim Centrum Superkomputerowo-Sieciowym.

GhostRacer - towards self-driving cars

by **Jakub Tomczak**, GHOST

W ostatnich latach autonomiczne pojazdy stały się głównym celem badań firm zajmujących się przemysłem samochodowym. BMW, Toyota, Ford, Volvo, a także wiele innych znanych marek jak np. Google przeznacza duże nakłady finansowe na rozwój technologii związanych z autonomicznymi pojazdami. Gwałtownemu postępowi w tej dziedzinie towarzyszy również zwiększone zapotrzebowanie na wykwalifikowanych inżynierów, którzy posiadają specjalistyczną wiedzę i umiejętności do konstrukcji tego typu pojazdów. Grupa DeepRacer działająca w ramach Koła Naukowego GHOST zajmuje się budową miniatury autonomicznego samochodu oraz poznawaniem algorytmów wykorzystywanych w takich pojazdach. W czasie prezentacji podzielimy się naszymi doświadczeniami, opowiemy czym jest ROS oraz wyjaśnimy, jak autonomiczne pojazdy "widzą" otaczającą rzeczywistość.

Jakub Tomczak to student II roku studiów magisterskich na kierunku Informatyka na Wydziale Informatyki Politechniki Poznańskiej oraz członek Koła Naukowego GHOST. Obecnie pracuje jako AI Engineer w Antmicro.

Posters

Emotion recognition using psychological signals from wearable devices

by **Weronika Michalska** (Wrocław University of Science and Technology)

The aim of presented paper is to summarize and compare recent and most important studies where recognise different emotions occurs. Our research presents important topics related to the emotion recognition. This process is very complicated and consist of many aspects and problems. We are focused on wearable devices like bands, smartwatches etc. Smart wearables allow measure many biologic parameters such as HRV, GST, ECG, EEG. The other way to recognize emotion we can analyze image from camera or body movement. We presented the most common devices used in related studies. We compared methods and types of collecting samples to give an overview in this subject. We took into account also different emotion stimulation types and common emotion models and scales. Signal preprocessing, feature extraction and learning algorithms play an important role in each study. Accuracy of correct emotion recognition depends of the amount of participants, quality of the database, number of emotions, classification and validation method and many more aspects.

Stock price prediction leveraging sentiment data and stacked CNN-LSTM architecture

by **Piotr Gloger** (Netguru)

We present our work on stacked CNN-LSTM models for stock price prediction. According to Efficient Market Hypothesis, stock prices reflect all information and consistent alpha generation is impossible, meaning there is no value in basing our predictions on historical information. Researchers in the past indicated that previous price of an asset in most cases is the best prediction for the next time-step. The results obtained with the CNN-LSTM trained models beat the previous price benchmark considerably. While the predictions still do not follow the actual price perfectly, they do very well in most cases, especially in the trending market.

Unsupervised Machine Translation

by **Kamil Pluciński** (GHOST, Poznan University of Technology)

Due to the tremendous growth in human-made content, the market value of machine translation has increased significantly in recent years. The most popular supervised approaches require large parallel corporas to achieve high performance. This is a serious limitation for its usage on under-resourced languages pairs. For such languages, the translation is often performed by introducing additional translation step to some intermediary resource-rich language (often English). However, this additional step introduces new errors and decreases the translator's performance.

Recently, the idea of unsupervised machine translation gained significant research attention. Such approaches learn (usually) linear mapping between two word-embedding spaces enabling translator training in a cross-lingual embedding space. Both word embeddings and the mapping itself can be learned without the usage of parallel corpora and any other form of supervision, besides the need for large, possibly entirely independent, collections of texts in both languages. In this work, we experimentally evaluate such an approach, previously presented in the literature. We test its generalization abilities to new language pairs, especially focusing on Polish. We also verify the impact of replacing (commonly used in the literature) word2vec embeddings with more advanced approaches i.e. FastText. Other research questions such as the impact of vocabulary size, of different initialization methods and the rotation hypothesis, are also experimentally assessed.

Intelligent Framework for Designing Urban Bus Network using Memetic Algorithm

by **Hanan Ba Ali** (Jagiellonian University)

Urban Bus Network Design (UBND) problem represents a challenge in designing routes, serving passengers and operators benefits. In this study, a suggested framework is used to solve (UBND) problem in two phases. In the first phase, a bus network with a set of routes and frequencies are generated. The predefined set of routes satisfy the constraints that are involved: route length, lack of loops, number of bus routes, capacity of route and bus, all bus stop have at least one route, fleet size. In the second phase, Memetic Algorithm (MA) is used to regenerate all possible routes structure from the predefined set. As well as, the proposed MA finds

the best structure that represents the flawless Bus Network Design. The nominated Bus Network satisfies the objectives that are involved the least value in passengers waiting time, journey time, walking time, operator and passenger costs and number of transfers. The suggested algorithm used hill climbing local search algorithm as an extra operator in Genetic Algorithm to improve routes structure during the global search. The proposed MA provides promising results compared to other studies.

Decision processes of neural networks via topological methods - the Mapper algorithm

by **Sławomir Pioroński** (GHOST, Adam Mickiewicz University)

Celem plakatu jest zaprezentowanie algorytmu Mapper i możliwości jego wykorzystania w analizie procesów decyzyjnych sieci neuronowych, co zostało zainspirowane pracą Saula i Arendta pt. „Machine Learning Explanations with Topological Data Analysis". Mapper jest algorytmem o naturze topologicznej, który obecnie znajduje coraz szersze zastosowanie w analizie danych. Posłużymy się przykładem prostej sieci neuronowej klasyfikującej cyfry z bazy MNIST Digits. Korzystając z Mappera dla wyników klasyfikacji tworzymy graf, który zachowuje informacje o podobieństwie predykcji sieci neuronowej. Używając pojęcia tzw. drogi ucieczki interpretujemy na jakie cechy obiektów sieć zwraca uwagę. Kod w oparciu o który powstały prezentowane wyniki udostępniony jest pod adresem: `https://github.com/slapio/mapper-thesis`.

Using neural network and multimode environment representation for mobile robot control

by **Milena Molska** (Poznan University of Technology)

Mobile robots use various methods to avoid obstacles and navigate in the indoor environment. Most popular are a graph and sampling-based path planning and trajectory optimization. Recently, neural networks have been proven to be efficient in path planning and control of mobile robots. In this work, I applied a neural network and supervised learning to teach the robot obstacles avoidance and navigate in the natural indoor environment. I extended the state of the art method for model-less obstacle avoidance. I improved the neural network architecture and collected data to train

the robot. I provide the depth images to the input of the network. The output is the control signal for the wheels. . In contrast to methods available in the literature, I propose the neural network which considers also a history of the previous control signal to avoid ""deadlock"" and navigate the robot smoothly in the environment. I also performed extended experimental verification which shows that the proposed neural network can be used by the mobile robot to navigate and avoid obstacles efficiently in the indoor environment.

Faster-RCNN object detection for page segmentation of document images

by **Karolina Cwojdzińska** (Alphamoon)

We propose an object detection model based on the Faster-RCNN architecture for image-based page layout segmentation. Our method allows for automated document processing regardless of the document format. This visual-based approach is able to effectively split a document into segments containing text, tables, graphics, pictures, and others. To train our model we have manually prepared a dataset from various sources, including financial documents, research papers, invoices, and legal acts, which resulted in 3800 training examples. We achieved mean AP of 85% with 0.6 IOU over 10 classes on validation set.

Efficient Algorithm for Set-Valued Prediction in Multi-Class Classification

by **Marek Wydmuch** (Poznań University of Technology)

In cases of uncertainty, a multi-class classifier preferably returns a set of candidate classes instead of predicting a single class label with little guarantee. More precisely, the classifier should strive for an optimal balance between the correctness (the true class is among the candidates) and the precision (the candidates are not too many) of its prediction. We formalize this problem within a general decision-theoretic framework that unifies most of the existing work in this area. In this framework, uncertainty is quantified in terms of conditional class probabilities, and the quality of a predicted set is measured in terms of a utility function. We address the problem of finding the Bayes-optimal prediction, i.e., the subset of class

labels with the highest expected utility. For this problem, which is computationally challenging as there are exponentially (in the number of classes) many predictions to choose from, we propose an efficient algorithm that can be applied to a broad family of utility scores.

The application of unsupervised learning to the AC susceptibility data of High-Temperature Superconductors

by **Marcin Kowalik** (AGH University of Science and Technology)

High-temperature superconductors are materials which exhibit properties like zero electrical resistance and expulsion of an external magnetic field from the interior of a superconductor. These unique properties are used in the state of the art applications for the medicine (MRI/NMR machines), science (particle accelerators), transportation (Maglev - magnetic levitation train) and electrical industry (electric power transmission, fault current limiters). The greatest drawback of superconductors is that they only function in low temperatures. The scientists are still pursuing the discovery of room-temperature superconductor. The phenomenon of high-temperature superconductivity is still not fully understood.

Our work aims to provide insights if clustering technique applied to the dataset consisting of the measurements of High-Temperature Superconductors using the AC susceptibility method, will allow recovering know relationships between different types of high-temperature superconductors.

Reconstruction of particle tracks using Deep Neural Networks and reconstruction of estimation errors

by **Mateusz Slysz** (GHOST, Poznan University of Technology)

Currently to reconstruct particle tracks in detectors (like the one in CERN) scientists use slow, sequential algotithms like the Kalman filter. Computation time grows worse than quadratic with increasing luminosity and whith higher and higher luminosities in particle detectors, this will soon become a bottleneck for discovering new particles. We tried to use a different aproach, by using Deep Neural Networks to solve this problem. Once the network is trained, the recognition is much faster. Another imporatant thing for physicists are measurment errors, which are not supplied by currently used algorithms. We used a Mixed Denisty Network layer to estimate these er-

rors. The basic concept of how this layer works, is that instead of single numbers, the netowrk returns a probability distribution.

Machine Learning Study Groups

by **Mateusz Podlasin**

For the last half a year, I am leading a study group, where people can meet together via videoconference tool and discuss together material from the fields of Data Science & Machine Learning. It started as a Facebook group, which already has 1000+ participants - `https://www.facebook.com/groups/2281635012163777/`. We already worked through the whole book "Introduction to Statistical Learning" and now we are starting "Deep Learning" series of meetings. All these meetings are recorded and posted on YouTube.

The poster would promote the group and the idea of learning together, while creating free educational resources at the same time. It would also encourage people who lack academic background, to start learning Data Science & Machine Learning. I am a great example, since I was able to transition professionally from front-end developer to junior data scientist.


GHOST (99%)
WHAT IS YOUR
SUPERPOWER?
GHOST DAY
APPLIED MACHINE LEARNING
CONFERENCE