The 3rd Annual Student Research Project Conference

Stiftung Universität Hildesheim
December 5, 2019

Contents

SESSION 1

Chair: Mofassir ul Islam Arif

Computer Aided Detection System For Lung Cancer using CT Scan Images

Team: Anurag Pandey, Gulraiz Ali, Nourhan Ahmed, Sai Ram Reddy Pulcanti
Supervisor: Lukas Brinkmeyer

Abstract

Brain tumor segmentation is an important task in medical image analysis. Early diagnosis of different brain tumors plays an important role in increasing the survival rate of the patients. Currently, automatic segmentation using deep learning methods became popular since they achieve the state-of-the-art results and can address this problem better than other methods. 3D medical images are sometimes too large to be processed as the full 3D volume due to memory limitations, thus often requires down sampling which can obscure small anomalies. Fully convolutional networks (FCN) such as U-Net have been shown to achieve high accuracy and robust performance in different tasks including MRI scans for image segmentation tasks. In this thesis, we provide an alternative 2D architecture for 3D image segmentation of brain MRI scans, based on a combination of different variants of U-net architecture and recurrent neural network (RNN) to include the spatial information in the proposed 2D approach. The proposed architectures are evaluated on BraTS dataset for benchmarking. Experimental results show that the proposed frameworks consistently improve the prediction performance of traditional U-Net. We also show that the proposed framework provides an impressive performance on detecting tumors from MRI scans compared to the known DL-based 3D segmentation architectures.

References

- Makaju, S., P. W. C. Prasad, A. Alsadoon, A. K. Singh, and A. Elchouemi. 2018. "Lung cancer detection using ct scan images." Procedia Computer Science.Elsevier. 107-114.
- Chen, J., L. Yang, Y. Zhang, M. lber, and D. Z. Chen. 2016. "Combining fully convolutional and recurrent neural networks for 3d biomedical image segmentation." Advances in neural information processing systems3036-3044.
- Firmino, M., A. H. Morais, R. M. Mendoça, M. R. Dantas, H. R. Hekis, and R. Valentim. 2014. "Computer-aided detection system for lung cancer in computed tomography scans: review and future prospects." Biomedical engineering online(13) 1: 41.

Box Office Return Prediction (BORP) using Movie Trailers

Team: Diana Artiom, Raaghav Radhakrishnan, Kalaiselvan Panneerselvam, Muhammad Mehmood Ali, Muhammad Inzimam Saghir
Supervisor: Mofassir ul Islam Arif

Abstract

Predicting movie success in its early stages is a topic of great concern for production companies. A reliable box office return prediction before the theatrical release can greatly reduce financial risk. Existing approaches have tried predicting the box office return based on metadata, poster and/or user interaction and feedback on movie or it's trailer. In this research project we predict box office return using movie trailers from YouTube and metadata information fetched from The Movie Database and Box Office Mojo. To preserve the movie temporal information, we consider two models: Long-term Temporal Convolutions (LTC) and Long Short-Term Memory(LSTM), as well as an ensemble of the two. We show that considering movie trailers and metadata together performs better than metadata alone.

References

- Krizhevsky, A., Sutskever, I., and Hinton, G. E., "Imagenet classification with deep convolutional neural networks," in Advances in neural inform-ation processing systems, pp. 1097–1105, 2012
- Simoes, Gabriel S., et al. "Movie genre classification with convolutional neural networks." 2016 International Joint Conference on Neural Net-works (IJCNN). IEEE, 2016
- Krushikanth R. Apala, Merin Jose, Supreme Motnam, C.-C. Chan, KathyJ. Liszka, and Federico de Gregorio, "Prediction of Movies Box Office Per-formance Using Social Media" in 2013 IEEE/ACM International Confer-ence on Advances in Social Networks Analysis and Mining

SESSION 2

Chair: Hadi Samer Jomaa

Cancer Detection Assistant

Team: Alexander Busch, Alexandre Carey, Jorge Manuel Orozco Prado, Tolga Akar
Supervisor: Hadi Samer Jomaa

Abstract

Cancer is a major disease in modern days. Its accurate and prompt diagnosis is imperative to procure an effective treatment. The current methods for breast cancer detection, specifically, involve a tedious and painful procedure for pathologists that entail searching for minuscule cell formations in incredibly large, irregular, and noisy microscopy-generated images. Added to this, the process is performed under tight time constraints, making it prone to human error. Because a correct diagnosis is critical for this delicate activity, a Deep Learning architecture is proposed to aid medical staff. It combines the use of image-modification heuristics to detect areas of interest, or patches, and Convolutional Neural Networks to properly classify mitosis from non-mitosis.

References

- Das K., Pranab K. Dutta. Efficient automated detection of mitotic cells from breast histological images using deep convolutional neural networks with wavelet decomposed patches. Computers in biology and medicine 104 (2019): 29-42.
- Li C., Wang X., Wenyu L., and Longin Jan L. DeepMitosis: Mitosis Detection via deep detection, verification and segmentation networks. Medical image analysis 45 (2018): 121-133.

Exploring the influence of data aggregation in parking prediction

Team: Daniela Thyssens, Shereen Elsayed, Arslan Tariq Hafiz, Shabanaz Chamurally
Supervisor: Hadi Samer Jomaa

Abstract

We investigate the effect of external features on the prediction of car par- king occupancy rates based on different amounts of training data (history) in a time series forecasting framework. More specifically, the aim is to ans- wer the following questions; What are the important external features? And how does the effect of these external features vary if we change the amount of data history the models are trained on? These questions are eva- luated by applying four different approaches on two open-source datasets, and consequently we found that the external features have a throughout positive effect and that the extent of this effect varies across training his- tories and models.

References

- Rong, Y., et al. Du-parking: Spatio-temporal big data tells you realtime parking availability. Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery Data Mining. ACM, 2018
- Camero, A., Toutouh, J., Stolfi, D. H., & Alba, E. (2018, June). Evolutionary deep learning for car park occupancy prediction in smart cities. InInternational Conference on Learning and Intelligent Optimization(pp. 386-401). Springer, Cham.
- Bath North East Somerset Council. BANES Live Car Park Occu- pancy, 2014. <https://data.bathhacked.org/Transport/BANES-Live-Car-Park- Occupancy/u3w2-9yme>

SESSION 3

Chair: Ahmed Rasheed

Multi-Relational Classification for Image-based Recommendations

Team: Baha Tanveer, Kiyoto Ikuma, Sathish Kumar Chandrasekaran, Sharmila Raghunadhan, Victor Uwaje, Jia-Jen Yang
Supervisor: Ahmed Rasheed

Abstract

Recommendation systems in general has an important application to the e-commerce and the online media industry. There has always been a demand from these sectors to improve their recommendation engines, because a major portion of their revenues depend on the outcome of their recommendation system. The recommendation system finds the relationship between users and items with employing large user purchase records, thereby recommending only the most relevant item to a user.However, not many research have considered the visual features of the item along side with its style and item information to recommend a product to the user.
In this project we aim to enhance the item recommendation task by utilizing the visual features of the item with its style relationship and item information.This research project tries to develop a scalable SVBPR, SVBPR+price information and SVBPR+Brand name information of the items using the state of the art Hit Ratio and NDCG as it measure of accuracy.

References

- J. McAuley, C. Targett, Q. Shi, A. van den HengelImage-based Recommendations on Styles And Substitutes. ACM, 2015.
- R. He and J. McAuley 2005.VBPR: Visual Bayesian Personalized Ranking from Implicit Feedback. AAI, 2016.

Volkswagen Mileage Forecast

Team: Mario David Torres Aragon, Mohit Bansal, Miguel Angel Caldas Villamizar, Gökce Sucu, Nurbakyt Kulbatshayeva
Supervisor: Rafael Drumond

Abstract

The Analytics Services and Business Development department of Volkswagen AG, acts in the name of the dealership to offer analytical insights to clients. The aim of this department is to identify individual customer needs in a personalized and timely manner. Our objective is to estimate the annual mileage (km/a) of a customers last year and to project a new mileage into the future. This information can then be used to estimate the future service event for a car. So far, the mileage has been calculated only on the basis of a linear function of the initial registration date and the last reported mileage. However, this method is very inaccurate. In this project, it is aimed to have more accurate service event prediction by comparing the performance of traditional forecasting techniques such as ARIMA against advanced machine learning models such as Long Short Term Memory (LSTM) and Convolutional Neural Network (CNN).

References

- Williams, T. A., Methodologies Used to Estimate and Forecast Vehicle Miles Traveled(VMT). No. PRC 15−40 F. Texas A&M Transportation Institute, 2016.
- Klatko, T. J., Estimation and Prediction of Statewide Vehicle Miles Traveled (VMT) byHighway Category and Vehicle Classification, 2016

Academic Performance Prediction via Collaborative Filtering

Team: Adel Morgan, Muhammad Bilal, Haider Shabbir, Sehadije Bakija, Rameez Ahmed Khan
Supervisor: Ahmed Rasheed

Abstract

Nowadays, the world's most valuable resource is data. We must transform this raw data into useful information to be able to use it and take decisions based on patterns from this data.In this project, our motivation was to predict the accepted students' performance based on their previous studies and profile. As the previous educational excellence of the admitted students will have a great impact on their academic performance. In our project we are using two types of features, static and dynamic features. Static Features are set of features each having a fixed value. Those values do not change over time and are fixed for each sample.On the other hand, in dynamic features we have textual information and varying number of grades in each profile. Recently, there has been great advancement in the field of deep learning because of increased computational performance of the systems. In our project we used multi-layer perceptron combined with auto encoders which outperformed Random Forest , Xgboost,SVM.

References

- He, X., Liao, L., Zhang, H., Nie, L., Hu, X., & Chua, T. S. (2017, April). Neural collaborative filtering. InProceedings of the 26th International Conference on World Wide Web(pp. 173-182). International World Wide Web Conferences Steering Committee.
- Conneau, A., Schwenk, H., Barrault, L., & Lecun, Y. (2016). Very deep convolutional networks for text classification.arXiv preprint arXiv:1606.01781.
- Chen, J., Zhang, H., He, X., Nie, L., Liu, W., & Chua, T. S. (2017, August). Attentive collaborative filtering: Multimedia recommendation with item-and component-level attention. InProceedings of the 40th International ACM SIGIR conference on Research and Development in Information Retrieval(pp. 335-344). ACM.
- Zhou, C., Sun, C., Liu, Z., & Lau, F. (2015). A C-LSTM neural network for text classification.arXiv preprint arXiv:1511.08630.

SESSION 4

Chair: Randolf Schulz

Differentiable Neural Dictionaries In Supervised Learning

Team: Alperen Dedeoglu, Philip Kurzendorfer, David Obando, Sebastian Pineda, Friedemann Schestag
Supervisor: Randolf Schulz

Abstract

In Neural Episodic Control [1], a memory module called differentiable neural dictionary was proposed. In a reinforcement learning setting, such memory modules yield higher scores in early episodes. The objective of this project is to apply differentiable neural dictionaries to supervised learning, in order to improve sample efficiency. The results show that memory modules lead to a lift in accuracy, especially for lower training percentages in supervised learning when they are combined with deep neural networks.

References

- Alexander Pritzel et al. “Neural Episodic Control”. In: CoRRabs/1703.01988 (2017).
- Lukasz Kaiser et al. “Learning to Remember Rare Events”. In: CoRRabs/1703.03129 (2017).
- A. Emin Orhan. “A Simple Cache Model for Image Recognition”. In: CoRRabs/1805.08709 (2018).

Learning the early stopping criterion using Reinforcement Learning

Team: Maryam Bahrami, Olasubomi Saheed Kasali, Niklas Rowohl
Supervisor: Randolf Schulz

Abstract

Training neural networks often is expensive and prone to overfitting due to the high complexity of deep networks. One way to avoid long training and overfitting is to stop training early. In an idealized environment the optimal stopping epoch is easily found by training until the error on a validation set increases. However do the error trajectories oscillate in reality, making it difficult to know the optimal epoch at which to stop during training. Learning to stop at the right epoch can therefore ameliorate the effort it takes to train a neural network while still getting a small loss. Previous work by Prechelt, et al. focuses on using three heuristic methods to determine the ideal stopping epoch. Using the training and validation error, metrics like the generalization loss were created and the training stopped once these metrics violated a certain threshold. We created a linear combination of two metrics to evaluate our results. The first metric E1 is the difference in terms of the loss between the optimal epoch and the epoch we stopped at. The second metric E2 is the difference in epochs between the optimal epoch and the epoch we stopped at.

References

- Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT press. Pp.241-245.
- Wei, Y., Yang, F., & Wainwright, M. J. (2017). Early stopping for kernel boosting algorithms: A general analysis with localized complexities.In Advances in Neural Information Processing Systems (pp. 6065-6075).
- Mahsereci, M., Balles, L., Lassner, C., & Hennig, P. (2017). Early stopping without a validation set.
- Prechelt, L. (1998). Early stopping-but when?. In Neural Networks: Tricks of the trade (pp. 55-69). Springer, Berlin, Heidelberg.
- Duvenaud, D., Maclaurin, D., & Adams, R. (2016, May). Early stopping as nonpara-metric variational inference.In Artificial Intelligence and Statistics(pp. 1070-1077)

Autonomous car control in different weather and road conditions using deep reinforcement learning

Team: Mehran Amiri, John Robert, Godwin K. Namwamba, Famakin Olawole, Jayamalathi Selvakumar
Supervisor: Jonas Falkner

Abstract

Nowadays, there exist many autonomous cars being able to perform different tasks from car control to ramp merging built upon reinforcement learning. They are mainly operating on a single type of road which is asphalt. However, this is not realistic since there exist different types of road which a car must be able to drive on, e.g. sand, ice, etc. Moreover, different weather conditions like rain and snow usually result in different conditions of roads. In this paper, we have used Deep Deterministic Policy Gradient (DDPG) algorithm to train an agent to drive on roads consist of both asphalt and sand. The algorithm exploits deep learning and uses actor critic method to learn the policy. Obtained results show that the agent learns to drive on roads consist of sand and asphalt.

References

- L. Kaiser et al., One model to learn them all. arXiv Prepr. arXiv1706.05137, 2017.
- D. C. K. Ngai and N. H. C. Yung, A multiple-goal reinforcement learning method for complex vehicle overtaking maneuvers. IEEE Trans. Intell. Transp. Syst., vol. 12, no. 2, 2011.
- Z. Wang et al., “Sample efficient actor-critic with experience replay,” arXiv Prepr. arXiv1611.01224, 2016.

PITCHES

Data augmentation strategies for object detection

Daniel Pototzky (Bosch)

Abstract

Sophisticated data augmentation strategies have shown to significantly boost object detection performance. However, little research has been done on using data augmentation for solving common problems like crowd scenes, heavy occlusion and distinguishing similar classes like cows and horses. Expanding upon one of the existing methods (like a bag of freebies) might improve performance on those tasks.

References

- Barret Zoph, Ekin D. Cubuk, Golnaz Ghiasi, Tsung-Yi Lin, Jonathon Shlens, Quoc V. Le (2019): Learning Data Augmentation Strategies for Object Detection. https://arxiv.org/abs/1906.11172
- Zhi Zhang, Tong He, Hang Zhang, Zhongyue Zhang, Junyuan Xie, Mu Li (2019): Bag of Freebies for Training Object Detection Neural Networks. https://arxiv.org/abs/1902.04103
- Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, Yi Yang (2017): Random Erasing Data Augmentation. https://arxiv.org/abs/1708.04896

PITCHES

Speaker identification with Noise Filtering

Rafael Drumond

Abstract

Speaker identification is an essential task in the field of security and speech processing. In the literature, we find many works that are able to identify a speaker in clear audio speech conditions fully. However, it is not feasible to consider clear audio data in a real-life situation where audio is usually distorted by background noise or device interference. In this research, your goal is to be able to filter the audio, or it's extracted features to facilitate speaker identification. Such techniques might include masking, de-noising auto-encoders, and generative adversarial networks. You must understand the most recent techniques and theory behind it and be able to improve the performance of the state-of-art SID models.

References

- May, Tobias. "Influence of binary mask estimation errors on robust speaker identification." Speech Communication 87 (2017): 40-48.
- May, Tobias, and Timo Gerkmann. "Generalization of supervised learning for binary mask estimation." Acoustic Signal Enhancement (IWAENC), 2014 14th International Workshop on. IEEE, 2014.
- Tirumala, Sreenivas Sremath, and Seyed Reza Shahamiri. "A review on Deep Learning approaches in Speaker Identification." Proceedings of the 8th international conference on signal processing systems. ACM, 2016.

PITCHES

Motion Reconstruction from Scarce/Sparse Data

Rafael Drumond

Abstract

In the field of entertainment such as movies and game development, there are several techniques to bring a user body motion to a digitized environment. Several research papers have been working in mapping users limbs to a digital model. This process is not always possible when data representing the actor's body is not "complete" (i.e., not all body joints or limbs are mapped). This research task aims to recover or estimate missing data given the visible part as well as to provide an analysis of the most representative body joints using accelerometers and Motion Capture data-sets.

References

- Ahmadi, Amin, et al. "3D human gait reconstruction and monitoring using body-worn inertial sensors and kinematic modeling." IEEE Sensors Journal 16.24 (2016): 8823-8831.
- Andrews, Sheldon, et al. "Real-time physics-based motion capture with sparse sensors." Proceedings of the 13th European Conference on Visual Media Production (CVMP 2016). ACM, 2016.
- Tautges, Jochen, et al. "Motion reconstruction using sparse accelerometer data." ACM Transactions on Graphics (ToG) 30.3 (2011): 18.

PITCHES

Off-Policy learning for Combinatorial Optimization Problems

Jonas Falkner

Abstract

In recent years the ML-based optimization of combinatorial problems like the Jobshop, Knapsack or TSP have received increasing interest. Most methods employ encoder-decoder architectures that embed the problem graph and decode it sequentially based on some context. Usually the model is trained with RL. However, most of the proposed models employ on-policy methods based on REINFORCE and a policy-gradient procedure, either with a baseline or in a actor-critic framework. This research project aims to develop an efficient off-policy training procedure for this kind of problems and to compare it to respective SotA on-policy approaches.

References

- Elias Khalil, Hanjun Dai, Yuyu Zhang, Bistra Dilkina, and Le Song. Learning Combinatorial Optimization Algorithms over Graphs. In Advances in Neural Information Processing Systems 30, pages 6348–6358. Curran Associates, Inc., 2017.
- Michel Deudon, Pierre Cournut, Alexandre Lacoste, Yossiri Adulyasak, and Louis-Martin Rousseau. Learning Heuristics for the TSP by Policy Gradient. In Willem-Jan van Hoeve, editor, Integration of Constraint Programming, Artificial Intelligence, and Operations Research, volume 10848, pages 170–181. Springer International Publishing, Cham, 2018.
- Kool, Wouter, Herke van Hoof, and Max Welling. "Attention, Learn to Solve Routing Problems!." arXiv preprint arXiv:1803.08475 (2018).

PITCHES

Data Analytics for Cricket

Shayan Jaweed

Abstract

There appears to be a current trend of applying machine learning models to Cricket games, for various types of supervised regression and classification tasks. The dynamics of cricket are particularly suited to such tasks as they give rise to complex data. One of, if not the most important task is quantify the winning chances of a team. This has applications in betting but more importantly bringing the heart rate down of enthusiastic fans. An "as early as possible" confident prediction that their team would win could thus go a long way stopping [1] for example. The idea is to implement and extend the state-of-the-art for early time series classification for the cricket games. Think about a better suited regularization for the algorithm from [2] as a possible extension. However, this stands as just one example, from a team of up-to 3 members it is expected that a standard implementation and novel extension be proposed. One limitation that I would impose would be for the team to stick to the temporal aspect of the game, leveraging advances however from reinforcement learning and/or multi-task learning. A stepping stone could be to mine the data available at [3].
Screen reader support enabled.

References

- Pakistan and Afghanistan cricket fans involved in fight outside Headingley. https://bit.ly/34ytxXN
- Adaptive-Halting Policy Network for Early Classification. KDD 2019
- Cricsheet freely available ball-by-ball data for international cricket matches. https://cricsheet.org/

PITCHES

Latent Embedding optimization for Few-Shot Segmentation

Lukas Brinkmeyer

Abstract

Few-shot classification has become an immensely popular problem in the deep learning community. It deals with the classification of images of a new task with novel labels after observing only a handful of samples of the respective task. Various meta-learning approaches drastically improved the state-of-the-art in recent years. However, only a few approaches focus on few-shot segmentation. The task is to analyze current state-of-the-art approaches for few-shot classification and evaluate them for their applicability to segmentation tasks while innovating additional strategies to adapt these methods more closely to a segmentation problem.

References

- Finn, Chelsea, Pieter Abbeel, and Sergey Levine. "Model-agnostic meta-learning for fast adaptation of deep networks." Proceedings of the 34th International Conference on Machine Learning-Volume 70. JMLR. org, 2017.
- A. A. Rusu, D. Rao, J. Sygnowski, O. Vinyals, R. Pascanu, S. Osindero, and R. Hadsell. Meta-learning with latent embedding optimization. In International Conference on Learning Representations, 2019.
- Rakelly, Kate, et al. "Few-shot segmentation propagation with guided networks." arXiv preprint arXiv:1806.07373 (2018).

PITCHES

Dynamic Graph Embedding

Ahmed Rasheed

Abstract

The task of link prediction in graphs spans a wide range of domains such as in document classification in citation networks, online websites, communication networks such as emails, and even in biology such as in protein labeling in protein interaction graphs. Current state-of-the-art models rely on learning per-entity latent representations by mining the whole structure of the relations' graph. However, one major challenge that faces them is that most of the graphs are dynamic which means links emerge or disappear over time.
The aim of this research project is to analyze the current state-of-the-art dynamic graph/network embedding models and investigating potential ways to increase their scalability or prediction performance.

References

- Dynamic Network Embedding : An Extended Approach for Skip-gram based Network Embedding. Lun Du, Yun Wang, Guojie Song, Zhicong Lu, Junshan Wang. IJCAI 2018.
- Attributed Network Embedding for Learning in a Dynamic Environment. Jundong Li, Harsh Dani, Xia Hu, Jiliang Tang, Yi Chang, Huan Liu. CIKM 2017.
- DyRep: Learning Representations over Dynamic Graphs. Rakshit Trivedi, Mehrdad Farajtabar, Prasenjeet Biswal, Hongyuan Zha. ICLR 2019.
- Embedding Temporal Network via Neighborhood Formation. Yuan Zuo, Guannan Liu, Hao Lin, Jia Guo, Xiaoqian Hu, Junjie Wu. KDD 2018.
- Node Embedding over Temporal Graphs. Uriel Singer, Ido Guy, Kira Radinsky. IJCAI 2019.

PITCHES

Inductive learning for TextGCN

Vijaya Krishna

Abstract

Recently TextGCN has been found to be a promising solution for the text classification. However, the algorithm requires complete text data that includes testing data for training the embeddings which makes it transductive. However, in real time it is not possible to obtain test data during training and one needs to develop inductive models. The aim of the project is to develop an inductive model that can learn with mini batches for the text classification

References

- Yao, Liang, Chengsheng Mao, and Yuan Luo. "Graph convolutional networks for text classification." Proceedings of the AAAI Conference on Artificial Intelligence. Vol. 33. 2019.

PITCHES

Loss Curve Forecasting Across Datasets

Hadi Samer Jomaa

Abstract

The objective of a machine learning algorithm is to minimize a defined objective function (loss) in order to achieve good overall performance. However, this can be time consuming, especially when it comes to tuning the hyper-parameters of a model. In this project, you will learn a cross-dataset model which can forecast the training loss of an algorithm given a limited budget (epochs). By using predefined (or trained) meta-features, you will be able to scale-up the approach across other domains (datasets), and be able to determine the best (or top N best) algorithm configurations that might be worth training for the full budget.

References

- van Rijn, Jan N., et al. "Fast algorithm selection using learning curves." International symposium on intelligent data analysis. Springer, Cham, 2015.

PITCHES

Disturb Label Regularization

Mofassir ul Islam Arif

Abstract

The paper suggests a very simple regularization technique and yields very competitive results when compared to the other state of the art model. The simplicity of the idea leads to the question whether this technique can be carried over to other domains. The idea of this project would be to expand Disturb Label and make it compatible with not just computer vision domain but to any other domain that benefits from regularization. This paper also lays out mechanisms for regularizing the loss layer rather than the parameters, which can be further expanded.

References

- Xie, L., Wang, J., Wei, Z., Wang, M., & Tian, Q. (2016). Disturblabel: Regularizing cnn on the loss layer. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 4753-4762).

PITCHES

Personalized Recommendation of Parking Places

Eya Boumaiza

Abstract

Several smart cities around the world have begun monitoring parking areas in order to estimate free spots and help drivers that are looking for parking. Previous works focus only on predicting parking places as a time series problem. In this project, we would like to adapt the idea of integrating context in recommender system to the domain of parking. Instead of predicting the occupancy of parking places given the occupancy history and the user current location and time, we want to generate personalized recommendations of parking places based on driver preferences and history, occupancy of parking places, as well as context information such as the current location and time. The aim of this research project is to investigate the use of different state-of-the-art context-aware recommender systems in the domain of parking.

References

- Yang S., Ma W., Pi X., Sean Q, (2019) A deep learning approach to real-time parking occupancy prediction in spatio-temporal networks incorporating multiple spatio-temporal data sources [Yang S. et al., 2019]
- Context-Aware Recommender Systems in Mobile Scenarios [Woerndl W. et al.,2009]
- Context-aware recommender system: A review of recent developmental process and future research direction [Haruna K. et al., 2017]
- Time-aware recommender systems: a comprehensive survey[Campos et al., 2014]
- Personalized Recommendation Using Time Series Analysis [Zhang et al., 2015]

PITCHES

Distance Metric Learning for identifying similar hyperparameter combinations

Randolf Schulz

Abstract

Distance Metrics play a crucial role in many Machine Learning applications such as clustering or distance based models such as K nearest neighbors. However it is observed that high dimensional data often is not distributed uniformly in the feature space, but rather embedded on a lower dimensional submanifold. Consequently a metric should be chosen which respects the non-linear structure of said submanifold. To do this location dependent metrics need to be introduced (as opposed to metrics which only depend on the vector difference such as dist(x,y) = ||x-y||_p). In NeurIPS 2019 a novel framework called Curvilinear Distance Metric Learning was introduced which facilitates location dependent metrics. Based on this paper one possible application of distance metric learning is the idea of learning similarities between hyperparameters. Usually, changing a single hyperparameter has an easy to understand effect on the behaviour of the model whereas the effect of changing multiple hyperparameters simultaneously is not well understood. By learning an appropriate curvilinear distance metric we hope to uncover similarities between hyperparameter combinations that help us guide the optimal hyperparameter search.

References

- Curvilinear Distance Metric Learning
  https://papers.nips.cc/paper/8675-curvilinear-distance-metric-learning