To William Demant Fonden

Project plan for external research stay in MLG, Cambridge

To Whom it May Concern,

For my external research stay I have ideas for combining deep generative models (DGM), uncertainty estimation, multi-task learning and active learning (AL) for molecular exploration and detection through Raman spectra as a molecular fingerprint. This is mainly inspired by the work of my host-supervisor, Assistant Professor, José Miguel Hernández-Lobato, on ChemVAE, GVAE, BNN-LV, PBP, the recent work on Deep Learning Spectroscopy by my supervisor, Associate Prof. Mikkel N. Schmidt and my own work w. Prof. Ole Winther developing a Gaussian-mixture-VAE for gene expressions and exploration of cell types in latent space. These ideas lead to Miguel, my supervisors and me agreeing on and planning a 4 months research project in the Machine Learning Group, University of Cambridge, which now needs additional external funding from you.

4 June 2019
MV

The project for this 4 months external stay is a part of my PhD project in Active Deep Learning for Nano-Sensor Systems, which is described in detail below, and will therefore be focused on developing deep learning models for Raman spectroscopy, which is the most commonly used technique for detection of molecules. Therefore I will here motivate Raman spectroscopy and explain why a data-driven approach will benefit the research and industry in pharmaceutical engineering.

The first research question is: Can determine and detect the fingerprint of the building blocks of nature itself in a smarter more data efficient way? The hypothesis is: We can use deep neural networks for learning the underlying process behind how light is scattered by linking the Raman spectra and various representations of molecular structure. Learning how to generate the fingerprint of molecules, will make it possible to detect the molecules, their concentrations and properties as medicine contained in micro-containers by non-invasive methods, e.g. lasers.

Raman scattering stems from driven molecular vibration coming from symmetric and asymmetric stretching of bonds between atoms in molecules, when the incident laser hits the molecule. The intensity is below 0.01% of the direct Raleigh scattered light, so it needs filtering and the signal-to-noise-ratio is low, which also argues for using Surface-Enhanced Raman Spectroscopy (SERS) or Coherent anti-Stokes Raman scattering spectroscopy (CARS). SERS leads to a very complex behaviour and makes it hard to simulate with DFT. The frequency of the Raman scattered light (Raman bands) will depend on the strength of the atomic bonds and atomic masses and can be modelled in the time domain with an ordi-

DTU Compute
Department of Applied Mathematics
Computer Science
www.compute.dtu.dk

Building 321, Room 106
Richard Petersens Pl,
2800 Kgs. Lyngby, Denmark

Dir. +45 61603013
Tel.. +45 45 25 30 31 and
maxvo@dtu.dk

nary differential equation like Hooke's law and Newton's 2nd law of motion. Through quantum mechanical equations the differential functional theory (DFT) can be used to simulate Raman spectra, but for complicated molecular compounds and environment like SERS, the simulation rarely fits with real world measurements and will often be re-adjusted manually to fit this. This suggests using a data-driven approach by generating spectra through deep graph neural networks, which are commonly used for linking molecular properties to the molecular structure.
Here is the representations of molecules and what we need for predicting molecular properties and Raman spectra. We need both molecular structure, atomic mass and band strength to determine vibrational modes and differential equations, so here is a list of these commonly used measures:

- SMILES (Text representation of molecules)
- Coulomb matrix (Energy interaction based on distances between all atoms)
- Bag of bonds
- Histograms
- Radial distribution functions
- Chemical environment
- ACSF

We can find these resources at:

- https://www.researchgate.net/post/Free_Database_with_Raman_spectra
- https://www.ccdc.cam.ac.uk/solutions/csd-system/components/csd/
- https://serc.carleton.edu/research_education/crystallography/xldatabases.html
- http://oqmd.org/
- http://quantum-machine.org/datasets/

At the IDUN Center of Excellence, researchers can as well provide huge amounts of data for our project through their ongoing experiments with Raman spectroscopy on micro-container drug-delivery. With computational resources at both University of Cambridge and DTU Compute, we can scale up to big data and thereby the representational strength of our deep learning models tremendously.

For this project the goal is for me to be the main author on two publications in high-impact conferences or journals and will be conducted in collaboration with:

Host-supervisor:	José Miguel Hernández-Lobato - Assistant Professor
Principal supervisor:	Mikkel Nørgaard Schmidt - Associate Professor
Co-supervisor:	Tommy Sonne Alstrøm - Senior Researcher

Best regards,

Maximillian Fornitz Vording
PhD Student
Section for Cognitive Systems

Study Plan - Project Description

PhD Project:
Active Deep Learning for Nano-Sensor Systems

Maximillian Fornitz Vording (maxvo@dtu.dk)
27 September 2018

Institution:
The Technical University of Denmark

Principal supervisor:
Associate Professor, Mikkel N. Schmidt

Co-supervisor:
Senior Researcher, Tommy Sonne Alstrøm

Introduction

This 3-year PhD project starting August 1st 2017 (2017.Q3) at DTU Compute's Section for Cognitive Systems is a part of IDUN center of excellence, which is a multi-disciplinary project lead by DTU Nanotech with participation of DTU Compute and The University of Copenhagen, and funded by the Danish National Research Foundation and the Villum Foundation. The center is divided into IDUN Drug focusing on drug delivery and IDUN Sensor focusing on nano-mechanical sensors.

The PhD project is associated with IDUN Sensor research that focuses on development and exploration of nano-mechanical biosensors. Data processing and modelling are indispensable multipurpose tools for sensor development and evaluation and analysis of results in demonstration activities [14, 1]. In sensor development, data modelling tools provide other views and insight into the physical and chemical properties of the sensor as well as sensing principle; hence, improving sensor development in terms of time-use, but also the ability to robustly confirm hypotheses about the sensor's functionality. In relation to sensor demonstration activities, data modelling is important for obtaining robust sensor performance by suppression of noise caused by undesired physical and chemical properties of the sensor as well as uncontrollable experimental factors.

We are living in the information age, where data is the raw material. But access to data can be sparse and expensive, especially when human annotation and expert supervision is needed in the data collection process [12, 31]. Therefore optimising data efficiency in machine learning and developing optimal experimental designs is the two methodological objectives in this PhD project. The PhD project aim to prove that combining active and deep learning can provide:

- Efficient ways of interactive usage, understanding, training and integration of deep learning models in experiments for both researchers and common users.
- A more friction- and effortless teacher-student interaction so as to speed up online learning in everyday life and lower the expenses on data collection.
- Impact on the information society as a whole with more personalised, adjusting and interacting digital devices.

Background

Recent advances within the deep learning research field has shown remarkable performance in a great variety of data processing tasks. The PhD project will focus on developing new active learning methods [28] for deep neural network models [10]. Such methods can provide optimal experimental design and hypothesis testing for sensor development, and further reduce the need for user labels in connection with demonstration of detection sensing capabilities. The methodological objectives and research relates to, and will leverage from, current advances in Bayesian optimization [23, 4, 22]; Bayesian experimental design [30]; reinforcement learning [15, 16]; curiosity-driven exploration [13, 24]; one-shot-learning [7]; unsupervised and semi-supervised generative models [18, 19, 20, 21, 29, 3, 11, 26]; human-in-the-loop [31, 25, 8, 5], where the user is the sensor developer and/or an end-user providing labeled information.

Hypothesis

The main hypothesis driving the project is that generalisation performance, robustness and data effiency of machine learning can be enhanced increasingly with more intelligent data collection and interaction with human annotators through a general purpose active deep learning framework. This framework will integrate Bayesian optimisation and reinforcement learning methods for optimising the sequence of chosen state-of-the-art active learning techniques, new features and unlabeled samples added to the training set. An underlying hypothesis supporting the general purpose active deep learning framework is that, more data efficient acquisition functions can be obtained through better uncertainty estimators, e.g Monte Carlo dropout [9, 10], in deep deterministic and bayesian neural networks. Learned latent representations and posterior class probabilities in semi-supervised generative models [18, 19, 20, 21, 29] can be used for more efficient acquisition functions utilizing both unlabeled and labeled data. These latent represenatations and intepretable model inspection through sensisivity maps and layer-wise relevance propagation [27, 2] can provide more interpretable data visualisation for interaction with human users in the labeling process. Exploring all of these hypotheses in the delivery of a general purpose tool for automated optimal experimental design will be the overarching goal.

Objectives and deliverables

The novel methods, combinations and applications in the active deep learning framework will be reviewed during the first 3 quarters (2017.Q3-Q4 and 2018.Q1) of the project while also solving the task of cell image segmentation using active learning [28] with bayesian convolutional neural networks [10]. As proof of concepts and hypothesis, the methods and application found in the literature review will be implemented for handling real-world problems and experiments in the research field of developing nano-sensor systems. In the quarters 2018.Q1-Q4, applications of the methods for automatic molecule detection using Raman spectroscopy and in the design of morphological and chemical properties of micro-containers for optimal direct drug delivery will be investigated. I will arrange for a an external stay at the Machine Learning Group, University of Cambridge, in 2019.Q1-Q2, to work with Prof. Carl Edward Rasmussen on applying active learning in Bayesian models, Gaussian processes and Semi-supervised Gaussian mixture variational auto-encoders [17, 21, 6], with focus on propagating uncertaintanty in Bayesian neural networks [10]. These concepts will be implemented and maintained in the experimental frameworks during 2019.Q1-Q4, so the researchers themselves can control and monitor the experiments, and provide data for improving the models

in the most efficient active learning loop. In the last three quarters 2020.Q1-Q3 the resulting papers from the application of active deep learning in the various nano-sensor and drug-delivery systems will be combined into the thesis evaluating these in the scope of the overall hypothesis.

References

[1] Alstrøm, T. S., Schmidt, M. N., Rindzevicius, T., Boisen, A., and Larsen, J. (2017). A pseudo-voigt component model for high-resolution recovery of constituent spectra in raman spectroscopy. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE.

[2] Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.-R., and Samek, W. (2015). On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS ONE*, **1**0(7), e0130140.

[3] Bowman, S. R., Vilnis, L., Vinyals, O., Dai, A. M., Jozefowicz, R., and Bengio, S. (2015). Generating sentences from a continuous space. *arXiv preprint*.

[4] Calandra, R., Seyfarth, A., Peters, J., and Deisenroth, M. P. (2016). Bayesian optimization for learning gaits under uncertainty. *Annals of Mathematics and Artificial Intelligence*, **7**6(1), 5–23.

[5] Cranor, L. F. (2008). A framework for reasoning about the human in the loop. *UPSEC*, **8**(2008), 1–15.

[6] Dilokthanakul, N., Mediano, P. A. M., Garnelo, M., Lee, M. C. H., Salimbeni, H., Arulkumaran, K., and Shanahan, M. (2016). Deep unsupervised clustering with gaussian mixture variational autoencoders. *arXiv preprint*.

[7] Fei-Fei, L., Fergus, R., and Perona, P. (2006). One-shot learning of object categories. *IEEE transactions on pattern analysis and machine intelligence*, **2**8(4), 594–611.

[8] Fung, P. T., Norgate, G., Dilts, T. A., Jones, A. S., and Ravindran, R. (1992). Human-in-the-loop machine control loop. US Patent 5,116,180.

[9] Gal, Y. (2016). *Uncertainty in Deep Learning*. Ph.D. thesis, Department of Engineering - University of Cambridge.

[10] Gal, Y., Islam, R., and Ghahramani, Z. (2017). Deep bayesian active learning with image data. *CoRR*, **a**bs/1703.02910.

[11] Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. *arXiv preprint*.

[12] Honnibal, M. and Montani, I. (2017). Prodigy: A new tool for radically efficient machine teaching.

[13] Houthooft, R., Chen, X., Duan, Y., Schulman, J., De Turck, F., and Abbeel, P. (2016). Curiosity-driven exploration in deep reinforcement learning via bayesian neural networks. *arXiv preprint arXiv:1605.09674*.

[14] Jaumot, J., Gargallo, R., de Juan, A., and Tauler, R. (2005). A graphical user-friendly interface for mcr-als: a new tool for multivariate curve resolution in matlab. *Chemometrics and Intelligent Laboratory Systems*, **7**6, 101–110.

[15] Kaelbling, L. P., Littman, M. L., and Moore, A. W. (1996). Reinforcement learning: A survey. *Journal of artificial intelligence research*, **4**, 237–285.

[16] Karpathy, A. (2016). Deep reinforcement learning: Pong from pixels.

[17] Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint*.

[18] Kingma, D. P. and Welling, M. (2013). Auto-encoding variational Bayes. *arXiv preprint*.

[19] Kingma, D. P., Rezende, D. J., Mohamed, S., and Welling, M. (2014). Semi-supervised learning with deep generative models. *arXiv preprint*.

[20] Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. (2016). Improving variational inference with inverse autoregressive flow. *arXiv preprint*.

[21] Maaløe, L., Fraccaro, M., and Winther, O. (2017). Semi-supervised generation with cluster-aware generative models. *arXiv preprint*.

[22] Martinez-Cantin, R., de Freitas, N., Brochu, E., Castellanos, J., and Doucet, A. (2009). A bayesian exploration-exploitation approach for optimal online sensing and planning with a visually guided mobile robot. *Autonomous Robots*, **2**7(2), 93–103.

[23] Mockus, J. (1989). *Bayesian approach to global optimization: theory and applications*. Mathematics and its applications (Kluwer Academic Publishers).: Soviet series. Kluwer Academic.

[24] Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. (2017). Curiosity-driven exploration by self-supervised prediction. *arXiv preprint arXiv:1705.05363*.

[25] Pretlove, J. and Skourup, C. (2007). Human in the loop. *ABB Review*, **1**, 6–10.

[26] Salimans, T., Karpathy, A., Chen, X., and Kingma, D. P. (2017). PixelCNN++: Improving the PixelCNN with discretized logistic mixture likelihood and other modifications. *arXiv preprint*.

[27] Samek, W., Wiegand, T., and Müller, K.-R. (2017). Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models. *ITU Journal: ICT Discoveries - Special Issue 1 - The Impact of Artificial Intelligence (AI) on Communication Networks and Services*, **1**, 1–10.

[28] Settles, B. (2012). *Active Learning*. Morgan & Claypool.

[29] Sønderby, C. K., Raiko, T., Maaløe, L., Sønderby, S. K., and Winther, O. (2016). Ladder variational autoencoders. *arXiv preprint*.

[30] Vanlier, J., Tiemann, C. A., Hilbers, P. A. J., and van Riel, N. A. W. (2012). A bayesian approach to targeted experiment design. *Bioinformatics*, 28(8), 1136–1142.

[31] Wang, Y. (2016). The power of human-in-the-loop: Combine human intelligence with machine learning.