Spectralligence

Spectral Analysis in life sciences and materials sciences through Artificial Intelligence

Deliverable 1.2 – Public Version

AI and Spectroscopy: needs and deep learning network candidates

Project Coordinator	Johan van den Brink, Philips Medical Systems B.V.		
Start date Project	November 1st 2021	**Duration**	36 months
Version	1.0		
Status	Final		
Date of issue	22/02/2024		
Dissemination level	Public		

Authors' data

Author	Beneficiary
Final editor's address	Johan van den Brink Philips Medical Systems B.V. Veenpluis 6 5684PC Best / Netherlands

Table of Contents

1. Introduction

1.1. Purpose of the Document

In this document an overview will be given of the candidate deep learning (DL) networks that are suitable for spectral analysis. Within Spectralligence, multiple spectral domains are investigated by multiple partners. For this deliverable, three spectral domains are considered by the corresponding partners (see Table 1). The final aim is to combine the DL candidates of all domains and to summarize similarities and cross-domain interests.

Table 1: Overview of the spectral domains and the corresponding partners who investigated the DL candidates.

Spectral Domain	Abbreviation	Partners
(Neutron Induced) Gamma Ray Spectroscopy	GRS	Dynaxion
Atomic Emission Spectroscopy	AES	Sensmet
Magnetic Resonance Spectroscopy	MRS	Philips and TU/e

1.2. Spectroscopy Workflow

To facilitate a general framework of summarizing DL candidates per domain, an overall spectroscopy workflow is constructed. This workflow will apply to all spectral domains and will enable the comparison of all domains and the categorization of the type of DL candidate. Figure 1 shows the spectroscopy workflow that is used for this document. Different steps in the spectroscopy workflow are accompanied by various machine learning (ML) applications. Per domain, all possible DL candidates are categorized with the use of this workflow.

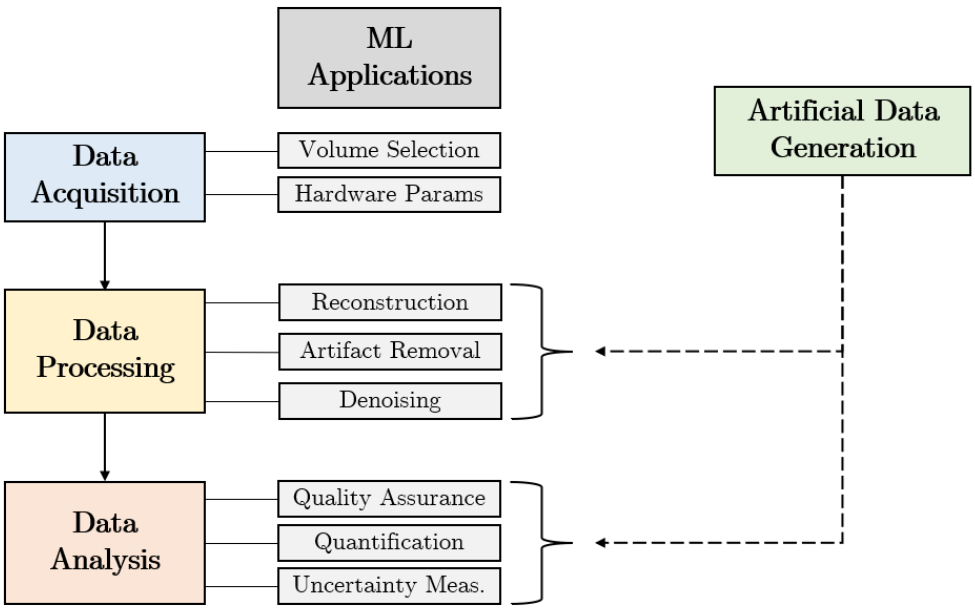

Figure 1: Schematic representation of the overall spectroscopy workflow. The workflow is divided in Data Acquisition, Data Processing and Data Analysis. Each block has its own machine learning (ML) applications, which are indicated in the grey boxes. An additional block for Artificial Data Generation is added to represent possible ML methods that focus on generating synthetic data for training the ML applications.

2. Spectral Domains and Needs

All materials have characteristic energy states reflecting atomic or molecular structures, states, vibrational modes and interactions. Spectroscopy is the domain of science and technology focused on probing the distribution of these energy states, either actively or passively, with a wealth of sensors and embedded systems, varying from micro devices and Lab-on-Chip to CERN's accelerators and radio telescopes.

Table 1 provides an overview of different types of spectroscopy that are in common use in pharmaceutical and (bio-)chemical labs. The referenced webpage provides a very basic overview of workflow steps for each modality.

Table 2: Overview of types of spectroscopy (https://microbenotes.com/types-of-spectroscopy/)

1. Absorption spectroscopy
2. Astronomical spectroscopy
3. Atomic absorption spectroscopy
4. Circular dichroism spectroscopy
5. Electrochemical impedance spectroscopy (EIS)
6. Electron spin resonance (ESR) spectroscopy
7. Emission spectroscopy
8. Energy dispersive spectroscopy
9. Fluorescence spectroscopy
10. Fourier-transform infrared (FTIR) spectroscopy
11. Gamma-ray spectroscopy
12. Infrared (IR) spectroscopy/ Vibrational spectroscopy
13. Magnetic resonance spectroscopy
14. Mass spectroscopy
15. Molecular spectroscopy
16. Mossbauer spectroscopy
17. Nuclear magnetic resonance (NMR) spectroscopy
18. Photoelectron spectroscopy
19. Raman spectroscopy
20. UV spectroscopy
21. Ultraviolet and visible (UV/Vis) spectroscopy
22. X-ray photoelectron spectroscopy

References

Sources

The detected electromagnetic waves range from radio (kHz and MHz, including NMR and MRI) and radar (GHz) through (near-)optical (THz, UV/VIS/IR, Raman) to hard radiation (X-ray, Neutron).

Applications cover military, space exploration, GIS (hyperspectral) for pollution and crops monitoring, chemistry and pharma, biology and medicine, environmental monitoring, and others. Advances in sensor miniaturization and energy efficiency have broadened its application to beyond the lab, for example through multi-spectral optical imaging from satellites, planes and drones.

The general challenge in all of these applications is to provide ultra-simple procedures to accurately detect, classify and quantify of signals of components in mixtures, and from very noise signals.

To address these challenges, several groups have engaged in including AI/ML for denoising and classification tasks across the full spectrum.

Several recent reviews (for example those listed below) cover the scientific potential of AI/ML in spectral analysis, and some products are available based on physics-informed spectral training and analysis (e.g. Bruker's high-resolution NMR spectroscopy SW package). Translation from scientific progress with generative AI towards validated real-world applications is an active field of research, as also recognized by IAEA (2022) in their white paper on Artificial Intelligence for Accelerating Nuclear Applications, Science and Technology. Further investments in miniaturized sensing, on-chip analysis, and reduced expert-dependence are necessary to further harness the robustness and validation of AI/ML in real-world settings with strict QA/QC settings. The most significant challenge is training models in the absence of ground-truth big data, and providing proof that models trained on simulated or synthetic data operate accurately in the target application.

The references below provide an examplary set of recent overview papers for several spectral domains:

[1] Artificial Intelligence in Analytical Spectroscopy (2023), Part I: Basic Concepts and Discussion (spectroscopyonline.com)

[2] Artificial Intelligence in Analytical Spectroscopy (2023), Part II: Examples in Spectroscopy (spectroscopyonline.com)

[3] A review of machine learning applications for the proton MR spectroscopy workflow - Sande - 2023 - Magnetic Resonance in Medicine - Wiley Online Library

[4] Current advances in imaging spectroscopy and its state-of-the-art applications – (2024) ScienceDirect

[5] Zajnulina (2022).pdf (arxiv.org): Advances of Artificial Intelligence in Classical and Novel Spectroscopy[1]Based Approaches for Cancer Diagnostics. A Review

[6] Deep Learning for Biospectroscopy and Biospectral Imaging: State-of-the-Art and Perspectives | Analytical Chemistry (acs.org)

[7] Application of hyperspectral imaging systems and artificial intelligence for quality assessment of fruit, vegetables and mushrooms: A review - ScienceDirect

[8] Ozdemir (2020) Deep Learning Applications for Hyperspectral Imaging: A Systemic Review.pdf (iecscience.org)

[9] Biomolecular NMR spectroscopy in the era of artificial intelligence – (2023) ScienceDirect

[10] AtomAI framework for deep learning analysis of image and spectroscopy data in electron and scanning probe microscopy (2022) Nature Machine Intelligence

3. Deep Learning Candidates per Spectral Domain

3.1. Introduction

In this section, we provide the DL network candidates that will be further investigated for the identified workflow steps for the spectral domains covered in the Spectralligence project. For each step in the spectroscopy workflow relevant DL architectures will be summarized and categorized. At the end, a domain-specific conclusion will be given.

3.2. (Neutron Induced) Gamma Ray Spectroscopy

This overview is created by Dynaxion.

3.2.1. Data Acquisition

Application	DL Architecture	Explanation	Reference
Volume selection	-	We will get information from first level screening about the volume to be scanned	-
Hardware params	-	We use simulations to predict which hardware parameters are optimal, for example: timing window.	-

3.2.2. Data Processing

Application	DL Architecture	Explanation	Reference
Denoising	-	Using Savitzky Golay filter	-
			-

3.2.3. Data Analysis

Application	DL Architecture	Explanation	Reference
Quantification	CNN	After converting our neutron induced gamma spectra to suitable format, Google-net is used in combination with transfer-learning to classify materials.	[1]
	Custom Autoencoder	The Custom Autoencoder and models is used to estimate the atomic composition from the neutron induced gamma spectra	[2]

3.2.4. Artificial Data Generation

Application	DL Architecture	Explanation	Reference
Data Augmentation	ReVAE	ReVAE model is used for conditional data generation. This model can generate the spectra for the different materials and combinations. This network is trained on the neutron induced gamma spectra.	[3]

3.2.5. Conclusion

An accurate deep learning classification network is being developed by Dynaxion for the classification of different materials. This network extends and harnesses their proof of concept, where 16 different materials were measured, and later accurately classified by networks trained on both Monte-Carlo simulation data, and measurement data. Furthermore, Dynaxion is developing the custom autoencoder model to accurately estimate the atomic composition from measured spectra of these 16 different materials. Currently, Dynaxion is also exploring Mixture of Experts Ensemble models [4] to estimate the atomic composition from the neutron induced gamma spectra.

In addition, Dynaxion has experimented with three different types of Variation Autoencoders, namely Variational Autoencoder, Ladder – Variational Autoencoder, and Reparametrized VAE (ReVAE) for artificial data generation. The data generated by the first two networks was not in line with the neutron induced gamma spectra. So, Dynaxion moved to ReVAE for generating the artificial data to expand the training data. This network will be developed further for use in generating data that matches with the properties of neutron induced gamma spectra of additional materials.

Dynaxion does not apply any other ML applications that are connected to the general framework sketched in Figure 1.

3.2.6. References

[1] E. T. Moore, J. L. Turk, W. P. Ford, N. J. Hoteling, and L. S. McLean, “Transfer Learning in Automated Gamma Spectral Identification.” arXiv, Mar. 23, 2020. doi: 10.48550/arXiv.2003.10524.

[2] R. Dong *et al.*, “DeepXRD, a Deep Learning Model for Predicting XRD spectrum from Material Composition,” *ACS Appl. Mater. Interfaces*, vol. 14, no. 35, pp. 40102–40115, Sep. 2022, doi: 10.1021/acsami.2c05812.

[3] T. Joy, S. Schmon, P. Torr, S. Narayanaswamy, and T. Rainforth, *Rethinking Semi-Supervised Learning in VAEs*. 2020.

[4] Rokach, L. (2010). *Pattern Classification Using Ensemble Methods*. World Scientific.

3.3. Atomic Emission Spectroscopy

This overview is created by Sensmet.

3.3.1. Data Acquisition

Application	DL Architecture	Explanation	Reference
Volume selection	-	Volume selection is not applicable to AES of liquid samples as the samples are homogenous.	-
Hardware params	Two-headed	Intermediate input layer to enhance quantification by inputting device hardware parameters such as sample conductivity, temperature, etc.	-

3.3.2. Data Processing

Application	DL Architecture	Explanation	Reference
Reconstruction	Fully connected	Feature selection by applying fully connected neural network to optimize emission lines to be used.	[1], [2]
Artifact removal	-	DL not being applied at the moment.	-
Denoising	-	DL not being applied at the moment.	-

3.3.3. Data Analysis

Application	DL Architecture	Explanation	Reference
Quality assurance	RNN	Recurrent neural networks (RNNs) to detect anomalies during device operation.	[1]
Quantification	CNN, two-headed	Quantitative calculation model that extracts information from emission spectra and incorporates it with device parameters to calculate concentrations with minimized matrix dependence.	[1]
Uncertainty Meas.	-	Real-time definition of measurement uncertainty by applying statistical process control according to laboratory guidelines. Future possibility to apply recurrent neural networks to detect measurement anomalies and raise measurement uncertainty flags.	[3]

3.3.4. Artificial Data Generation

Application	DL Architecture	Explanation	Reference
Data Augmentation	-	DL not being applied at the moment.	-

3.3.5. Conclusion

A deep learning network is being developed by Sensmet for the calibration of metal concentration measurements of liquid samples employing online AES. The network has been first used in the proof of concept, where two different metal concentrations were monitored in process liquids. The quantitative two-headed CNN was used to process raw emission spectra in combination with device operating parameters to produce quantitative concentration results of the sample metal concentrations. Furthermore, Sensmet has also experimented with simpler fully connected neural network that was used with existing traditional calculation methods to correct the concentration results for matrix effects. In the other mentioned processing steps, Sensmet does not use ML methods.

3.3.6. References

[1] C. Sun *et al.*, “Machine Learning Allows Calibration Models to Predict Trace Element Concentration in Soils with Generalized LIBS Spectra,” *Sci. Rep.*, vol. 9, no. 1, Art. no. 1, Aug. 2019, doi: 10.1038/s41598-019-47751-y.

[2] L. Zou *et al.*, “Online simultaneous determination of H2O and KCl in potash with LIBS coupled to convolutional and back-propagation neural networks,” *J. Anal. At. Spectrom.*, vol. 36, no. 2, pp. 303–313, Feb. 2021, doi: 10.1039/D0JA00431F.

[3] B. Magnusson, T. Näykki, H. Hovind, and M. Krysell, “Handbook for Calculation of Measurement Uncertainty in Environmental Laboratories,” *Nord. TR 537 Ed. 31*, Jan. 2012.

3.4. Magnetic Resonance Spectroscopy

This overview is created by Eindhoven University of Technology (TU/e) and Philips

3.4.1. Data Acquisition

Application	DL Architecture	Explanation	Reference
Volume selection	CNN	A CNN was used to segment MRI images. An objective function used the segmentations to predict the best voxel placement. Similar networks are already available at Philips for automatic accurate placement of imaging slabs, based on a fast scout scan.	-
Hardware params.	Ensemble of CNNs + MLP	Orthogonal shim values (x,y,z) were predicted using an ensemble of CNNs that are combined with the use of averaging or by adding an MLP. The inputs were raw H-FID signals.	-

3.4.2. Data Processing

Application	DL Architecture	Explanation	Reference
Reconstruction	LSTM	An LSTM network was proposed taking (under-sampled) sparse time-domain signals and corresponding sampling schedule as input and outputting the reconstructed time-domain FID.	[1]
	U-Net	A U-Net architecture was trained with truncated FIDs to reconstruct fully sampled spectra.	[2]
Artifact removal	MLP, CNN	Both MLP and CNN architectures have been applied to correctly predict phase and frequency offsets of an input spectrum in a supervised regression.	[2] [3]
	MLP, CNN, Auto-encoder,	An MLP and CNN were proposed for identification of ghosting artifacts (spurious echoes) by means of classification. Also, an auto-encoder network was proposed to identify and correct ghosting artifacts within spectra, by training with (synthetically) contaminated spectra as input and clean spectra as ground truths.	[4]
	U-Net	A U-Net was trained with low resolution MRSI data as input and high resolution data as ground truth.	[5]
Denoising	LSTM,	An LSTM network is trained in supervised fashion with low average spectra to predict high average (therefore higher SNR) spectra.	[6]
	Auto-encoder	Another work used an auto-encoder network.	[7]

3.4.3. Data Analysis

Application	DL Architecture	Explanation	Reference
Quality assurance	CNN	ML methods are trained to classify (by experts) labeled MRS spectra based on their quality. A CNN architecture as well as multiple ML methods (SVM, LDA, RUSBoost in combination with ICA and PCA feature extractors) were proposed.	[8]
	SVM, LDA, RUSBoost		[9]
Quantification	Random forest, CNN,	Quantification via supervised regression taking spectra and outputting metabolite concentrations have been proposed using random forest and CNN architectures.	[10] [11]
	CNN encoder and physics-model decoder,	A CNN encoder and physics-model decoder was proposed, taking the real part of a spectrum as input, deconvolving it into signal parameters which in turn are reconstructed by the physics models into spectra, allowing self-supervised training.	[12]
	Multiple CNNs, U-Nets	Per metabolite a CNN was trained to output a metabolite-only spectrum from the real part of a processed spectrum. Another work utilized U-Nets for the same setup.	[13]
Uncertainty Meas.	CNN	A CNN-based approach to quantify metabolite concentrations and obtain an uncertainty estimate was proposed. By constructing a 3D space of the quantitative errors for each target metabolite as a function of the SNR, linewidth, and SBR.	[14]
		Further Bayesian CNN methods used Monte Carlo dropout for uncertainty estimation of the NN.	[15] [16]

3.4.4. Artificial Data Generation

Application	DL Architecture	Explanation	Reference
Data Generation	(DC) GAN	Different GAN models were trained to generate H-MRS spectra for healthy, low-grade and high-grade tissue based on a real clinical dataset.	[17]

3.4.5. Integrated Analysis and Visualization

Application	DL Architecture	Explanation	Reference
Clinical Workflow Support	multiple	A comprehensive integration of AI/ML based pre-processing, classification and analysis is provided for MR Spectroscopy, from a competing university.	[18]

3.4.6. Conclusion

TU/e and Philips, in collaboration with MUMC+ and UMCU, are actively monitoring developments for AI/ML in human MR Spectroscopy. We found that different model architectures are applied at different stages in the workflow. CNN architectures generally found application for tasks involving feature extraction, such as convolving spectra into parameters of interest. MLPs have been used for such scenarios, though, CNNs seem to outperform them. U-Nets and Autoencoders were mainly applied in scenarios mapping spectral shapes to spectral shapes, or to increase dimensionality. Overall, NNs have

been implemented and trained mainly in a black-box manner, replacing specific problem areas as a whole. Further extensions of classification networks with uncertainty prediction and explainability, e.g., feature visualization are also under investigation.

3.4.7. References

[1] D. F. Hansen, “Using Deep Neural Networks to Reconstruct Non-uniformly Sampled NMR Spectra,” *J. Biomol. NMR*, vol. 73, no. 10, pp. 577–585, Nov. 2019, doi: 10.1007/s10858-019-00265-1.

[2] S. Tapper *et al.*, “Frequency and phase correction of J-difference edited MR spectra using deep learning,” *Magn. Reson. Med.*, vol. 85, no. 4, pp. 1755–1765, Apr. 2021, doi: 10.1002/mrm.28525.

[3] D. J. Ma *et al.*, “MR spectroscopy frequency and phase correction using convolutional neural networks,” *Magn. Reson. Med.*, vol. 87, no. 4, pp. 1700–1710, Apr. 2022, doi: 10.1002/mrm.29103.

[4] S. P. Kyathanahally, A. Döring, and R. Kreis, “Deep learning approaches for detection and removal of ghosting artifacts in MR spectroscopy,” *Magn. Reson. Med.*, vol. 80, no. 3, pp. 851–863, 2018, doi: 10.1002/mrm.27096.

[5] Z. Iqbal, D. Nguyen, G. Hangel, S. Motyka, W. Bogner, and S. Jiang, “Super-Resolution 1H Magnetic Resonance Spectroscopic Imaging Utilizing Deep Learning,” *Front. Oncol.*, vol. 9, p. 1010, Oct. 2019, doi: 10.3389/fonc.2019.01010.

[6] W. Hu *et al.*, “Denoising Single Voxel Magnetic Resonance Spectroscopy with Deep Learning on Repeatedly Sampled In Vivo Data,” *ArXiv210111442 Phys.*, Jan. 2021, Accessed: Jan. 31, 2022. [Online]. Available: http://arxiv.org/abs/2101.11442

[7] Y. Lei, B. Ji, T. Liu, W. J. Curran, H. Mao, and X. Yang, “Deep learning-based denoising for magnetic resonance spectroscopy signals,” in *Medical Imaging 2021: Biomedical Applications in Molecular, Structural, and Functional Imaging*, Online Only, United States, Feb. 2021, p. 3. doi: 10.1117/12.2580988.

[8] S. S. Gurbani *et al.*, “A convolutional neural network to filter artifacts in spectroscopic MRI,” *Magn. Reson. Med.*, vol. 80, no. 5, pp. 1765–1775, 2018, doi: 10.1002/mrm.27166.

[9] S. P. Kyathanahally *et al.*, “Quality of clinical brain tumor MR spectra judged by humans and machine learning tools,” *Magn. Reson. Med.*, vol. 79, no. 5, pp. 2500–2510, 2018, doi: 10.1002/mrm.26948.

[10] D. Das, E. Coello, R. F. Schulte, and B. H. Menze, “Quantification of Metabolites in Magnetic Resonance Spectroscopic Imaging Using Machine Learning,” in *Medical Image Computing and Computer Assisted Intervention − MICCAI 2017*, Cham, 2017, pp. 462–470. doi: 10.1007/978-3-319-66179-7_53.

[11] N. Hatami, M. Sdika, and H. Ratiney, “Magnetic Resonance Spectroscopy Quantification Using Deep Learning,” in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, Cham, 2018, pp. 467–475. doi: 10.1007/978-3-030-00928-1_53.

[12] S. S. Gurbani, S. Sheriff, A. A. Maudsley, H. Shim, and L. A. D. Cooper, “Incorporation of a spectral model in a convolutional neural network for accelerated spectral fitting,” *Magn. Reson. Med.*, vol. 81, no. 5, pp. 3346–3357, 2019, doi: 10.1002/mrm.27641.

[13] H. H. Lee and H. Kim, “Intact metabolite spectrum mining by deep learning in proton magnetic resonance spectroscopy of the brain,” *Magn. Reson. Med.*, vol. 82, no. 1, pp. 33–48, 2019, doi: 10.1002/mrm.27727.

[14] H. H. Lee and H. Kim, “Deep learning-based target metabolite isolation and big data-driven measurement uncertainty estimation in proton magnetic resonance spectroscopy of the brain,” *Magn. Reson. Med.*, vol. 84, no. 4, pp. 1689–1706, 2020, doi: 10.1002/mrm.28234.

[15] H. H. Lee and H. Kim, “Bayesian deep learning–based 1H-MRS of the brain: Metabolite quantification with uncertainty estimation using Monte Carlo dropout,” *Magn. Reson. Med.*, vol. 88, no. 1, pp. 38–52, Jul. 2022, doi: 10.1002/mrm.29214.

[16] R. Rizzo, M. Dziadosz, S. P. Kyathanahally, M. Reyes, and R. Kreis, “Reliability of Quantification Estimates in MR Spectroscopy: CNNs vs Traditional Model Fitting,” in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*, vol. 13438, L. Wang, Q. Dou, P. T. Fletcher, S. Speidel, and S. Li, Eds. Cham: Springer Nature Switzerland, 2022, pp. 715–724. doi: 10.1007/978-3-031-16452-1_68.

[17] N. Olliverre, G. Yang, G. Slabaugh, C. C. Reyes-Aldasoro, and E. Alonso, “Generating Magnetic Resonance Spectroscopy Imaging Data of Brain Tumours from Linear, Non-linear and Deep Learning Models,” in *Simulation and Synthesis in Medical Imaging*, Cham, 2018, pp. 130–138. doi: 10.1007/978-3-030-00536-8_14.

[18] X. Chen, J. Li, D. Chen, Y. Zhou, Z. Tu, M. Lin, T. Kang, J. Lin, T. Gong, L. Zhu, J. Zhou, O. Lin, J. Guo, J. Dong, D. Guo and X. Qu, “CloudBrain-MRS: An intelligent cloud computing platform for in vivo magnetic resonance spectroscopy preprocessing, quantification, and analysis”, Journal of Magnetic Resonance, Volume 358, 2024,https://doi.org/10.1016/j.jmr.2023.107601.

4. Overall Conclusion

4.1.1. Workflow Focus

Considering all the three spectral domains and partners of Spectralligence, there are some differences between the current focus areas of the DL methods for spectral analysis. For GRS, Dynaxion is currently focusing on classification and artificial data generation using CNNs and VAEs respectively. Meanwhile, Sensmet focuses on quantification for AES. The DL quantification method uses a two-headed CNN with both the hardware parameters and the spectral data as input. Finally, TU/e and Philips are focusing on MRS and currently exploring the whole spectroscopy workflow. They have found literature for all MRS workflow steps and many architectures include CNNs. While all partners have a common ground on the spectroscopy workflow, the investigated DL applications are different.

4.1.2. Network Architectures

Looking at the found and used DL architectures, all spectral domains use CNNs. These types of architectures are mainly used to process spectral data to a lower dimension, which occurs during quantification. Additionally, (V)AEs are also present in both the MRS and GRS pipeline. VAEs are used for artificial data generation and AEs for removing artifacts and denoising purposes. Other found architectures like MLP, RNN, U-Net and GAN are used in some workflow steps but are not implemented by more than one partner. More research and testing need to be done in order to find out if these architectures could be useful for spectral analysis in a cross-domain fashion. Also, new developments in DL architectures, e.g., transformers [1] should be followed closely as well as ways to explain and understand the working principles of selected DL models.

4.1.3. Outlook

By understanding the DL approaches used in each spectral domain, the partners of Spectralligence are actively collaborating to learn from each other's work and goals. Further research can be conducted to define the next steps in developing DL methods for spectral analysis in each domain and to identify opportunities for cross-domain DL models that can leverage similarities in applications and architectures. Identification of commonalities and emerging architectures in adjacent AI/ML application domains will be continued by the project partners.

4.1.4. References

[1] A. Vaswani *et al.*, “Attention is All you Need,” in *Advances in Neural Information Processing Systems*, 2017, vol. 30. Accessed: Dec. 20, 2022. [Online]. Available: https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html