Aus der Klinik und Poliklinik für Kinder- und Jugendpsychiatrie,
Psychosomatik und Psychotherapie
Klinikum der Ludwig-Maximilians-Universität München
Direktor: Prof. Dr. Gerd Schulte-Körne

Geometric Deep Learning for Alzheimer's Disease Analysis

Dissertation
zum Erwerb des Doktorgrades der Naturwissenschaften
an der Medizinischen Fakultät der
Ludwig-Maximilians-Universität München

vorgelegt von
Ignacio Sarasúa Cañedo-Argüelles
aus
Barcelona, Spanien

Jahr
2023

Mit Genehmigung der Medizinischen Fakultät
der Universität München

Betreuer(in):	Prof. Dr. rer. nat. Christian Wachinger
Zweitgutachter(in):	Prof. Dr. rer. nat. Michael Ingrisch
Dekan:	Prof. Dr. med. Thomas Gudermann
Tag der mündlichen Prüfung:	27. Oktober 2023

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

Dean's Office
Faculty of Medicine

MMRS MUNICH MEDICAL RESEARCH SCHOOL

Affidavit

Sarasua Cañedo-Argüelles, Ignacio

Surname, first name

Street

Zip code, town

Country

I hereby declare, that the submitted thesis entitled

Geometric Deep Learning for Alzheimer's Disease Analysis

is my own work. I have only used the sources indicated and have not made unauthorised use of services of a third party. Where the work of others has been quoted or reproduced, the source is always given.

I further declare that the dissertation presented here has not been submitted in the same or similar form to any other institution for the purpose of obtaining an academic degree.

Munich, 31.10.2023

Place, date

IGNACIO SARASUA CAÑEDO-ARGÜELLES

Signature doctoral candidate

"... Los hay que luchan una vida, esos son los IMPRESCINDIBLES".

Esto es para ti, pater.

Contents

List of publications

- **CASHformer: Cognition Aware SHape Transformer for Longitudinal Analysis Sarasua, I.**, Pölsterl, S., Wachinger, C. Medical Image Computing and Computer Assisted Interventions 2022.
- **Is a PET all you need? A multi-modal study for Alzheimer's disease using 3D CNNs.** Narazani, M., **Sarasua, I.**, Pölsterl, S., Lizarraga, A., Yakushev, I., Wachinger, C. Medical Image Computing and Computer Assisted Interventions 2022.
- **TransforMesh: A Transformer Network for Longitudinal modeling of Anatomical Meshes. Sarasua I**, Pölsterl S, Wachinger C. International Workshop on Machine Learning in Medical Imaging 2021
- **Geometric Deep Learning on Anatomical Meshes for the Prediction of Alzheimer's Disease Sarasua I**, Lee J, Wachinger C. International Symposium on Biomedical Imaging 2021
- **Recalibration of Neural Networks for Point Cloud Analysis**. **Sarasua I**, Pölsterl S, Wachinger C. International Conference on 3D Vision 2020
- **Discriminative and generative models for anatomical shape analysis on point clouds with deep neural networks**. **Sarasua I**, Gutiérrez-Becker B, Wachinger C. Medical Image Analysis. 2020
- **Recalibrating 3d convnets with project & excite** Rickmann AM, Roy AG, **Sarasua I**, Wachinger C. IEEE transactions on medical imaging. 2020.
- **Prediction of Fluid Intelligence from T1-Weighted Magnetic Resonance Images.** Pölsterl, S., Gutiérrez-Becker, B., **Sarasua, I.**, Guha Roy, A., Wachinger, C. (2019, October). In Challenge in Adolescent Brain Cognitive Development Neurocognitive Prediction (pp. 35-46).
- **A wide and deep neural network for survival analysis from anatomical shape and tabular clinical data**. Pölsterl S, **Sarasua I**, Gutiérrez-Becker B, Wachinger C. European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases 2019.

List of Abbreviations

- AD: Alzheimer's Disease
- APP: Amyloid precursor protein
- A: amyloid- peptide
- AOs: A oligomers
- CNN: convolutional neural networks
- CSF:cerebrospinal fluid
- DL: Deep Learning
- GeomDL: Geometric Deep Learning
- GPU: Graphics Processing Unit
- HC: Healthy Controls
- MCI: Mild Cognitive Impairment
- ML: Machine Learning
- MLP: Multi-layer perceptron
- MRI: Magnetic Resonance Image
- NLP: Natural Language Processing
- NFT: neurofibrillary tangles
- NNs: Neural Networks
- PET: Positron Emission Tomography
- RNNs: Recurrent Neural Networks
- ROI: Region of interest
- SotA: State-of-the-Art

Contribution to publications

3.1 Contribution to Publication I: Discriminative and generative models for anatomical shape analysis on point clouds with deep neural networks

Ignacio Sarasua was in charge of gathering the experimental results and writing the manuscript. Besides organizing the different sections of the article and contributing their writing, Ignacio also took care of the design of the architecture figures and tables reporting the results. Ignacio was also in charge of the reviewing phase, where he handle the different requests of the reviewers and modified the manuscript accordingly.

3.2 Contribution to Publication II: Hippocampal representations for deep learning on Alzheimer's disease

Ignacio Sarasua took care of performing the experiments and analyzing the results. In particular, Ignacio's work was to implement the different models, fine-tune their hyperparameters, and evaluate them on the different cross-validation sets. In addition, Ignacio also took care of generating the attribution maps for each shape representation. Ignacio also contributed to writing the manuscript and designing all the figures in it.

3.3 Contribution to Appendix A.1: Geometric Deep Learning on Anatomical Meshesfor the Prediction of Alzheimer's Disease (Conference Publication)

Ignacio was in charge of designing and analyizing the experiments. He supervised the whole project that lead to this publication. He also took care of writing the manuscript and adding the comments obtained during the review process.

3.4 Contribution to Appendix A.2: Recalibration of Neural Networks for Point Cloud Analysis (Conference Publication)

Ignacio was responsible for designing, implementing and analyzing the results of this work. In particular, he implemented each re-calibration block and analyzed the contribution to the overall performance for each task. He also wrote the manuscript as well as incorporated all the comments from the reviewers during the rebuttal phase.

3.5 Contribution to Appendix A.3: CASHformer: Cognition Aware SHape Transformer for Longitudinal Analysis (Conference Publication)

Ignacio designed, implemented, and analyzed all the experiments in this article. The main tasks included: finding the most suitable architecture and the optimal hyperparameters and performing a set of ablation studies to evaluate the contribution of each proposed block to the overall performance. He was also in charge of reviewing the State-of-the-Art as well as deciding which baselines were suitable to compare to. Ignacio also took care of the majority of the writing process and the creation of the figures.

3.6 Contribution to Appendix A.4: Is a PET all you need? A multi-modal study for Alzheimer's disease using 3D CNNs (Conference Publication)

Ignacio conceptualized this work and togther with Marla Narazani analyzed the outcome from it. He also worked on writing the manuscript. He also supervised the whole project that lead to this publication. Particularly, he attended several meetings with medical researchers and capture the main concerns in the article.

Introduction

In 1906, Dr. Alois Alzheimer noticed abnormalities in a patient that had passed away of a rare mental disease (1). The symptoms included unpredictable behavior, language problems, and memory loss. During the autopsy, unusual clumps and tangled bundles of fibers were found in the brain - nowadays known as amyloid plaques and neurofibrillary tangles (NFTs), respectively. Years later, this neurodegenerative disorder is known as Alzheimer's Disease (AD) and it has been labeled as the most common type of dementia (between 50-70% of the cases) [1]. The World Alzheimer Report (2), released by Alzheimer's Disease International, stated that around 50 million suffer from Dementia worldwide bringing the overall cost to 818 billion dollars. Patients suffering from AD experience a progressive and irreversible cognitive decline that in the long run disables the patient's capabilities of carrying out day-to-day activities (3).

4.1 Alzheimer's Disease pathology and stages

NFTs are mostly formed by abnormally phosphorylated and aggregated tau protein, which destabilizes the microtubules and affects axonal transport (4). More recent studies have also shown a link between cognitive deficits in early AD and tau pathology (5). Another very key factor of AD development is the sequential cleavage of the amyloid precursor protein (APP) by the $\gamma-$ and $\beta-$scretase enzymes, which causes amyloid-β peptide (Aβ) accumulation (6). In addition to plaque deposition, Aβ oligomers (AβOs) seem to play an important role in the cognitive decline produced during AD (7). These toxins accumulate in the cerebrospinal fluid (CSF) of the brain, affecting the synapse structure and composition and causing memory loss (8).
Subjects suffering from AD start experiencing pathophysiological changes in their brain long before the first symptoms appear. During the preclinical stage, patients might not notice any cognitive decline, however, the progressive amyloid deposition will start affecting their short-time memory. This stage of the disease is known as prodromal AD stage or Mild Cognitive Impairment (MCI) due to AD. In future stages of the disease, other brain areas get affected as well as the cognitive functions associated to them. This results in a lack of autonomy for the patient due to severe memory loss and metabolic derangements. While the spectrum of AD development has been deeply studied and depicted, the boundaries between the different stages are hard to draw (e.g. asymptomatic AD patients are hard to differentiate from subjects that are healthily aging).

[1] https://www.alzheimers.org.uk/

4.1.1 The role of the hippocampus in Alzheimer's Disease

The hippocampus is a subcortical structure that plays a key role in forming new memories. Its different sub-regions help with the generation of episodic memory (9). The CA3 area of the hippocampus is responsible for producing sharp-wave ripples (SWR). SWR propagates new memory traces into the neocortex and consolidates the memory (10). In the process of developing AD, the hippocampus is one of the first parts of the brain to be affected by the disease. In the early stages of the disease, while other structures might not experience many changes, the hippocampus shows rapid atrophy. NFTs first accumulate in the CA1 area of the hippocampus. Later in the disease, they extend and affect the subiculum, CA2, CA3, and DG (11). In addition, Aβ deposition has been observed to reduce the hippocampus inputs (12). Therefore, studying the deformation of the hippocampus over time can be a key biomarker to detect the early stages of AD.

4.2 Deep Learning on Alzheimer's Disease

As seen in the previous section, AD is a very complex neurodegenerative process, involving several areas of the brain. Researchers can use several biomarkers to detect it. However, given the large volumes of data, Machine Learning (ML) and in particular Deep Learning (DL) could be used to aid that decision. In addition, current tools could help identify patterns in the evolution of patients suffering from AD.

4.2.1 Detecting AD based on Hippocampus shapes

While these changes happen in many areas of the brain, analyzing them on a structural level can help as painting a better picture of this process. As discussed in section 4.1.1, medical research has been able to observe acute changes in the hippocampus in patients developing AD. For that reason, an interesting research question would be to analyze whether hippocampal atrophy can also be detected by DL models, and can be used as a classification method to differentiate AD patients from healthy controls. This could confirm the importance of this structure for automatic diagnosis.
One thing to take into account when looking at single parts of the brain is that these can be encoded as a 3D shape(i.e. the 3d form of an object defined by boundary lines, angles, and surfaces). In section 5.1, we explain different ways of representing a shape and introduce DL methods that can work on this type of input. In section 5.3.2, we explain how to obtain these representations from MRI scans.

4.2.2 Early detection using longitudinal data

Once the ability of Neural Networks (NNs) to capture changes in the hippocampus is confirmed (as can be seen in chapters 8 and 9), we can apply these findings to early detection off the disease. As mentioned in section 4.1, AD is a continuum of changes that happen over time. Therefore, observing the patients evolution can provide a better view on their prognosis. A great way to capture the cognitive decline is by monitoring subjects during a period of time as part of a study. In these longitudinal studies, the subjects enroll

at a particular date, known as baseline visit. Images, as well as other biomarkers and questionnaires, are taken in each visit as part of the screening process. The time between visits vary depending on the study, but it is usually consistent among all the patients. One of the main challenges in this type of studies is that not all the subjects attend each visit, leaving some of the entries in the database empty. This makes analyzing the data and identifying patterns a cumbersome task. Predicting these missing entries would not only help with this issue, but it could also help understanding better how the disease progresses under different circumstances. Latest work in the field of Natural Language Processing (NLP), such as Transformers (13), and its recent application to other types of temporal data (e.g. frame prediction in video sequences), has shown very promising results on predicting future states of longitudinal data. However, while their popularity has grown in many fields beyond NLP, to the best of our knowledge, their application for medical data imputation has not been explored. In section 5.2 we provide an overview on how these type of networks work.

4.3 Goal and layout

In short, the goal of this thesis is to analyze the importance of the hippocampus in the development of AD, from a Deep Learning perspective, and with a particular focus on early detection based on longitudinal trajectories. First, we will study the abilities of Geometric Deep Learning methods to detect the disease based on the hippocampus (chapter 8), as well as best practices for representing this structure (chapter 9). Then our findings will be applied to the clinical problem of forecasting/imputing missing trajectories and how this can help identify patients that are converting from MCI to AD (appendix A.3). Before explaining the contributions of this thesis, in chapter 5 we will introduce certain concepts that are relevant for a better understanding of this work, as well as insights into the data used for the different experiments.

Background

5.1 Geometric Deep Learning

Since the early 2000s, deep learning models, such as convolutional neural networks (CNNs), have revolutionized many fields. Given their high number of learnable parameters, these models were able to outperform the previous state-of-the-art in vision tasks such as image classification, segmentation, or object detection. While the range of applications is wide, the majority of these models were limited to Euclidean data, i.e. input data that could be represented in a grid-like structure (e.g. images). Graphs, meshes, or point clouds are only a few examples of data representations that do not fall into this category, and therefore, cannot benefit from regular CNNs. In 2017, the concept of Geometric Deep Learning (GeomDL) was introduced for the first time (14). Its purpose was to serve as an umbrella term that established a mathematical framework for existing neural network architectures that were able to work on non-Euclidean data, as well as for developing future ones. This type of model is particularly interesting for medical shape analysis since it allows us to study fine-grained changes in anatomical structures (e.g. hippocampus). In particular, there are two types of shape representations that have been relevant for anatomical shape analysis: point clouds and meshes.

5.1.1 Point Clouds

Point clouds (Figure 5.1) are very lightweight representations, composed of unordered sets of points that represent a 3D surface. They are usually represented as a vector of 3-D Cartesian coordinates (x, y, z). Point clouds are commonly used in 3D Computer Vision tasks where laser scanners (e.g. LiDAR(15)) generate a point for each measurement of the laser scan. However, they can also be sampled from a continuous surface (e.g. a segmentation mask) or a mesh.

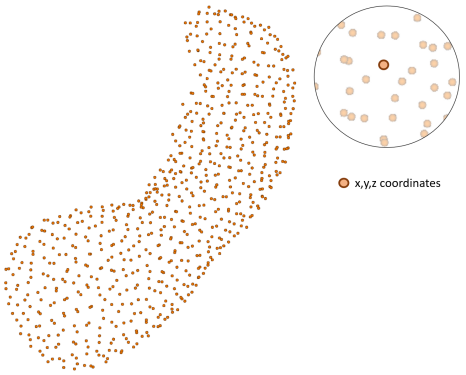

Figure 5.1: Hippocampus point cloud.

PointNet (16) was one of the first deep learning architectures proposed for 3D point cloud analysis. The main idea behind this method is to

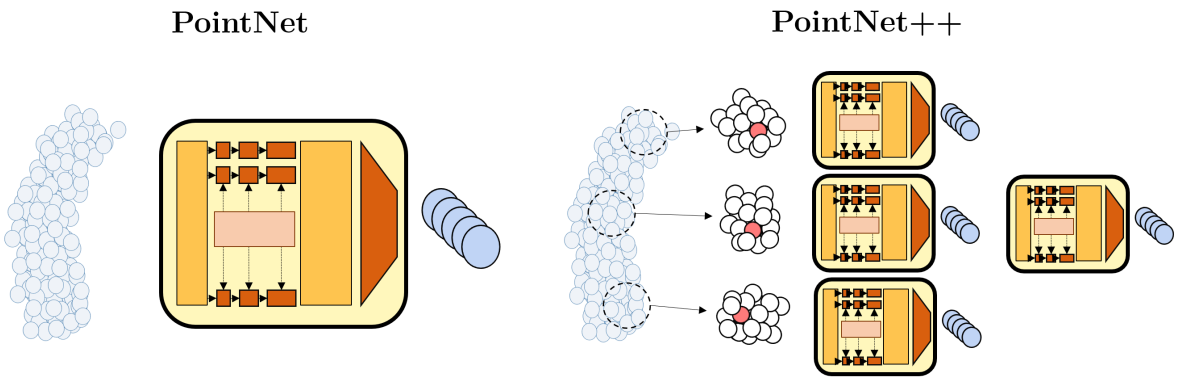

Figure 5.2: Left: Original PointNet architecture. Shared MLPs are used for obtaining features for each point. A max-pooling layer is used to obtain the global descriptor of the point cloud. Right: The input point cloud is split in point cloud subsets. For of these points, a PointNet is used to extract the cluster descriptors. This hierarchical behavior continues until the last layer, where the local descriptors are pooled into a global one.

extract features for every point in the point cloud by passing them through a Multi-layer Perceptron(MLP) that shares its weights across all points. In order to extract a global descriptor, all the individual feature vectors are collapsed using a max-pooling layer, which is a key factor to make the network order invariant, since it will capture the most relevant features without taking into account their position in the point cloud. In chapter 8, we explore the application of this network to AD analysis.
While PointNet was able to expand the applications of DL to point cloud analysis, its design limits the network to solely describe the overall shape, being unable to capture fine-grained local structural information. Many methods have been proposed to overcome this limitation(17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29). For example, the same authors introduced PointNet++(17), which builds on top of the original network, but captures local changes by clustering the structure in smaller subsets of point clouds and extracting features for each cluster (similar to regular CNNs). This concept of capturing features at increasingly larger scales along a multi-resolution hierarchy has also been explored in (18, 19). Other works introduce convolutional operations, similar to the ones in CNNs. However, in contrast to kernels for 2D grid structures, defining these operators is not as straightforward for 3D point clouds.The work in (20, 21, 22, 23, 24) define convolutional kernels on a continuous space, by considering the spatial distribution of the subset of points when weighting the contribution of each of them. The main difference between these works is the approach used to define the subset of points considered for the convolution. Finally, another group of methods approaches local information limitation by considering each point in the point cloud as a vertex of a graph(25, 26, 27, 28, 29). For example, DGCNN(25) defines a set of connections/edges between each point/vertex in the point cloud. However, the structure of the graph is not fixed along the network layers, since it is not defined by the proximity of the points in the Euclidean space, but by the similarity of their descriptors. In appendix A.2, we compare some of these networks and explore a set of recalibration blocks to improve their performance.

An advantage of these architectures is that they do not need inter-subject correspondences, which sometimes are harder to get when the topology of the different structures varies along the dataset. On the other hand, not having correspondences can increase the overall number of trainable parameters of the network, since pre-computations are not possible. In addition, interpreting the decision process of the model is more challenging, since the order of each point within the point cloud is different.

5.1.2 Meshes

Meshes are defined by a set of vertices (similar to Point Clouds) and edges that connect them (Figure 5.3). Having connectivity information establishes a more comprehensive representation of the underlying anatomical surface of the organ, compared to point clouds. Moreover, meshes can flexibly adapt to the complexity of the geometry, i.e. flatter regions of the shape can be represented with a small number of vertices and edges, while more complex areas can be captured by increasing the resolution. For brain structures, meshes can be directly obtained from segmentation softwares (e.g. FSL-FIRST) or by applying algorithms (e.g. marching cubes (30)) to segmentation masks.

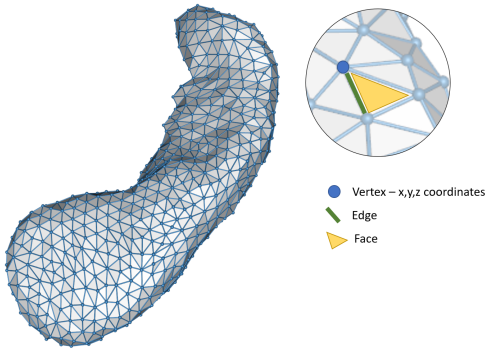

Figure 5.3: Hippocampus mesh.

Deep learning methods for meshes can be divided according to their need to register to a template mesh that establishes correspondences between all the meshes in the dataset. This means that all mesh samples need to have the same number of vertices and the same connectivity. The main advantage of this type of models is the possibility of pre-computing certain operations only once for the template, and then applying them to the samples in the dataset. This implies a significant reduction of the model complexity as well as the inference time. Some examples of template-based methods are Convolutional Mesh Autoncoder (CoMA) (31) and SpiralNet++ (32). The former uses of fast spectral convolutions, which operate in the frequency domain, as the main operators. In order to reduce operational complexity, fast spectral convolutions approximate the product between the convolution kernel and the Fourier transform of the vertices using recursive Chebyshev polynomials (33). SpiralNet++, on the other hand, defines the operators in the spatial domain. For each vertex, a spiral sequence along its neighbors is defined. This method takes advantage of the template, by pre-defining the spiral sequences on the template and extrapolating them to the rest of the samples. The convolutions are performed by concatenating the features within the spiral sequence and passing them through a MLP. Both CoMA and SpiralNet++ apply the same method for deciding which vertices to pool after each convolutional block: vertex pairs are iteratively contracted so they would minimize the quadric error between the original mesh and the down-sampled one.

While template-based models present multiple computational advantages, defining a com-

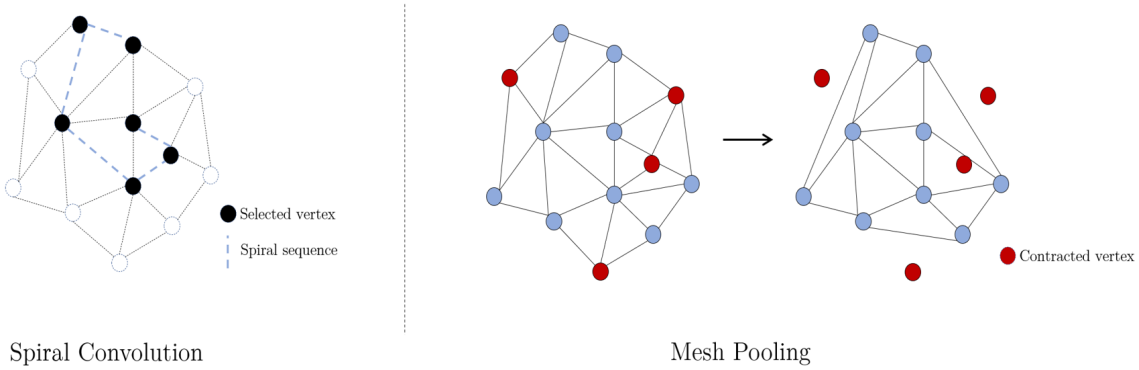

Figure 5.4: SpiralNet++ (32) main operations. Left: Spiral sequence used for computing the convolution. Right: Pooling operation based on smallest quadric error. Vertices that lead to the smallest difference to the surface are pooled.

mon topology across all the meshes in the dataset might not be possible for certain applications. Thus, methods that do not require of a reference shape, and fixed structure for all the samples, are also worth studying. Two examples of non-template-based methods are MeshCNN (34) and MeshNet (35). MeshCNN associates a set of descriptors for each edge in the mesh and directly operates on them. Therefore, the convolution operation is computed using invariant kernels. Edges associated with the lowest magnitude in the feature space are pooled after each convolutional block. On the other hand, MeshNet defines the faces of the mesh as the main feature unit and operates on them. Face-based features provide more regularity and order invariance, compared to other methods.
In our work added to appendix A.1), we compared the four methods with PointNet on the task of AD prediction. We observed that SpiralNet++ yields the best performance while being very efficient in terms of trainable parameters.

5.1.3 Euclidean representations

Despite the benefits of the previously discussed shape representations, it is worth studying Euclidean alternatives. As previously mentioned, the majority of CNNs methods have been developed for this type of data, presenting a clear advantage for them. In chapter 9, we compare point clouds and meshes (with their corresponding SotA methods) to three Euclidean forms of representing the hippocampus: masks, texturized masks, and region of interest (ROI). Masks are the binary output from any type of segmentation software, where voxels belonging to the structure (e.g. hippocampus) are set to 1 and the rest to 0. Texturized masks are obtained by substituting the 1-entries in the binary mask for the original values in the MRI. Finally, we can also define a 3D bounding box around the segmentation mask, and include all the values inside the bounding box as a ROI (see Figure 5.5).

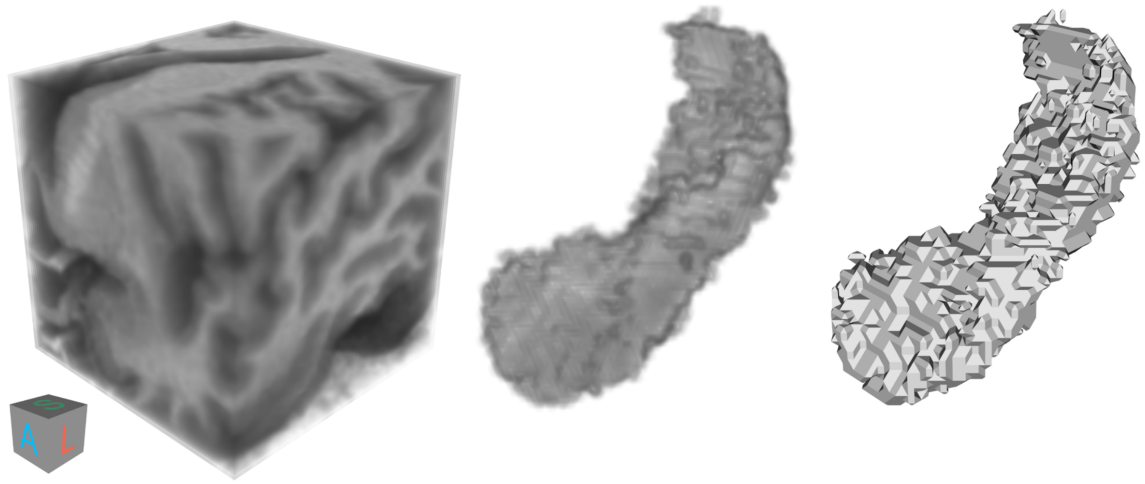

ROI Texture Mask

Figure 5.5: Euclidean representations

5.2 Transformers

Transformers(13) have become the new SotA for NLP applications. Until their appearance, language models were based on Recurrent NNs (RNNs), which keep the temporal information of the sequence in their hidden states. While this was able to capture relational patterns among the input elements, it enforces a sequential nature, which limits their parallelization, something that is particularly critical for long sequences. Transformer models overcome this issue by fully relying on attention mechanisms that can capture relationships between each pair of elements in the sequence. These operations can be computed in parallel, since it does not depend past states. Transformers have not only become a breakthrough for NLP, but for many other many fields such as Computer Vision (36) or Medical Image Segmentation (37).

5.2.1 Inputs

The input to the Transformers is a set of embeddings. Every element in the sequence (e.g. a word) is projected to a latent space. This operation can be done by a pre-defined dictionary, where each word is associated with a vector, or it can be computed by a neural network (e.g. (36)). Therefore the sequence would be translated into an input matrix, $\mathbf{X} \in \mathbb{R}^{N \times d}$, where N is the sequence length and d is the embedding dimension.

5.2.2 Encoder and Decoder stacks

The original Transformer architecture proposed (13) is formed by a stack of encoder blocks and another one of decoder blocks. Encoder and decoder blocks are practically identical in design. They are formed by a multi-head self-attention module, followed by a position-

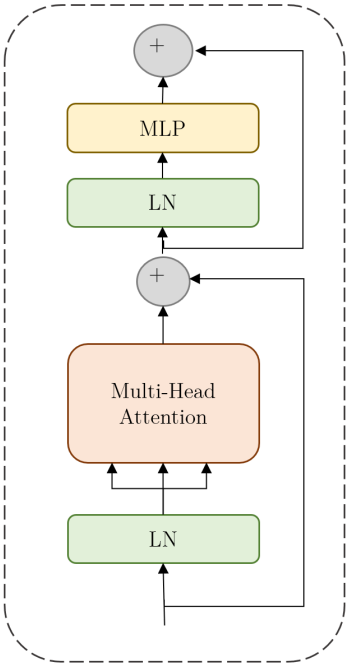

Figure 5.6: Transformer encoder block.

wise fully connected forward module. Residual connections are added to this modules as well as a layer normalization block (38) (See Figure 5.6). Decoder blocks also add an additional attention layer that incorporates the information from the encoder blocks to the decoding path.

5.2.3 Multi-Head Attention

The input matrix, $\mathbf{X} \in \mathbb{R}^{N \times d}$ (being N the length of the sequence) , is multiplied by three learnable matrices $\mathbf{W^Q} \in \mathbb{R}^{d \times d_Q}$, $\mathbf{W^K} \in \mathbb{R}^{d \times d_K}$ and $\mathbf{W^V} \in \mathbb{R}^{d \times d_V}$, resulting into three new matrices: Query $\mathbf{Q}$, Keys $\mathbf{K}$ and Values $\mathbf{V}$. The new projections $\mathbf{Q}$ and $\mathbf{K}$ of the matrix are used to computed the self-attention coefficients that will multiply the values, $\mathbf{V}$. The output of the Self-Attention module is computed as:

$$Attention(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = softmax(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}})\mathbf{V} \tag{5.1}$$

where d_k is the dimension of each row element in $\mathbf{K}$.
However, instead of only performing one single attention output, the operation is performed (in parallel) h times, where each of these attention heads will project the data on to a different latent space and perform the attention in it. Multi-head self attention will allow to build different spaces in which jointly attend to the information coming from the Queries and the Keys. The new attention output is computed by concatenating all the self-attention heads:

$$MultiHeadAttention(head_1, ..., head_h, \mathbf{W^O}) = Concat(head_1, ..., head_h)\mathbf{W^O} \tag{5.2}$$

with

$$head_i = Attention(\mathbf{XW_i^Q}, \mathbf{XW_i^K}, \mathbf{XW_i^V}) \tag{5.3}$$

where $\mathbf{W_i^Q} \in \mathbb{R}^{d \times d_Q}, \mathbf{W_i^K} \in \mathbb{R}^{d \times d_K}, \mathbf{W_i^V} \in \mathbb{R}^{d \times d_V}$ are the projection matrices for each head, and $\mathbf{W^O} \in \mathbb{R}^{hd_v \times d}$ is a weight matrix jointly trained with the attention heads.

5.2.4 Positional Encoding

Transformer models do not contain any structural information, since they do not use recurrent mechanisms. Therefore, in order to encode the order of the elements in the sequence, a set of positional embeddings ($\mathbf{P} \in \mathbb{R}^{N \times d}$) can be defined, such as that they are unique for every position. These embeddings can be fixed following a particular pattern (e.g. cosine functions of different frequencies (13)) or learnt together with the rest of the model (36). Transformers have also been applied to medical data. From predicting entries from electric health records (39) to analyzing medical images (40). In apendix A.3, we explore their application to predicting changes in the hippocampus.

5.3 Data

5.3.1 Datasets

In our experiments, we use data from The Alzheimer's Disease Neuroimaging Initiative (ADNI; (41)) and The Australian Imaging, Biomarker & Lifestyle Flagship Study of Ageing (AIBL; (42)).

ADNI

ADNI is a longitudinal study started in 2004 and lead by Dr.Michael W. Weiner. Its main goal is to detect AD at the earliest possible stage in order to support new advances on its prevention, intervention and treatment. Since its creation, ADNI has been able to enroll more than 1800 subjects between the ages 55 to 90 around 57 sites in the United States and Canada, divided in four groups: elderly controls, early MCI, late MCI and AD. After obtaining informed consent, participants undergo a set of initial tests which will be repeated in 6 months to 1 year intervals. These test might include clinical and neuropsychological evaluation, genetic testing, lumbar puncture and brain image acquisition (MRI and PET).

AIBL

AIBL was launched in 2006 in Australia in order to develop a better understanding of AD. With more than 1000 participants over the age of 60 across two centers, AIBL is the largest study of its kind in the country. Participants are grouped in three groups: AD, MCI and healthy volunteers. During the study, subjects undergo a set of tests including clinical and cognitive evaluation, biomarkers extraction, lifestyle questionnaires and neuroimaging (MRI and PET).

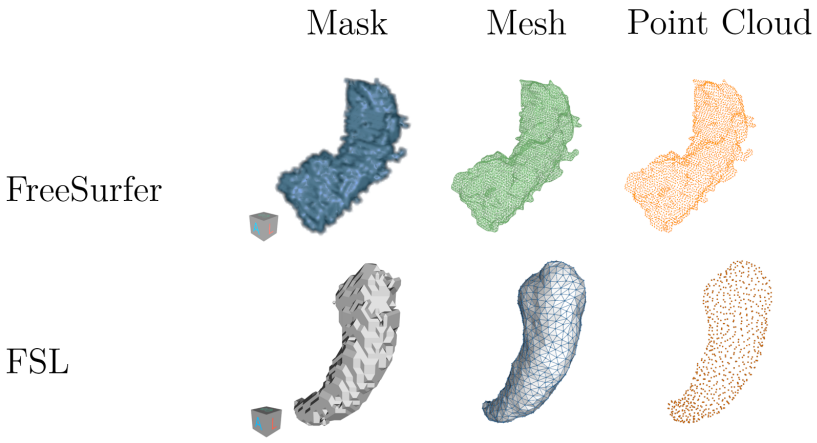

Figure 5.7: Output differences between FreeSurfer and FSL

5.3.2 Processing

The hippocampus shape are extracted from T1-weighted brain MRI scans. These are first conformed into a resolution $256 \times 256 \times 256$ and an isotropic voxel size of 1mm^3. After the conforming step, the N4ITK (43) algorithm is used for bias field correction. The resulting image is registered , using the SyN (44) affine registration algorithm implemented in ANTs, to the MNI space using the ICBM 2009c non-linear symmetric template (45). This pipeline guarantees that all the scans will be in the same space and have the same resolution, which will facilitate the following segmentation steps.

Segmentation and shape extraction

For segmenting the hippocampus we made use of two different softwares: FreeSurfer (46) and FSL-FIRST (47). FreeSurfer is an atlas-based method, i.e. a selected reference volume, which has been previously labeled, is deformed to better align with each patient. In contrast, FSL-FIRST proposes an Bayesian framework, utilizing the principles of Appearance Models and Active Shape , to build probabilistic relationships between shape and intensity.
The outputs of each segmentation software are also different (see Figure 5.7). FreeSurfer generates a voxel-wise segmentation map. By using marching cubes algorithm (48), we are able to extract 3D surfaces. After applying Laplacian smoothing to adjust the vertex coordinates, we can sample points from the reconstructed surface in order to generate a point cloud. The main limitation of this method is that the output meshes are not registered among each other, and therefore cannot be used by the template-based methods described in section 5.1.2. FSL-FIRST, on the other hand, can directly output 3D triangular meshes with vertex-wise correspondences along all the subjects and a template. Therefore, not only we can sample point clouds from the output surfaces to train the point cloud networks, but the generated meshes can be used for training template-based mesh neural networks described in section 5.1.2.

Abstract

Alzheimer's Disease (AD) represents between 50-70% of the cases of dementia, which translates in around 25-35 million people affected by this disease. During its development, patients suffering from AD experience an irreversible cognitive decline, which limits their autonomy on their daily lives. While many of the causes of AD are still unknown, researchers have noticed a abnormal amyloid deposition and neurofibrillary tangles that will start affecting the short-term memory of the patient, together with other cognitive functions. In fact, these pathophysiological changes start taking place even before the patient experiences the first symptoms. One of the structures that is first affected by the disease is the hippocampus. During the development of AD, this part of the brain experiences an irregular deformation that affects its capabilities of forming new memories. Therefore, many clinical work has set a focus on studying this structure and its evolution along the disease. Identifying the changes it suffers can help us understand better the causes of the patient's cognitive decline.
Given the complexity that characterizes AD, identifying patterns during its development is still a cumbersome task for physicians.Thus, aiding the diagnosis and prognosis of the disease using Deep Learning methods can be highly beneficial, as seen for other medical applications (49). In particular, if the focus is set on single structures (e.g. the hippocampus) Geometric Deep Learning offers a set of models that are best suited for 3D shape representations. We believe these methods can help doctors identify abnormalities in the structure that can lead to AD in the future.
In this work, we first study the capabilities of current Geometric Deep Learning methods in diagnosing patients suffering from AD, by only looking at the hippocampus. We start by studying one of the simplest 3d representations, point clouds. We continue by comparing this representation to other non-Euclidean representations, such as meshes, and also Euclidean ones (e.g. 3d masks). We observe that meshes are one of the optimal ways of representing 3d structures for capturing fine-grained changes, but they carry additional pre-processing steps that Euclidean representations do not require. Finally, once we have confirmed that Geometric Deep Learning, particularly mesh neural networks, can properly capture the effects of AD on the hippocampus, we expand their application to longitudinal analysis of the structure. We propose a new temporal model based on Spiral Resnet and Transformers that sets a new state-of-the-art for the task of predicting longitudinal trajectories of the hippocampus. We also evaluated the effect that imputing missing longitudinal data has on detecting subjects that are developping to AD. Our experiments show an increase of a 3% in distinguishing between converting and stable trajectories.

Zusammenfassung

Die Alzheimer-Krankheit macht zwischen 50 und 70% der Fälle von Demenz aus, was in etwa 25-35 Millionen Menschen entspricht. Im Verlauf der Krankheit kommt es bei Patienten mit Alzheimer zu einem irreversiblen Verlust kognitiver Fähigkeiten, wodurch sie den Alltag nicht mehr eigenständig bewältigen können. Während die genauen Ursachen der Alzheimer-Krankheit noch unbekannt sind, haben Forscher festgestellt dass eine abnorme Aggregation von Amyloid-Peptiden und Neurofibrillen das Kurzzeitgedächtnis der Patienten sowie andere kognitive Funktionen beeinträchtigen. Diese pathophysiologischen Veränderungen beginnen bereits bevor der Patient die ersten Symptome zeigt. Ein Teil des Gehirns, der zuerst von der Krankheit betroffen ist, ist der Hippocampus. Während der Entwicklung der Alzheimer-Krankheit wird dieser Teil des Gehirns beschädigt, was die Gedächtnisbildung beeinträchtigt. Daher haben sich viele klinische Arbeiten zu der Alzheimer-Krankheit auf die Untersuchung des Hippocampus konzentriert. Angesichts der Komplexität der Alzheimer-Krankheit kann der Einsatz von Deep-Learning-Methoden bei der Diagnose und Prognose der Krankheit von großem Nutzen sein. Insbesondere wenn der Fokus auf einzelne Gehirnstrukturen– wie dem Hippocampus– gelegt wird, bietet das geometrische Deep Learning eine Reihe von Modellen, die sich speziell für dreidimensionale Geometrien eignen. Wir glauben, dass diese Methoden Ärzten helfen können, Anomalien im Hippocampus zu identifizieren, die in Zukunft zu Alzheimer führen können. In dieser Dissertation untersuchen wir zunächst die Fähigkeiten aktueller geometrischer Deep Learning-Methoden bei der Diagnose von Patienten mit Alzheimer, indem wir nur den Hippocampus betrachten. Wir beginnen mit der Untersuchung einer der einfachsten dreidimensionalen Darstellungen, den Punktwolken. Anschließend vergleichen wir diese Darstellung mit anderen nicht-Euklidischen Darstellungen, wie z. B. Gitternetzen, und auch Euklidischen Darstellungen (z. B. 3D-Masken). Wir stellen fest, dass Gitternetze eine der besten Möglichkeiten zur Darstellung von 3D Strukturen sind um feine Änderungen zu erfassen. Allerdings erfordern sie zusätzliche Vorverarbeitungsschritte, die bei Euklidischen Darstellungen nicht erforderlich sind. Nachdem wir bestätigt haben, dass geometrisches Deep Learning, insbesondere neuronale Netze mit Gitternetzen, die Auswirkungen von Alzheimer auf den Hippocampus richtig erfassen können, erweitern wir ihre Anwendung auf die longitudinale Analyse des Hippocampus. Wir schlagen ein neues temporales Modell auf der Grundlage von Spiral Resnet und Transformers vor, das einen neuen Stand der Technik für die Vorhersage von longitudinalen Trajektorien des Hippocampus darstellt. Wir haben auch ausgewertet, welchen Nutzen das Imputieren fehlender longitudinaler Daten auf die Erkennung von Personen hat, die eine Alzheimer-Krankheit erleiden. Unsere Experimente zeigen eine Steigerung von 3% bei der Differen-

zierung zwischen konvertierenden und stabilen Trajektorien.

Publication I: Discriminative and generative models for anatomical shape analysis on point clouds with deep neural networks

Gutierrez-Becker, Benjamin, Ignacio Sarasua, and Christian Wachinger. "Discriminative and generative models for anatomical shape analysis on point clouds with deep neural networks." Medical Image Analysis 67 (2021): 101852.

Publication II: Hippocampal representations for deep learning on Alzheimer's disease

Sarasua, Ignacio, Sebastian Pölsterl, and Christian Wachinger. "Hippocampal representations for deep learning on Alzheimer's disease." Scientific reports 12.1 (2022): 8619.

Conclusion

This work has focused on analzying the role of the hippocampus during the development of AD from a DL perspective. First, we evaluated the ability of Geometric DL networks to detect differences between HC and AD patients. Later on, we applied our findings in studying longitudinal trajectories combining GeomDL networks and Transformers.

10.1 Geometric DL for AD diagnosis based on hippocampus shapes

We have confirmed the medical findings on the importance of this structure during AD development. In particular, we have observed that the most simple representations, point clouds, can already provide good accuracy in detecting AD. However, if we compare its performance to other shape representations, we can see that meshes are the best non-Euclidean form of representing the hippocampus. Nevertheless, while we have confirmed that these models can be applied to the target task, and even outperform regular CNNs that work on 3D images, they might not be an optimal choice for detecting AD. Working with non-Euclidean representations currently still presents several challenges that their Euclidean counterparts do not (e.g. pre-processing). In addition, contrary to shape representations, working directly on 3D volumes allow us to define a continuous region of interest, that can include several structures or even the whole brain. While fitting a whole 3D volume in a GPU was challenging in the past, recent technological advances in the field have made larger memories more affordable, which can help fit larger amounts of data. In addition, given that AD affects so many parts of the brain at the same time, this feature can be quite decisive, especially in early detection tasks. Moreover, one of our latest works (appendix A.4), has shown that PET is a more fitting modality for detecting AD. We, therefore, believe that working on larger regions of the brain, and on PET data (if possible), yields the best performance when the task is to solely detect AD based on medical imaging data.

10.2 Longitudinal analysis using Transformers on Hippocampus meshes

When studying the overall evolution of subjects developing AD, using large regions of brain scans might not be as suitable anymore. While registration algorithms can provide a good alignment between scans, having perfect correspondences between areas of the brain is still not possible. This issue is particularly present when working on PET scans, due to their low resolution. In those instances, shape analysis, applied to segmentations coming from MRI, has always been a better way of approaching the problem. In particular, software that provides perfectly registered meshes, like FSL-FIRST, is useful since they enable vertex-wise differences between two different meshes. When observing longitudinal data, this feature is especially convenient, since it allows us to identify more patterns in the disease progression. In addition, neural networks working on this type of mesh data can lower the complexity by pre-computing many of their operations. Given the light weight of these networks, they can be easily be integrated in heavier temporal models, such as Transformers. In appendix A.3 we have observed that combining spiral convolutions with pre-trained transformers, can help to solve the problem of data imputation as well as predict future stages of the disease, even in datasets with limited amount of samples. Imputing and forecasting shape trajectories can be particularly useful to detect patients that are converting from MCI to AD.

10.3 Future directions

AD is a very complex phenomenon, which involves several functional and structural changes in the brain. In addition, many of their symptoms can be confused with other diseases (such as depression). Currently, researchers approach its diagnosis from a multimodal perspective, where they combine the information from different biomarkers. We believe DL research should follow the same direction and study the following topics:

10.3.1 Detection of different types of Dementia and AD

While AD is the most prevalent type of dementia, there are other types of dementia that are also affecting a large portion of society. The most common ones after AD are Vascular, Lewy Body Disease, and Frontotemporal. Even though their symptoms might be similar, the physiological effects in the brain differ. Early detection and differentiation of the different types will lead to a better aid for the treatment.
It is also worth noting that, while AD is characterized as one type of dementia, different subtypes of this disease have been observed(50). Accounting for these differences when building our models will make them more robust and easier to generalize, which is strongly desirable before deploying them in clinical practice.

10.3.2 Integration with multimodal data and explainability

AD (as well as other types of dementia) is a very complex disease that involves various processes in the brain. As seen in section 4.1, there are several biomarkers that could be included in forming a decision. One of the advantages of DL is the capacity of including large inputs without having to preselect the features and letting the model decide which ones are the most relevant for the target task. However, this black-box behavior is not ideal in clinical practice. Physicians need to ensure that the decision has been made based on factors that correlate with previous medical findings. We believe a possible extension of this work would be to incorporate several sources of data (e.g. shape, imagining, genetics) and together with explainability algorithms, to get a more informed decision.

10.3.3 Longitudinal analysis of multiple shapes

The hippocampus is not the only structure affected by AD. Our findings should motivate future work on incorporating further shapes into the analysis to build a more holistic picture of the disease evolution. This is particularly beneficial for those types of AD where the hippocampus is not affected (i.e. hippocampus sparing (50)). In addition, combining several structures at the same time will help finding correlations between the changes observed in the different structures at each stage of the disease.

Bibliography

[1] Maurer K, Volk S, Gerbaldo H. Auguste D and Alzheimer's disease. The lancet. 1997;349(9064):1546-9.

[2] Patterson C. World alzheimer report 2018. 2018.

[3] Öhman A, Josephsson S, Nygaard L. Awareness through interaction in everyday occupations: experiences of people with Alzheimer's disease. Scandinavian journal of occupational therapy. 2008;15(1):43-51.

[4] Morris M, Maeda S, Vossel K, Mucke L. The many faces of tau. Neuron. 2011;70(3):410-26.

[5] Fu H, Rodriguez GA, Herman M, Emrani S, Nahmani E, Barrett G, et al. Tau pathology induces excitatory neuron loss, grid cell dysfunction, and spatial memory deficits reminiscent of early Alzheimer's disease. Neuron. 2017;93(3):533-41.

[6] Selkoe DJ, Hardy J. The amyloid hypothesis of Alzheimer's disease at 25 years. EMBO molecular medicine. 2016;8(6):595-608.

[7] Yang T, Li S, Xu H, Walsh DM, Selkoe DJ. Large soluble oligomers of amyloid β-protein from Alzheimer brain are far less neuroactive than the smaller oligomers to which they dissociate. Journal of Neuroscience. 2017;37(1):152-63.

[8] Bilousova T, Miller CA, Poon WW, Vinters HV, Corrada M, Kawas C, et al. Synaptic amyloid-β oligomers precede p-Tau and differentiate high pathology control cases. The American journal of pathology. 2016;186(1):185-98.

[9] Collin SH, Milivojevic B, Doeller CF. Memory hierarchies map onto the hippocampal long axis in humans. Nature neuroscience. 2015;18(11):1562-4.

[10] Abadchi JK, Nazari-Ahangarkolaee M, Gattas S, Bermudez-Contreras E, Luczak A, McNaughton BL, et al. Spatiotemporal patterns of neocortical activity around hippocampal sharp-wave ripples. Elife. 2020;9:e51972.

[11] De Flores R, La Joie R, Chételat G. Structural imaging of hippocampal subfields in healthy aging and Alzheimer's disease. Neuroscience. 2015;309:29-50.

[12] Lace G, Savva G, Forster G, De Silva R, Brayne C, Matthews F, et al. Hippocampal tau pathology is related to neuroanatomical connections: an ageing population-based study. Brain. 2009;132(5):1324-34.

[13] Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need. Advances in neural information processing systems. 2017;30.

[14] Bronstein MM, Bruna J, LeCun Y, Szlam A, Vandergheynst P. Geometric deep learning: going beyond euclidean data. IEEE Signal Processing Magazine. 2017;34(4):18-42.

[15] Dong P, Chen Q. LiDAR remote sensing and applications. CRC Press; 2017.

[16] Qi CR, Su H, Mo K, Guibas LJ. Pointnet: Deep learning on point sets for 3d classification and segmentation. Proc Computer Vision and Pattern Recognition (CVPR), IEEE. 2017;1(2):4.

[17] Qi CR, Yi L, Su H, Guibas LJ. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: Advances in neural information processing systems; 2017. p. 5099-108.

[18] Zhao H, Jiang L, Fu CW, Jia J. PointWeb: Enhancing Local Neighborhood Features for Point Cloud Processing. In: CVPR; 2019. .

[19] Yan X, Zheng C, Li Z, Wang S, Cui S. PointASNL: Robust Point Clouds Processing using Nonlocal Neural Networks with Adaptive Sampling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition; 2020. p. 5589-98.

[20] Hua BS, Tran MK, Yeung SK. Pointwise Convolutional Neural Networks. In: CVPR; 2018. .

[21] Wu W, Qi Z, Fuxin L. PointConv: Deep convolutional networks on 3D point clouds. In: CVPR; 2019. .

[22] Xu Y, Fan T, Xu M, Zeng L, Qiao Y. SpiderCNN: Deep learning on point sets with parameterized convolutional filters. In: ECCV; 2018. .

[23] Li Y, Bu R, Sun M, Wu W, Di X, Chen B. PointCNN: Convolution on x-transformed points. In: NeurIPS; 2018. .

[24] Liu Y, Fan B, Xiang S, Pan C. Relation-Shape Convolutional Neural Network for Point Cloud Analysis. In: CVPR; 2019. .

[25] Wang Y, Sun Y, Liu Z, Sarma SE, Bronstein MM, Solomon JM. Dynamic Graph CNN for Learning on Point Clouds. ACM TOG. 2019.

[26] Te G, Hu W, Zheng A, Guo Z. RGCNN: Regularized graph CNN for point cloud segmentation. In: ACM MM; 2018. .

[27] Zhang Y, Rabbat M. A Graph-CNN for 3D point cloud classification. In: ICASSP; 2018. .

[28] Liu J, Ni B, Li C, Yang J, Tian Q. Dynamic Points Agglomeration for Hierarchical Point Sets Learning. In: ICCV; 2019. .

[29] Chen C, Li G, Xu R, Chen T, Wang M, Lin L. ClusterNet: Deep Hierarchical Cluster Network with Rigorously Rotation-Invariant Representation for Point Cloud Analysis. In: CVPR; 2019. .

[30] Lorensen WE, Cline HE. Marching cubes: A high resolution 3D surface construction algorithm. ACM siggraph computer graphics. 1987;21(4):163-9.

[31] Ranjan A, Bolkart T, Sanyal S, Black MJ. Generating 3D faces using convolutional mesh autoencoders. In: European Conference on Computer Vision (ECCV); 2018. p. 704-20.

[32] Gong S, Chen L, Bronstein M, Zafeiriou S. SpiralNet++: A Fast and Highly Efficient Mesh Convolution Operator. In: IEEE International Conference on Computer Vision Workshops; 2019. .

[33] Defferrard M, Bresson X, Vandergheynst P. Convolutional neural networks on graphs with fast localized spectral filtering. Advances in neural information processing systems. 2016;29.

[34] Hanocka R, Hertz A, Fish N, Giryes R, Fleishman S, Cohen-Or D. MeshCNN: a network with an edge. ACM Transactions on Graphics (TOG). 2019;38(4):1-12.

[35] Feng Y, Feng Y, You H, Zhao X, Gao Y. MeshNet: mesh neural network for 3D shape representation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33; 2019. p. 8279-86.

[36] Dosovitskiy A, Beyer L, Kolesnikov A, Weissenborn D, Zhai X, Unterthiner T, et al. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In: International Conference on Learning Representations; 2020. .

[37] Hatamizadeh A, Tang Y, Nath V, Yang D, Myronenko A, Landman B, et al. Unetr: Transformers for 3d medical image segmentation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision; 2022. p. 574-84.

[38] Ba JL, Kiros JR, Hinton GE. Layer normalization. arXiv preprint arXiv:160706450. 2016.

[39] Li Y, Rao S, Solares JRA, Hassaine A, Ramakrishnan R, Canoy D, et al. BEHRT: transformer for electronic health records. Scientific reports. 2020;10(1):1-12.

[40] He K, Gan C, Li Z, Rekik I, Yin Z, Ji W, et al. Transformers in medical image analysis: A review. arXiv preprint arXiv:220212165. 2022.

[41] Jack CR, Bernstein MA, Fox NC, Thompson P, Alexander G, Harvey D, et al. The Alzheimer's disease neuroimaging initiative (ADNI): MRI methods. Journal of magnetic resonance imaging. 2008;27(4):685-91.

[42] Ellis K, Bush A, Darby D, et al. The Australian Imaging, Biomarkers and Lifestyle (AIBL) study of aging. International Psychogeriatrics. 2009;21(04):672-87.

[43] Tustison NJ, Avants BB, Cook PA, Zheng Y, Egan A, Yushkevich PA, et al. N4ITK: improved N3 bias correction. IEEE transactions on medical imaging. 2010;29(6):1310-20.

[44] Avants BB, Epstein CL, Grossman M, Gee JC. Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. Medical image analysis. 2008;12(1):26-41.

[45] Manera AL, Dadar M, Fonov V, Collins DL. CerebrA, registration and manual label correction of Mindboggle-101 atlas for MNI-ICBM152 template. Scientific Data. 2020;7(1):1-9.

[46] Fischl B. FreeSurfer. NeuroImage. 2012;62(2):774-81.

[47] Patenaude B, Smith SM, Kennedy DN, Jenkinson M. A Bayesian model of shape and appearance for subcortical brain segmentation. NeuroImage. 2011;56(3):907-22.

[48] Lorensen WE, Cline HE. Marching cubes: A high resolution 3D surface construction algorithm. ACM SIGGRAPH Computer Graphics. 1987;21(4):163-9.

[49] Cruz-Roa A, Gilmore H, Basavanhally A, Feldman M, Ganesan S, Shih NN, et al. Accurate and reproducible invasive breast cancer detection in whole-slide images: A Deep Learning approach for quantifying tumor extent. Scientific reports. 2017;7(1):1-14.

[50] Ferreira D, Verhagen C, Hernández-Cabrera JA, Cavallin L, Guo CJ, Ekman U, et al. Distinct subtypes of Alzheimer's disease based on patterns of brain atrophy: longitudinal trajectories and clinical applications. Scientific reports. 2017;7(1):1-13.

Appendix

A.1 Appendix 1: Geometric Deep Learning on Anatomical Meshes for the Prediction of Alzheimer's Disease

Sarasua, Ignacio, Jonwong Lee, and Christian Wachinger. "Geometric deep learning on anatomical meshes for the prediction of Alzheimer's disease." 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI). IEEE, 2021.

A.2 Appendix 2: Recalibration of Neural Networks for Point Cloud Analysis

Sarasua, Ignacio, Sebastian Pölsterl, and Christian Wachinger. "Recalibration of neural networks for point cloud analysis." 2020 International Conference on 3D Vision (3DV). IEEE, 2020.

A.3 Appendix 3: CASHformer: Cognition Aware SHape Transformer for Longitudinal Analysis

Sarasua, Ignacio, Sebastian Pölsterl, and Christian Wachinger. "CASHformer: Cognition Aware SHape Transformer for Longitudinal Analysis." International Conference on Medical Image Computing and Computer-Assisted Intervention. Cham: Springer Nature Switzerland, 2022.

A.4 Appendix 4: Is a PET all you need? A multi-modal study for Alzheimer's disease using 3D CNNs

Narazani, Marla, et al. "Is a PET all you need? a multi-modal study for Alzheimer's disease using 3D CNNs." International Conference on Medical Image Computing and Computer-Assisted Intervention. Cham: Springer Nature Switzerland, 2022.

Acknowledgement

There are many people who have made this work a reality, and I need to thank them for it.
First of all to my advisor, Christian Wachinger. His proximity and availability at all times are some of the main reasons this work has been possible. He has not only acted as a supervisor for this work, but also as mentor in my path to become a researcher. Also to my unofficial advisor, Sebastian Pölsterl, who over these last years he has always been there as great colleague, a mentor and a friend. He is probably the most crucial person for this dissertation, not only for his infinite feedback, but also because he is printing this. To the rest of my colleagues over these years: Benjamin, Abhijit, Anne, Fabi, Nuno and my students Lee and Marla. Without our discussions, constant brain-picking, and fancy Fridays, I would not have been able to make it.
To my parents, Papa y Mama. My siblings, Benja, la Nena, Xico, Armando and Celia. My nephews and nieces Xico, Lola, Juan, Cecilia, Alejandra, Jorge, Jaime, Javi and Alvaro. My aunts, uncles and cousins. Despite the distance, they have been my rock over the years and they made me feel like they were with me at all times. Their constant support has made it much easier. Os quiero mucho.
My Munich friends. Ghetto, Stustas, Despechaos. Thank you for creating our own bubble where things happen under our own rules and pace. Also thank you for being the family I chose. To all my roomates, but especially to Awino, Xavier, Meghan and Diana, you made me feel home at all times. Everyone I met through dancing, with particular mention to Riddim, Moove, and Kalego. I never thought that something that started as a simple hobby could have brought so much and so many people into my life. Also to all my friends in Barcelona, especially Ines, Paloma y Miki. For not letting me forget how much I was missed back home and always made me feel like they will wait for me. And my Telecos, in Gil, Axel and Adri, who despite being all over the globe, they made sure our friendship would never change.
And of course, to Lance, for being my everything through the last three years. You have encouraged me to do things I would have never dared to dream about. Thank you for being a constant and a "home" to me. I love you.