

# Atlas-Based Interpretable Age Prediction In Whole-Body MR Images

Sophie STARCK [HTTPS://ORCID.ORG/0000-0003-2495-6114](https://orcid.org/0000-0003-2495-6114)

sophie.starck@tum.de

Artificial Intelligence in Healthcare and Medicine, School of Computation, Information and Technology, Technical University of Munich, Munich, Germany

Yadunandan Vivekanand KINI [HTTPS://ORCID.ORG/0000-0003-1558-0602](https://orcid.org/0000-0003-1558-0602)

ge75yoh@mytum.de

Artificial Intelligence in Healthcare and Medicine, School of Computation, Information and Technology, Technical University of Munich, Munich, Germany

Jessica J. M. RITTER [HTTPS://ORCID.ORG/0000-0001-7106-7518](https://orcid.org/0000-0001-7106-7518)

jessica.ritter@tum.de

Institute of Diagnostic and Interventional Radiology, Technical University of Munich, School of Medicine, Munich, Germany

Rickmer BRAREN [HTTPS://ORCID.ORG/0000-0001-6039-6957](https://orcid.org/0000-0001-6039-6957)

rbraren@tum.de

Artificial Intelligence in Healthcare and Medicine, School of Computation, Information and Technology, Technical University of Munich, Munich, Germany

Institute of Diagnostic and Interventional Radiology, Technical University of Munich, School of Medicine, Munich, Germany

German Cancer Consortium (DKTK), Munich partner site, Heidelberg, Germany

Daniel RUECKERT [HTTPS://ORCID.ORG/0000-0002-5683-5889](https://orcid.org/0000-0002-5683-5889)

daniel.rueckert@tum.de

Artificial Intelligence in Healthcare and Medicine, School of Computation, Information and Technology, Technical University of Munich, Munich, Germany

BioMedIA, Department of Computing, Imperial College London, UK

Tamara T. MUELLER [HTTPS://ORCID.ORG/0000-0002-1818-1036](https://orcid.org/0000-0002-1818-1036)

tamara.mueller@tum.de

Artificial Intelligence in Healthcare and Medicine, School of Computation, Information and Technology, Technical University of Munich, Munich, Germany

## Abstract

Age prediction is an important part of medical assessments and research. It can aid in detecting diseases as well as abnormal ageing by highlighting potential discrepancies between chronological and biological age. To improve understanding of age-related changes in various body parts, we investigate the ageing of the human body on a large scale by using whole-body 3D images. We utilise the Grad-CAM method to determine the body areas most predictive of a person's age. In order to expand our analysis beyond individual subjects, we employ registration techniques to generate population-wide importance maps that show the most predictive areas in the body for a whole cohort of subjects. We show that the investigation of the full 3D volume of the whole body and the population-wide analysis can give important insights into which body parts play the most important roles in predicting a person's age. Our findings reveal three primary areas of interest: the spine, the autochthonous back muscles, and the cardiac region, which exhibits the highest importance. Finally, we investigate differences between subjects that show accelerated and decelerated ageing<sup>1</sup>.

**Keywords:** Age prediction, Medical atlases, UK Biobank

---

1. The source code for this work can be found at: [https://github.com/starcksophie/atlas\\_age\\_pred](https://github.com/starcksophie/atlas_age_pred)

## 1. Introduction

Deep learning (DL) methods have significantly advanced medical research by delivering insights into normal physiology and disease processes. It can provide imaging-derived biomarkers for non-invasive predictions and support physicians in their work (Wang et al., 2019; Piccialli et al., 2021). Given the high sensitivity of medical data and the potentially life-altering impact that can result from using DL models for medical diagnoses or interventions, it is important to understand how or why a model reaches its decision. By inspecting which parts of the input are most relevant to a model’s decision, one can examine whether the model actually uses (medically) meaningful information or whether confounders are part of the decision process.

The investigation of ageing, age-related diseases, and the identification of specific areas in the body affected by age have been prominent research areas in medicine. Age shows one of the strongest correlations with the development of diseases and well-being in general (Niccoli and Partridge, 2012; Seale et al., 2022). Therefore, acquiring more knowledge about the ageing process can give insights into risk factors or abnormal ageing and serve as an early detection mechanism for several diseases (Fayosse et al., 2020). The utilisation of an accurate age prediction method can aid in (a) establishing a better understanding of the mechanisms of ageing in the human body and (b) finding discrepancies between an individual’s chronological and biological age. Chronological age refers to the time elapsed since birth, whereas biological age aims to describe the physiological age, e.g. how the body has aged. There might be deviations between the two, which is often referred to as *accelerated* (biological age > chronological age) or *decelerated* (biological age < chronological age) ageing. This has been investigated extensively for brain age estimation (Sajedi and Pardakhti, 2019) since brain structures are known to change over time (Esiri, 2007; Huizinga et al., 2018) and be highly correlated with neurodegenerative diseases such as Alzheimer’s or Parkinson’s disease (Luders et al., 2016). Brain magnetic resonance images (MRIs) are promising modalities to infer the biological brain age of a subject, often with the help of deep learning techniques (Sajedi and Pardakhti, 2019). Age estimation has also been performed on dental data (Verma et al., 2019), skeleton bones in the body, such as chest radiography (Monum et al., 2017), knee skeletons (Maggio, 2017), or hand skeletons (Darmawan et al., 2015). Despite significant changes in several abdominal organs and tissues, such as the liver (Tajiri and Shimizu, 2013), bone densities (Wishart et al., 1995), and the pancreas (Meier et al., 2007), whole-body age prediction has so far not been explored in great detail. However, some works have shown significant advancements in this direction, focusing on the abdominal Le Goallec et al. (2022) and whole-body scans Langner et al. (2019).

In this work, we investigate age prediction on the whole body (excluding the brain) to identify which areas show the highest information value about a person’s age, utilising the capacity of the whole 3D volumes. Towards this goal, we train a convolutional neural network (CNN) on 3D MR images that cover the full body from neck to knee. Subsequently, we apply Grad-CAM (Selvaraju et al., 2017), a well-established post-hoc interpretability method for CNNs, to identify areas in the body that are most important to the algorithm’s decision-making. Since we are specifically interested in the population-wide areas of highest interest for the model, we subsequently register the Grad-CAM results onto a medical atlas to acquire population-wide importance maps. Figure 1 shows an overview of the pipeline of

our work. We identify three main regions of interest in the extracted importance maps: the spine, the autochthonous back muscles, and the heart with its adjacent great vessels like the aorta. Figure 2 shows atlas-based importance maps for the healthy female group. We can see that the region along the spine and the area surrounding the heart show the most prominent Grad-CAM activation.

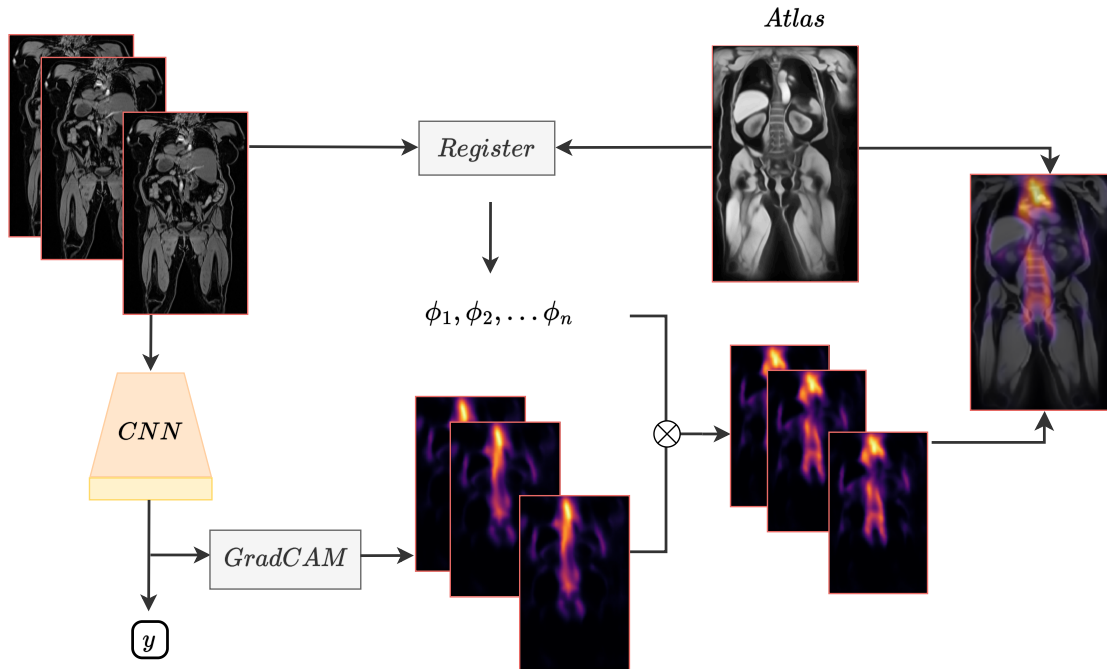


Figure 1: Overview of the pipeline used in this work: First, the CNN is trained to predict age. From the trained model, at inference, Grad-CAM visual explanations are extracted for each subject and then mapped to an atlas before being averaged into a population-wide importance map.

## 2. Background and Related Work

In this section, we summarise relevant background and related works that address interpretability in medical imaging, age prediction, and population-wide studies with medical atlases.

### 2.1 Grad-CAM

Interpretability methods, such as Grad-CAM can be applied to DL algorithms in order to get a better understanding of the decision-making process of neural networks (Carvalho et al., 2019). This is especially important in the medical domain, where critical patient diagnoses might depend on DL predictions, and both physicians and patients might want to understand how or why a model reaches a specific decision. One of the most commonly used interpretability methods is gradient-weighted class activation mapping (Grad-CAM)

(Selvaraju et al., 2017). Grad-CAM utilises the gradient information that flows into a convolutional layer of a CNN and applies global average pooling on these gradients to extract importance values for each input parameter (i.e. image voxel). Grad-CAM was originally designed for image classification, image captioning, and virtual question-answering tasks (Selvaraju et al., 2017); it has, since then, been applied for numerous tasks such as object detection or reinforcement learning (Joo and Kim, 2019; Dubost et al., 2020). However, it has been shown that it can also facilitate meaningful interpretations for regression tasks (Chen et al., 2020).

Grad-CAM has been used in several applications of DL to medical data (Xiao et al., 2021; Panwar et al., 2020; Daanouni et al., 2021) and also specifically in the context of age predictions (Langner et al., 2019; Le Goallec et al., 2022; Kerber et al., 2023; Bintsi et al., 2021; Raghu et al., 2021). However, one shortcoming of gradient-based interpretability methods such as Grad-CAM is that the results are subject-specific and do not allow for a population-wide investigation. In the medical context, subject-specific interpretation can be of interest in individual assessments, while a population-level map might hold more generalisable information. In this work, we are interested in population-wide importance maps, which we obtain by using registration methods.

## 2.2 Age prediction

Ageing is the main risk factor for disease development, and it is an important indicator of a person’s overall health Hou et al. (2019); Niccoli and Partridge (2012). MR images, in particular, hold great potential in the investigation of the physiological effects of ageing and subsequently identifying diseases. For instance, deep learning methods have been extensively applied to brain age estimation, achieving a highly accurate age prediction with an error of 2.14 years (Peng et al., 2021). Age prediction has also been investigated on different body regions such as the teeth (Verma et al., 2019), the chest (Monum et al., 2017), knees (Maggio, 2017), etc. Le Goallec et al. (2022), have been focusing on abdominal age prediction from liver and pancreas MR images and achieve a mean absolute error (MAE) of 2.94 years. The authors also utilise Grad-CAM to highlight the most relevant areas for the model’s prediction. However, here, a clear selection of specific regions is difficult as only subject-level areas have been investigated.

The probably most relevant related work to ours is Langner et al. (2019). The authors also perform interpretable age prediction on whole-body images and achieve a performance of 2.49 years on the UK Biobank (Sudlow et al., 2015). Here, constructed projections of the 3D image data into the 2D sagittal and coronal planes are used for training (see Appendix D, Figure 9). Throughout this paper, we refer to this method as 2.5D since the data is 2D in terms of dimensions, but encompasses information from the entire 3D volume. In Langner et al. (2019), the authors also use Grad-CAM to obtain interpretable maps, indicating which areas of the body are most important for the model to make its prediction. They aggregate the resulting saliency maps by co-registering the dataset onto a single representation. Using these projections comes with a major advantage of requiring fewer resources and training time. However, this method requires significantly more training samples to reach a comparable model performance. We compare our method to their approach and discuss these elements in more detail in Section 4. The probably strongest

shortcoming of their methods is that the resulting interpretability maps are 2D and highlight regions of the body where all projected slices are overlaid. This makes the actually most meaningful areas indistinguishable from data from other slices. In order to obtain more precise areas of interest in three dimensions, we here use the full capacity of the 3D volumes.

### 2.3 Population-wide Studies and Medical Atlases

Medical imaging is indispensable for medical research and assessment. However, medical images mostly come with high inter-subject variability that can stem from different morphologies or even just different positions in the scanner. Therefore, medical atlases are frequently used to allow for inter-subject or inter-population comparisons. They map several medical images into a common coordinate system, using registration techniques (Maintz and Viergever, 1998). The registered images are then averaged in order to acquire a template of a specific image modality. This is widely used for brain imaging, where atlases are used to generate an average representation of the human brain (Insel et al., 2013; Markram, 2012; Van Essen et al., 2013). Atlas generation on the whole body has been explored considerably less due to a much higher inter-subject variability compared to brain images. However, there are some works focusing on body MR atlas generation that have shown promising applications for these atlases (Sjöholm et al., 2019; Strand et al., 2017).

In this work, we utilise conditional atlases generated on a subset of the whole population, split by sex and BMI group (healthy, overweight, obese) (Starck et al., 2023). Consequently, we use six comprehensive whole-body atlases. For each individual, we apply Grad-CAM to generate subject-specific importance maps, which are subsequently aligned with these atlases, yielding population-wide importance maps.

## 3. Materials and Methods

### 3.1 Dataset

The UK Biobank (UKBB) dataset (Sudlow et al., 2015) is a large-scale longitudinal study that has been conducted in the UK since 2006. It contains information from approximately 100 000 participants, with a wide range of data such as genetics, biological samples and MR images from the brain, heart, and abdomen. In this work, we utilise the whole-body neck-to-knee MR images acquired with the Dixon technique for internal fat across six stations. We use the water contrast images and stitch the stations together using a publicly available tool (Lavdas et al., 2019). We select 3120 subjects with a balanced distribution across age, sex, and BMI. 1536 subjects were used for training, 384 for validation and 1200 for testing. The ages range from 46 to 81, and the mean age is 63.58 years. We ensure an equal representation of male and female subjects in all sets.

### 3.2 Training Pipeline

We train a 3D ResNet-18 model (He et al., 2016; Feichtenhofer et al., 2019) from torchvision (Paszke et al., 2019) with a hidden layer size of 256. The training is performed by using adaptive moment estimation (Adam) optimiser (Kingma and Ba, 2014) and by minimising the mean absolute error (MAE) of the age predictions. Furthermore, we use a gradient accumulation scheduler which sums and averages the gradients from 32 consecutive mini-




Figure 2: Visualisation of the population-wide Grad-CAM importance maps across several slices (columns) and of different planes (axial, coronal, sagittal) of the healthy female sub-population, overlaid on the respective atlas.

batches to update the model’s parameters. The initial learning rate is  $1e - 4$ , derived from manual tuning and reduced via scheduling when the validation error does not decrease for three epochs. The model was trained for 100 epochs, which lasted approximately 48 hours on an NVIDIA A40 GPU.

The application of Grad-CAM is independent of the training process. After training the model, we apply Grad-CAM on the third layer of the network using the implementation from (Gotkowski et al., 2020). We apply Grad-CAM at inference and on the test set to evaluate the essential body areas related to age prediction.

Additionally, following similar works in the brain, we apply a statistical bias correction method to the predicted ages to increase accuracy in the prediction and the downstream analysis. Indeed, many age estimation methods suffer from a recurring bias in the overestimation of the age in younger subjects and the underestimation in elders (Le et al., 2018; Smith et al., 2019). We use the real age of the validation data as a covariate to predict a bias-corrected age, which we then apply to the test data. An example of the predictions before and after bias correction is available in Appendix B.

### 3.3 Registration and Atlas Generation

We map all subject-level Grad-CAM maps onto an atlas to investigate the important regions for our age prediction model on a population level. Given the high variability of whole-body MR scans, we follow the pipeline proposed in Starck et al. (2023) and split all subjects into subgroups depending on their sex and BMI, following three commonly used BMI groups: healthy, overweight, and obese. The registration process is done by first registering all images of a sex and BMI group to the same target subject. We apply two methods: affine and deformable registration. Affine registration refers to a set of rigid transformations such as rotation, translation, shearing, and scaling. These types of transformations allow for a coarse alignment; they do not deform the anatomy of the given subject but only correct the overall position and orientation. The resulting images are then deformed with deformable registration for a more refined registration. This step is more localised and allows for a more detailed alignment. Both registration steps were performed using the publicly available registration tool `deepali` (Schuh et al.). All parameters are reproduced from Starck et al. (2023). Once all images are registered, the resulting deformation fields are applied to their corresponding activation map, as shown in Figure 1. Subsequently, an average map is generated from each subgroup of the dataset which serves as our population-wide importance map. We specifically generate a different overall importance map for the different sex and BMI groups, since there is high anatomical variability between these subgroups.

## 4. Results and Discussion

We here summarise the results obtained from our experiments, including the age prediction, the extraction of the Grad-CAM importance maps, and the generation of a population-wide importance map.

### 4.1 Age Prediction

We evaluate our 3D age prediction model, trained on 1536 training samples, by randomly selecting 1200 previously unseen subjects that are approximately equally distributed across all BMI and age groups. Our model achieves a mean absolute error (MAE) of **2.57** years on this test set. Table 1 summarises the model’s performance divided into the same groups that are used for the atlas generation. As baselines, we utilise a mean prediction for each group (“Mean Pred.”) and reproduce the approach from Langner et al. (2019) on our test set (“2.5D”). We here use a VGG16 and adapt the number of training samples to reach comparable performance. We can see that the model substantially outperforms the mean prediction (always predicting the mean age of the population), which indicates that it is learning meaningful information. Furthermore, we can see that the model performs best on healthy subjects and performance decreases slightly for the other BMI groups. However, the performance is pretty consistent across all BMI groups and sexes and we do not see a strong bias of the model towards different body composition values or a sex group. In comparison to the 2.5D approach, using the full 3D image volumes leads to slightly better, yet highly comparable, results for all categories apart from the overweight male subjects. We also compare the runtime and the number of training samples in Table 1. In order

Table 1: Summary of the age prediction results by sex and BMI group on the test set. All values are reported MAE scores in years. The *Overall* row reports the MAE of both sexes or all BMI groups, respectively. The score on the whole test set is underlined. The mean prediction refers to a baseline model, that always predicts the mean age of the respective subgroup.

Metric	Category	Sex	Mean Pred.	2.5D	Ours
MAE	Healthy	F	7.190	2.485	<b>2.460</b>
		M	7.672	2.614	<b>2.425</b>
	Overweight	F	7.485	2.661	<b>2.525</b>
		M	8.036	<b>2.327</b>	2.623
	Obese	F	7.045	2.863	<b>2.651</b>
		M	7.550	<b>2.669</b>	2.720
	Overall	M+F	7.499	2.613	<b>2.565</b>
	Nr. training samples	-	-	N/A	18,384
Runtime/epoch (min)	-	-	N/A	<b>1.88 ± 0.05</b>	51.33 ± 1.33
Inference/sample (s)	-	-	N/A	<b>0.024</b>	0.286

for the 2.5D approach to achieve comparable performance to the 3D method, it requires approximately 12 times more training samples. We here utilise 18,380 training samples to reach a model performance of 2.61 years on the same test set, compared to 1536 training samples for the 3D approach. However, its training time per epoch (“Runtime/epoch”) and (“Inference/sample”) is significantly faster due to the smaller data size.

A more detailed visualisation of our results is shown in Figure 4 with a scatter plot of the real age against the predicted one for the whole test set. More detailed plots on all individual groups as well as a table of the results are available in the Appendix, Section C.

## 4.2 Extraction of Grad-CAM Maps

To extract the Grad-CAM maps, we follow the original approach introduced by Selvaraju et al. (2017). We extract the importance maps from the inference runs of the 1200 test subjects (200 of each group) and register them to the subgroup atlases to obtain the population-wide attention maps shown in Figure 2. By visual assessment of these individual maps by expert radiologists, we identify three main areas of importance: (1) the spine, (2) the autochthonous muscles of the back, and (3) the heart region, including the myocardium (muscle tissue surrounding the heart) and the aortic arch. These regions are consistently highlighted over every atlas. Additional highlighted regions comprise the thyroid gland, as we can see in the top left visualisation of Figure 2, the knees (bottom left in Figure 2), the obturator muscles and the abdominal fatty tissue (middle right in Figure 2). These findings align with related work from Langner et al. (2019) as the same regions are consistently highlighted. The use of a 3D model, however, allows for leveraging the entire volume and, therefore, discovering another major region of interest: the spine. More specifically



the cervico-thoracic spine region shows strong activations. This could be due to changes in the curvature, i.e. disc degeneration and increased kyphosis of the cervico-thoracic spine, developing with age (Yukawa et al., 2012; Liu et al., 2019) (Liu et al., 2015). Additionally, structural changes in thyroid gland such as calcification and cysts are frequently seen in older patients and tend to increase in size and number with age, which could explain the activations we observe in the thyroid region. Also, degenerative changes of the main axis skeletal joints such as the knees are a frequent finding and scale with aging. The additional regions show lesser importance but are identifiable in all groups. In short, these findings concur with medical research, as these regions have demonstrated age-related impacts (Ignasiak et al., 2018; Paneni et al., 2017; Oei et al., 2022), which provides evidence that these population-wide activation maps hold great potential to investigate which areas in the body contribute most to the model’s prediction.

### 4.3 Age-specific Importance Maps

The groups selected to create the atlases, are chosen with respect to BMI and sex, following Starck et al. (2023). This categorisation aims to model a *global* representation of ageing. Our results align with the current understanding of how and where ageing is happening in the human body and we did not observe any notable differences between the different BMI and sex groups. However, we are also interested in whether there are any differences in the importance maps between different age groups. We explore this by generating population-wide importance maps for three different age groups: subjects below 60 years old, between 60 and 70 years old and above 70 years old (approximately 70 samples per group). Figure 3 visualises the importance maps for these three groups within the healthy male group. Here, we observe a noteworthy change across age groups, highlighting differences in characteristic regions of importance in each cohort. In particular, we note that the focus on the spine increases with age. Ageing comes indeed with various spine-related disorders such as degenerative scoliosis or osteoporosis (Fehlings et al., 2015), which potentially guides the model’s predictions for subjects in older age categories. Another finding is that the significance of the autochthonous back muscle region appears to grow with age, potentially due to the increased prevalence of sarcopenia in older individuals (Tournadre et al., 2019), a muscle disorder associated with ageing inducing lower muscle mass.

This analysis might indicate different ageing patterns over age, enabling the detection of age-related features specific to a subpopulation, e.g. frailty, and diseases. We see strong potential in the analysis of further subgroups and how ageing potentially impacts different regions of the body for specific diseases such as cardiovascular disorders or diabetes.

### 4.4 Chronological vs. Predicted Age Gaps

Individual importance maps provide insights into the regions that informed the model to make its prediction for a subject. We have leveraged this information to derive global features on a population level and have compared different groups of interest. These general maps emphasize strong predictors of ageing and remove any diffuse noise derived from Grad-CAM on the level of individual samples. A summary of the model performance is visible in Figure 4. The scatter plot of the predicted age versus the chronological age highlights the error as the distance to the identity function, i.e. the more distant a point is from the




Figure 3: Overview of age group interpretation. All results are shown for the healthy male subgroup. The first column is the population ages lower than 60, the second is between 60 and 70 years old, the third is above 70 and the last is across all groups. One can observe a tendency for stronger activation in the spine area for older subjects.

identity (black line), the higher the error. In some cases, the predicted age of a subject can differ significantly from their chronological age. This can be attributed to two reasons: either (a) the model fails to predict the actual age of a subject or (b) the subject exhibits signs of accelerated or decelerated ageing.

We investigate this by visualising the individual importance maps (Figure 4) for strong mis-predictions to assess whether model failure is visible in them. We visualise three examples for the most heavily accelerated (purple) and decelerated (yellow) data alongside “near perfect” predictions (red) for the overweight male group (Figure 4). The variability between these maps is quite high and aside from marginal noise in the abdominal region, no consistent deviation from the atlas is apparent in the maps. Given the fact that the importance maps do not show any sign that these are mis-predicted, we conclude that either




Figure 4: Overview of the model predictions (y-axis) versus actual age (x-axis) for all groups. Below is a visualisation of importance maps for three groups in the overweight male category: (a) subjects where the prediction was near perfect (MAE < 0.5) in red, (b) subjects where the age was underestimated (decelerated) in yellow, and (c) subject where the age was overestimated (accelerated) in purple. These examples are highlighted on the above scatter plot with their respective colours.

Grad-CAM is inadequate to reflect mis-predictions in this setting, or the model’s decision is informed by age-related features and the subjects are indeed accelerated agers.

Since consistent deviations from the atlas are difficult to detect on an individual basis, we aggregate the importance maps for the subjects where we observe accelerated and decelerated ageing. Comparing these maps, we do observe a difference for the accelerated age group; the spine and autochthonous muscles show stronger activations. This indicates that there might be physiological differences for accelerated ageing subjects compared to normal and decelerated ageing. However, validating this finding is challenging for several reasons. Firstly, the model was not exclusively trained on healthy data, which is typically done to ensure that the training subjects’ chronological age matches their biological age (Tian et al., 2023). To conduct a more thorough analysis of accelerated and decelerated ageing, a specifically designed training regime might be necessary. Additionally, the importance maps generated by Grad-CAM are not highly precise, and the strong signals from other parts of the dataset, such as the heart and spine, might obscure the detection of accelerated features. The model tends to utilise the simplest most predictive features in the input data, which is why secondary regions that still contain important information about the age of a subject might be ignored by the model. Therefore, further investigation would be needed to claim that the importance maps highlight accelerated or decelerated ageing features.

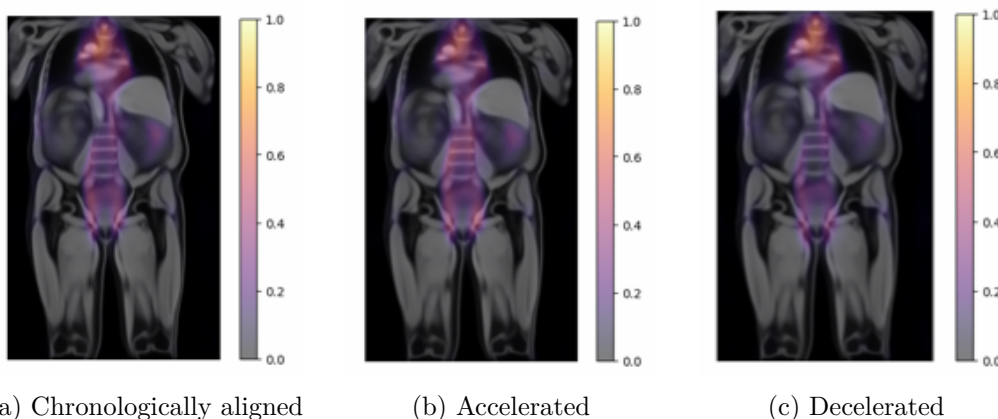


Figure 5: Visualisation of group-wise importance maps in the overweight male category for predictions that are aligned, accelerated and decelerated with respect to the chronological age. (5a) is the chronologically align atlas, (5b) the accelerated one and (5c) the decelerated one. We observe stronger activation in the spine for the accelerated age group compared to the decelerated age group.

## 5. Conclusion and Future Work

In this work, we investigate which areas in the body contribute most to our whole-body MRI age predictor. We train a 3D ResNet-18 model on 1536 neck-to-knee MR images from the UK Biobank (Sudlow et al., 2015). Our model performs whole-body age prediction with a mean absolute error of 2.57 years on the test set after bias correction. To investigate the most predictive parts of the body, we apply Grad-CAM to the gradients derived from each test subject. However, these importance maps are subject-specific and do not easily generalise to the whole population. We address this by registering all importance maps into the same coordinate space, aggregating them, and overlaying them with an atlas of whole-body MR images. We here use six distinct groups in the population, based on their sex and BMI. The aggregated importance maps for all subjects highlight three main regions of interest: the spine, the cardiac region, and the autochthonous back muscles. Despite mapping individual importance maps to the population atlas and across various groups, these areas stay consistent. We also investigate differences in importance maps across age groups and find that the model places a stronger focus on the spine and autochthonous back muscles in older subjects. This may be due to older individuals being more likely to suffer from spine-related issues, influencing the model’s predictions. In all experiments, the highlighted areas of the importance maps align with medical knowledge and previous studies on ageing. This alignment is encouraging, suggesting that examining the importance maps for pathological groups might reveal new insights into the association between DL-based age prediction and specific pathologies. Additionally, we analyse the generated importance maps for subjects that show where the deviation between model prediction and chronological age is high. We hereby distinguish between over- and under-estimated age predictions and compare the averaged importance maps of these two groups. We do not observe any significant changes in individual importance maps for these subjects. However, aggregating

them demonstrated slightly stronger activations in the spine for the accelerated ageing group. This could indicate physiological differences in accelerated ageing subjects compared to those with normal and decelerated ageing. However, further investigation would be required to confirm this finding, which we consider an interesting direction for future work.

We envision several other interesting directions of future work to further investigate highly relevant areas in the human body for DL-based age predictors. The here-generated importance maps primarily highlight specific regions of interest in the body, potentially neglecting other areas that may contain valuable information about ageing. An approach to address this, and focus on secondary body regions would be to mask the input images during training, either randomly or by deliberately omitting the most predictive areas (such as the spine, aortic arch, and back muscles). This would steer the model towards different features for its predictions, potentially discovering valuable medical insights. Furthermore, we identify the opportunity to extract even more qualitative importance maps as an interesting next step. We aim to further investigate different interpretability methods, such as perturbation-based methods (Zeiler and Fergus, 2014) or attention-based models, such as Vision Transformers (Dosovitskiy et al., 2020) and compare their results to the here utilised Grad-CAM method. Moreover, we intend to implement our method on comparable datasets like the German National Cohort (Bamberg et al., 2015) or in-house hospital data, and therefore a wider age range than the one represented in the UK Biobank, to validate the broader applicability of these findings. We believe that the here showcased move from individual insights into the decision-making process of DL methods to a population-level has great potential to further investigate the interplay between DL and medical research.

## Acknowledgments

SS and TTM were supported by the ERC (Deep4MI - 884622). This work has been conducted under the UK Biobank application 87802. SS has furthermore been supported by the Federal Ministry of Education and Research (BMBF). RB was funded by the Federal Ministry of Education and Research (BMBF, Grant Nr. 01ZZ2315B and 01KX2021), the Bavarian Cancer Research Center (BZKF, Lighthouse AI and Bioinformatics) and the German Cancer Consortium (DKTK, Joint Imaging Platform).

## Ethical Standards

The work follows appropriate ethical standards in conducting research and writing the manuscript, following all applicable laws and regulations regarding the treatment of animals or human subjects.

## Conflicts of Interest

We declare we don't have conflicts of interest.

## Data availability

The UK Biobank dataset is available on [www.ukbiobank.ac.uk](http://www.ukbiobank.ac.uk), upon registration. For this study, permission to access and analyse the UK Biobank data was approved under the application 87802

## References

- Fabian Bamberg, Hans-Ulrich Kauczor, Sabine Weckbach, Christopher L Schlett, Michael Forsting, Susanne C Ladd, Karin Halina Greiser, Marc-André Weber, Jeanette Schulz-Menger, Thoralf Niendorf, et al. Whole-body mr imaging in the german national cohort: rationale, design, and technical background. *Radiology*, 277(1):206–220, 2015.
- Kyriaki-Margarita Bintsi, Vasileios Baltatzis, Alexander Hammers, and Daniel Rueckert. Voxel-level importance maps for interpretable brain age estimation. In *Interpretability of Machine Intelligence in Medical Image Computing, and Topological Data Analysis and Its Applications for Medical Data: 4th International Workshop, iMIMIC 2021, and 1st International Workshop, TDA4MedicalData 2021, Held in Conjunction with MICCAI 2021, Strasbourg, France, September 27, 2021, Proceedings 4*, pages 65–74. Springer, 2021.
- Diogo V Carvalho, Eduardo M Pereira, and Jaime S Cardoso. Machine learning interpretability: A survey on methods and metrics. *Electronics*, 8(8):832, 2019.
- Lei Chen, Jianhui Chen, Hossein Hajimirsadeghi, and Greg Mori. Adapting grad-cam for embedding networks. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2794–2803, 2020.
- Othmane Daanouni, Bouchaib Cherradi, and Amal Tmiri. Automatic detection of diabetic retinopathy using custom cnn and grad-cam. In *Advances on Smart and Soft Computing: Proceedings of ICACIn 2020*, pages 15–26. Springer, 2021.
- MF Darmawan, Suhaila M Yusuf, MR Abdul Kadir, and H Haron. Age estimation based on bone length using 12 regression models of left hand x-ray images for asian children below 19 years old. *Legal Medicine*, 17(2):71–78, 2015.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Florian Dubost, Hieab Adams, Pinar Yilmaz, Gerda Bortsova, Gijs van Tulder, M Arfan Ikram, Wiro Niessen, Meike W Vernooij, and Marleen de Bruijne. Weakly supervised object detection with 2d and 3d regression neural networks. *Medical image analysis*, 65: 101767, 2020.
- Margaret M Esiri. Ageing and the brain. *The Journal of Pathology: A Journal of the Pathological Society of Great Britain and Ireland*, 211(2):181–187, 2007.

- Aurore Fayosse, Dinh-Phong Nguyen, Aline Dugravot, Julien Dumurgier, Adam G Tabak, Mika Kivimäki, Séverine Sabia, and Archana Singh-Manoux. Risk prediction models for dementia: role of age and cardiometabolic risk factors. *BMC medicine*, 18:1–10, 2020.
- Michael G Fehlings, Lindsay Tetreault, Anick Nater, Ted Choma, James Harrop, Tom Mroz, Carlo Santaguida, and Justin S Smith. The aging of the global population: the changing epidemiology of disease and spinal disorders, 2015.
- Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6202–6211, 2019.
- Karol Gotkowski, Camila Gonzalez, Andreas Bucher, and Anirban Mukhopadhyay. M3d-cam: A pytorch library to generate 3d data attention maps for medical deep learning, 2020.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- Yujun Hou, Xiuli Dan, Mansi Babbar, Yong Wei, Steen G Hasselbalch, Deborah L Croteau, and Vilhelm A Bohr. Ageing as a risk factor for neurodegenerative disease. *Nature Reviews Neurology*, 15(10):565–581, 2019.
- Wyke Huizinga, Dirk HJ Poot, Meike W Vernooij, Gennady V Roshchupkin, Esther E Bron, Mohammad Arfan Ikram, Daniel Rueckert, Wiro J Niessen, Stefan Klein, Alzheimer’s Disease Neuroimaging Initiative, et al. A spatio-temporal reference model of the aging brain. *NeuroImage*, 169:11–22, 2018.
- Dominika Ignasiak, Waldo Valenzuela, Mauricio Reyes, and Stephen J Ferguson. The effect of muscle ageing and sarcopenia on spinal segmental loads. *European Spine Journal*, 27:2650–2659, 2018.
- Thomas R Insel, Story C Landis, and Francis S Collins. The nih brain initiative. *Science*, 340(6133):687–688, 2013.
- Ho-Taek Joo and Kyung-Joong Kim. Visualization of deep reinforcement learning using grad-cam: how ai plays atari games? In *2019 IEEE Conference on Games (CoG)*, pages 1–2. IEEE, 2019.
- Bjarne Kerber, Tobias Hepp, Thomas Küstner, and Sergios Gatidis. Deep learning-based age estimation from clinical computed tomography image data of the thorax and abdomen in the adult population. *Plos one*, 18(11):e0292993, 2023.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
- Taro Langner, Johan Wikström, Tomas Bjerner, Håkan Ahlström, and Joel Kullberg. Identifying morphological indicators of aging with neural networks on large-scale whole-body mri. *IEEE transactions on medical imaging*, 39(5):1430–1437, 2019.

- I Lavdas, B Glocker, D Rueckert, SA Taylor, EO Aboagye, and AG Rockall. Machine learning in whole-body mri: experiences and challenges from an applied study using multicentre data. *Clinical radiology*, 74(5):346–356, 2019.
- Trang T Le, Rayus T Kuplicki, Brett A McKinney, Hung-Wen Yeh, Wesley K Thompson, Martin P Paulus, and Tulsa 1000 Investigators. A nonlinear simulation framework supports adjusting for age when analyzing brainage. *Frontiers in aging neuroscience*, 10:317, 2018.
- Alan Le Goallec, Samuel Diai, Sasha Collin, Jean-Baptiste Prost, Théo Vincent, and Chirag J Patel. Using deep learning to predict abdominal age from liver and pancreas magnetic resonance images. *Nature Communications*, 13(1):1979, 2022.
- Baoge Liu, Bingxuan Wu, Tom Van Hoof, Jean-Pierre Kalala Okito, Zhenyu Liu, and Zheng Zeng. Are the standard parameters of cervical spine alignment and range of motion related to age, sex, and cervical disc degeneration? *Journal of Neurosurgery: Spine*, 23(3):274–279, 2015.
- Jingpei Liu, Peng Liu, Zikun Ma, Jianhui Mou, Zhaolin Wang, Dong Sun, Jie Cheng, Dengwei Zhang, and Jianlin Xiao. The effects of aging on the profile of the cervical spine. *Medicine*, 98(7):e14425, 2019.
- Eileen Luders, Nicolas Cherbuin, and Christian Gaser. Estimating brain age using high-resolution pattern recognition: Younger brains in long-term meditation practitioners. *Neuroimage*, 134:508–513, 2016.
- Ariane Maggio. The skeletal age estimation potential of the knee: Current scholarship and future directions for research. *Journal of Forensic Radiology and Imaging*, 9:13–15, 2017.
- JB Antoine Maintz and Max A Viergever. A survey of medical image registration. *Medical image analysis*, 2(1):1–36, 1998.
- Henry Markram. The human brain project. *Scientific American*, 306(6):50–55, 2012.
- Jeffrey M Meier, Abass Alavi, Sireesha Iruvuri, Saad Alzeair, Rex Parker, Mohamed Houseni, Miguel Hernandez-Pampaloni, Andrew Mong, and Drew A Torigian. Assessment of age-related changes in abdominal organ structure and function with computed tomography and positron emission tomography. In *Seminars in nuclear medicine*, volume 37, pages 154–172. Elsevier, 2007.
- Tawachai Monum, Karnda Mekjaidee, Nuttaya Pattamapaspong, and Sukon Prasitwatanaseree. Age estimation by chest plate radiographs in a thai male population. *Science & Justice*, 57(3):169–173, 2017.
- Teresa Niccoli and Linda Partridge. Ageing as a risk factor for disease. *Current biology*, 22(17):R741–R752, 2012.
- Merrie W Oei, Ashley L Evens, Alok A Bhatt, and Hillary W Garner. Imaging of the aging spine. *Radiologic Clinics*, 60(4):629–640, 2022.



- Francesco Paneni, Candela Diaz Cañestro, Peter Libby, Thomas F Lüscher, and Giovanni G Camici. The aging cardiovascular system: understanding it at the cellular and clinical levels. *Journal of the American College of Cardiology*, 69(15):1952–1967, 2017.
- Harsh Panwar, PK Gupta, Mohammad Khubeb Siddiqui, Ruben Morales-Menendez, Prakhar Bhardwaj, and Vaishnavi Singh. A deep learning and grad-cam based color visualization approach for fast detection of covid-19 cases using chest x-ray and ct-scan images. *Chaos, Solitons & Fractals*, 140:110190, 2020.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- Han Peng, Weikang Gong, Christian F Beckmann, Andrea Vedaldi, and Stephen M Smith. Accurate brain age prediction with lightweight deep neural networks. *Medical image analysis*, 68:101871, 2021.
- Francesco Piccialli, Vittorio Di Somma, Fabio Giampaolo, Salvatore Cuomo, and Giancarlo Fortino. A survey on deep learning in medicine: Why, how and when? *Information Fusion*, 66:111–137, 2021.
- Vineet K Raghu, Jakob Weiss, Udo Hoffmann, Hugo JWL Aerts, and Michael T Lu. Deep learning to estimate biological age from chest radiographs. *Cardiovascular Imaging*, 14(11):2226–2236, 2021.
- Hedieh Sajedi and Nastaran Pardakhti. Age prediction based on brain mri image: a survey. *Journal of medical systems*, 43:1–30, 2019.
- Andreas Schuh, Huaqi Qiu, and HeartFlow Research. deepali: Image, point set, and surface registration in PyTorch. URL <https://github.com/BioMedIA/deepali>.
- Kirsten Seale, Steve Horvath, Andrew Teschendorff, Nir Eynon, and Sarah Voisin. Making sense of the ageing methylome. *Nature Reviews Genetics*, 23(10):585–605, 2022.
- Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- Therese Sjöholm, Simon Ekström, Robin Strand, Håkan Ahlström, Lars Lind, Filip Malmberg, and Joel Kullberg. A whole-body fdg pet/mr atlas for multiparametric voxel-based analysis. *Scientific Reports*, 9, 2019.
- Stephen M. Smith, Diego Vidaurre, Fidel Alfaro-Almagro, Thomas E. Nichols, and Karla L. Miller. Estimation of brain age delta from brain imaging. *NeuroImage*, 200:528–539, 2019. ISSN 1053-8119. . URL <https://www.sciencedirect.com/science/article/pii/S1053811919305026>.

- Sophie Starck, Vasiliki Sideri-Lampretsa, Jessica J. M. Ritter, Veronika A. Zimmer, Rickmer Braren, Tamara T. Mueller, and Daniel Rueckert. Constructing population-specific atlases from whole body mri: Application to the ukbb, 2023.
- Robin Strand, Filip Malmberg, Lars Johansson, Lars Lind, Magnus Sundbom, Håkan Ahlström, and Joel Kullberg. A concept for holistic whole body mri data analysis, imiomics. *PloS one*, 12(2):e0169966, 2017.
- Cathie Sudlow, John Gallacher, Naomi Allen, Valerie Beral, Paul Burton, John Danesh, Paul Downey, Paul Elliott, Jane Green, Martin Landray, et al. Uk biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS medicine*, 12(3):e1001779, 2015.
- Kazuto Tajiri and Yukihiro Shimizu. Liver physiology and liver diseases in the elderly. *World journal of gastroenterology: WJG*, 19(46):8459, 2013.
- Ye Ella Tian, Vanessa Cropley, Andrea B Maier, Nicola T Lautenschlager, Michael Breakspear, and Andrew Zalesky. Heterogeneous aging across multiple organ systems and prediction of chronic disease and mortality. *Nature medicine*, 29(5):1221–1231, 2023.
- Anne Tournadre, Gaelle Vial, Frédéric Capel, Martin Soubrier, and Yves Boirie. Sarcopenia. *Joint bone spine*, 86(3):309–314, 2019.
- David C Van Essen, Stephen M Smith, Deanna M Barch, Timothy EJ Behrens, Essa Yacoub, Kamil Ugurbil, Wu-Minn HCP Consortium, et al. The wu-minn human connectome project: an overview. *Neuroimage*, 80:62–79, 2013.
- Meenal Verma, Nikhil Verma, Rakhee Sharma, and Ashish Sharma. Dental age estimation methods in adult dentitions: An overview. *Journal of forensic dental sciences*, 11(2):57, 2019.
- Fei Wang, Lawrence Peter Casalino, and Dhruv Khullar. Deep learning in medicine—promise, progress, and challenges. *JAMA internal medicine*, 179(3):293–294, 2019.
- JM Wishart, AO Need, M Horowitz, HA Morris, and BEC Nordin. Effect of age on bone density and bone turnover in men. *Clinical endocrinology*, 42(2):141–146, 1995.
- Mengying Xiao, Liyuan Zhang, Weili Shi, Jianhua Liu, Wei He, and Zhengang Jiang. A visualization method based on the grad-cam for medical image segmentation model. In *2021 International Conference on Electronic Information Engineering and Computer Science (EIECS)*, pages 242–247. IEEE, 2021.
- Yasutsugu Yukawa, Fumihiko Kato, Kota Suda, Masatsune Yamagata, and Takayoshi Ueta. Age-related changes in osseous anatomy, alignment, and range of motion of the cervical spine. part i: Radiographic data from over 1,200 asymptomatic subjects. *European Spine Journal*, 21:1492–1498, 2012.
- Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part I 13*, pages 818–833. Springer, 2014.

## Appendix A. Atlases

Figure 6 shows the population-wide Grad-CAM maps for both sexes and all BMI groups. We here selected one coronal and one sagittal slice to show the main activations in both views. All activation maps highlight similar regions as discussed in the main part of the paper and we do not observe a meaningful difference between sexes or BMI groups.




Figure 6: Overview of the population-wide Grad-CAM activation maps across all categories, overlaid on the respective atlas.

## Appendix B. Bias correction

Similar to many other age prediction works Tian et al. (2023); Le et al. (2018); Smith et al. (2019), we apply a bias correction to the initial age predictions of our model. The scatter plots of the predictions before and after bias correction are visualised in Figure 7.




Figure 7: Visualisation of the error on the test set before and after bias correction. The left scatter plot shows the data before correction and the bias is clearly visible and the right plot shows the performance after correction.

## Appendix C. Summary performance

Figure 8 shows the scatter plots of the chronological age against the bias-corrected predicted age for all subgroups. These results are furthermore summarised in Table 1 in the main manuscript listing the respective mean absolute errors.

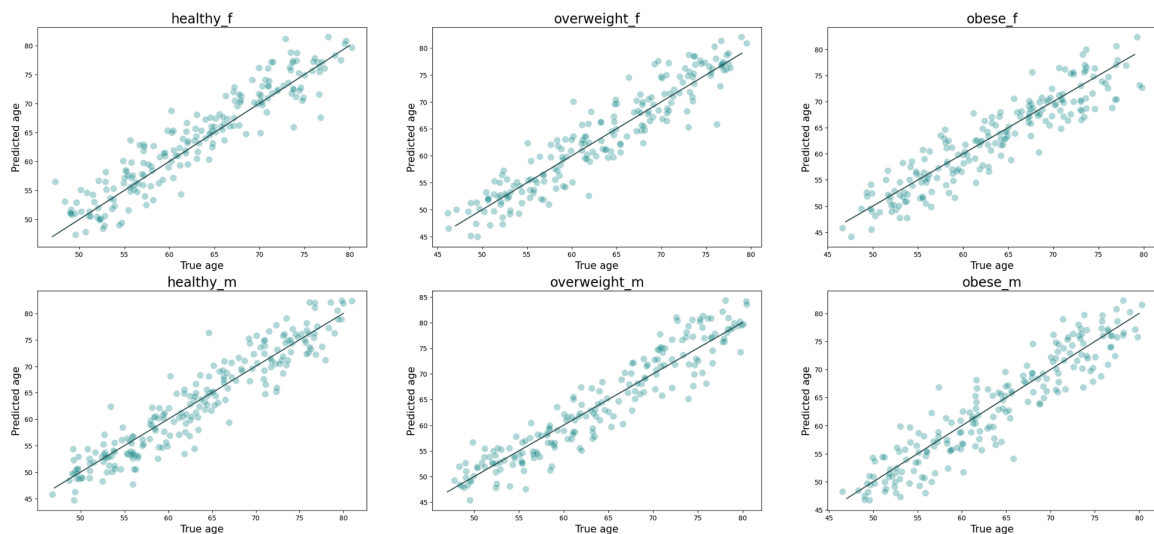


Figure 8: Overview of the bias-corrected age prediction of the model versus actual age for all groups. The top row shows the performance for the female group and the second for males for respectively the healthy, overweight and obese groups

## Appendix D. Comparison to 2.5D

We compare our approach to a less resource-hungry one that works on projected images from the 3D volume, which we call 2.5D. We follow the approach from Langner et al. (2019). Two example 2.5D images and their corresponding importance maps are visualised in Figure 9. We achieve very similar attention maps as reported in their work, however, highlight that the as important considered regions are difficult to assign to specific regions in 3D space due to the nature of the projected images.




Figure 9: Visualisation of two example test subjects and their interpretable Grad-CAM maps for the 2.5D approach. Here the heart region and the knee area seem to contribute mostly to the model’s prediction. However, it is difficult to identify which specific regions in 3D space are the most relevant, as they are all projected to the same plane.