

OPEN

Large-scale biometry with interpretable neural network regression on UK Biobank body MRI

Taro Langner¹✉, Robin Strand^{1,2}, Håkan Ahlström^{1,3} & Joel Kullberg^{1,3}

In a large-scale medical examination, the UK Biobank study has successfully imaged more than 32,000 volunteer participants with magnetic resonance imaging (MRI). Each scan is linked to extensive metadata, providing a comprehensive medical survey of imaged anatomy and related health states. Despite its potential for research, this vast amount of data presents a challenge to established methods of evaluation, which often rely on manual input. To date, the range of reference values for cardiovascular and metabolic risk factors is therefore incomplete. In this work, neural networks were trained for image-based regression to infer various biological metrics from the neck-to-knee body MRI automatically. The approach requires no manual intervention or direct access to reference segmentations for training. The examined fields span 64 variables derived from anthropometric measurements, dual-energy X-ray absorptiometry (DXA), atlas-based segmentations, and dedicated liver scans. With the ResNet50, the standardized framework achieves a close fit to the target values (median $R^2 > 0.97$) in cross-validation. Interpretation of aggregated saliency maps suggests that the network correctly targets specific body regions and limbs, and learned to emulate different modalities. On several body composition metrics, the quality of the predictions is within the range of variability observed between established gold standard techniques.

As part of the UK Biobank study¹ 100,000 volunteer participants are to be examined with magnetic resonance imaging (MRI). Among the scheduled imaging protocols is neck-to-knee body MRI, resulting in volumetric images with separate water and fat signal. These scans contain comprehensive information about the anatomy of each subject and are accompanied by a wide range of other collected metadata, spanning anthropometric measurements, questionnaires, biological samples, health outcomes, and more. Many of these properties also express themselves in the morphology of the human body and could potentially be inferred with machine learning. Techniques involving neural networks for image-based regression have been previously proposed for the analysis of brain MRI for detection of premature ageing², early symptoms of Alzheimers disease³ and mental disorders⁴. In heart MRI, related approaches were able to perform measurements of volumes and wall thicknesses of the heart⁵. Similarly, analyses of retinal fundus photographs showed that neural networks were able to leverage image features for the prediction of properties including age, gender, smoking status and blood pressure⁶. Many of these findings were unexpected as the underlying features are often not easily accessible even to human experts.

Research in metabolic and cardiovascular disease has led to increased interest in strategies for the automated analysis of body composition⁷. Individualized measurements of fat and muscle compartments in the body have the potential to provide new insight into the development of various medical conditions at greater detail than analyses based on anthropometric measures such as the body mass index (BMI)⁸. The amount of visceral adipose tissue in particular varies substantially between individuals and is directly related to cardiac and metabolic risk⁹. A more fine-grained analysis is of interest in research such as within the UK Biobank study itself¹⁰ but also as a potential tool for disease screening and individualized treatments. Several imaging techniques exist for the measurement of body fat, including computed tomography (CT) and dual-energy X-ray absorptiometry (DXA)¹¹ based on two-dimensional coronal projections. Chemical-shift encoded water-fat MRI acquires separate volumetric water and fat signal images which have the potential to allow for measurements without ionizing radiation, but can be challenging to evaluate. Various methods have been proposed for the delineation

¹Department of Surgical Sciences, Uppsala University, 751 85 Uppsala, Sweden. ²Department of Information Technology, Uppsala University, 751 85 Uppsala, Sweden. ³Antaros Medical AB, BioVenture Hub, 431 53 Mölndal, Sweden. ✉email: taro.langner@surgsci.uu.se

of individual adipose tissue depots in these images¹². Among other techniques, automated image analysis with convolutional neural networks for segmentation has become an established technique for images of this kind^{13,14} as well as for CT images^{15,16}. However, these systems learn to perform segmentation from training data in the form of reference segmentations, which must accordingly be carefully prepared, often with substantial amounts of manual guidance.

In this work, automated biometry is performed by training neural networks for image-based regression on UK Biobank neck-to-knee body MRI. The proposed approach extends a previously presented method for age estimation¹⁷ and requires no manual intervention or direct access to ground truth segmentation images. Instead, arbitrary numerical values can be inferred, ranging from anthropometric measurements to body composition metrics from dual-energy X-ray absorptiometry (DXA), multi-atlas-based MRI segmentations, dedicated liver scans and various other sources. The goal of this approach is to approximate all of these measurements with a fast and accurate, fully automated technique from the MRI data.

The following contributions are made:

- Extension of a framework for age estimation from UK Biobank neck-to-knee body MRI¹⁷
- Inference of 64 biological metrics (beyond just age)
- Design of an optimized and standardized configuration
- Extensive validation of both framework and predictions
- Aggregated saliency analysis¹⁷

To our knowledge, no comparable technique with convolutional neural network regression has been previously applied to neck-to-knee or whole-body MRI for inference of biological metrics other than age. Essential code, documentation and Supplementary Material has been made available for reproducibility and further use¹⁸.

Methods

A fixed configuration of a convolutional neural network for image-based regression was trained in cross-validation on two-dimensional representations of the neck-to-knee body MRI. For each of the 64 examined properties, the network was evaluated based on the generated predictions and saliency maps which highlight relevant image features.

Image data. Of the 100,000 MRI scans planned by the UK Biobank study, 32,323 were made available for the experiments in this work as part of application 14237. UK Biobank recruitment was organized by letter from the National Health Service and the vast majority of participants (94%) self-reported white British ethnicity in the initial assessment visit. All scans were acquired by the UK Biobank at three different centres in the United Kingdom in an imaging time of about six minutes each, using a dual-echo Dixon technique¹⁹ on a Siemens Aera 1.5T device. The resulting image data typically covers the body from neck to knee in six separate stations, whereas the arms and other parts of the body that extend laterally are usually not visible or subject to heavy distortion and artefacts²⁰. For the experiments in this work, those scans that contained water-fat swaps and other artefacts such as excessive noise, unusual positioning and artificial knee replacements were excluded by visual inspection of the projections by one operator, leaving 31,172 images for training and validation. The volumetric scan stations for a given subject were resampled to a resolution of 2.23 mm × 2.23 mm × 3 mm and fused into a volume of 370 × 224 × 174 voxels. This MRI volume was then cropped and compressed into a two-dimensional format of slightly lower resolution, showing a frontal and lateral projection of mean intensity, with a separate image channels for the water and fat signal. In this format, each subject was accordingly represented by a two-channel image of 256 × 256 pixels, as seen in Fig. 1, stored in 8bit format for easier processing by the neural network.

Biological metrics. From the thousands of non-imaging properties collected in the UK Biobank study, a subset of 64 fields with relevance for cardiovascular and metabolic disease was chosen. More than half of the chosen fields are measurements of body composition by DXA imaging^{11,22}, comprising mass and percentages of fat and lean tissue in the abdomen, trunk, arms and legs. The second largest group of measurements is based on multi-atlas segmentations of the neck-to-knee body MRI itself^{20,23,24} and describe volumes of adipose tissue depots and muscle groups in the abdomen, trunk and thighs. An additional group of fields contains the basic features of age, sex (1 for male, 0 for female), height, and weight. Due to privacy concerns, the age could only be calculated to an accuracy of about 15 days, based on the year (field 34) and month of birth (field 52) as well as the MRI scanning metadata (field 20201)¹⁷. The last group of fields contains values such as circumferences of the hip and waist, BMI, the percentage of fat accumulated in the liver, determined by dedicated liver MRI²⁵, the pulse rate on the imaging visit, and the measured grip strength of the right hand, which is often used as a biomarker for cardiovascular health. Of the 32,323 imaged subjects, only 3,048 have valid entries for all of the chosen fields. These subjects serve as a basis for the saliency analysis, described later in this chapter. A feature space of the 64 chosen metadata fields for these subjects is also visualized in Fig. 2 and showcases some of the underlying patterns relating to sex and body composition.

Using one standardized configuration, a dedicated neural network was trained to predict each of these 64 measurements separately. Each of them was evaluated in 7-fold cross-validation, so that all of those subjects with a valid entry for the given measurement were split into 7 subsets of equal size. By exempting each subset in turn from training and using it to make predictions which could then be compared to the reference, the network was effectively validated against all subjects without being able to memorize their values in training.

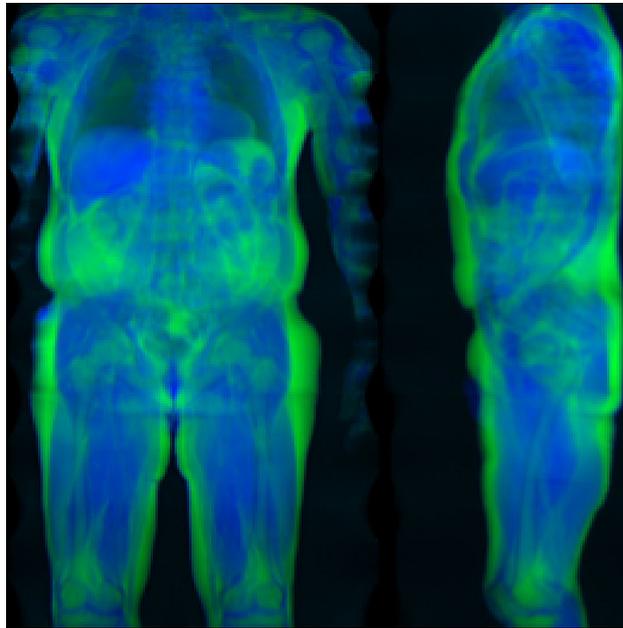


Figure 1. Two-dimensional representation of the volumetric MRI data, which serves as input to the neural network. The water (blue) and fat (green) signal images are projected along the coronal and sagittal plane and combined as color channels.

Network configuration. For each of the chosen fields a separate convolutional neural network was trained for regression in 7-fold cross-validation. The entire configuration of the network was fixed and no attempt was made to achieve better performance by tuning the network architecture or other parameters. Each unique training sample represents one subject and consists of two-dimensional format as extracted from the MRI data as input image and their field entry in the UK Biobank as numerical ground truth target value.

The neural network is a computational model that uses millions of variable parameter weights to convert an input image into one or more numerical output values. During training, it can learn to perform a certain task by making image-based predictions for samples with known reference values. The difference between prediction and reference is quantified by a loss function, and mathematical optimization involving its gradient adjusts the network parameters. In this way, parameter values can be learned that define convolutional image filters for extraction of relevant gradients, corners and edges from the image, which are subsequently formatted into increasingly abstract features that enable the network to infer the desired measurement. This process is entirely data driven and fully automated.

The previously presented regression pipeline¹⁷ for age estimation was optimized in several ways in order to process all of the chosen fields in a viable time frame. The main change consists in replacing the VGG16 architecture²⁶ with the more lightweight ResNet50²⁷. Furthermore, all numerical target values were standardized by subtracting the mean value and dividing by the standard deviation, as the ResNet50 proved more sensitive to variation in target scaling and shifts. This step resulted in faster convergence and improved stability, so that the total number of iterations could be vastly reduced from 80,000 iterations to just 6,000. To alleviate a tendency of the network to overfit in the final 1,000 iterations, the learning rate of 0.0001 in this phase is reduced by factor ten, typically resulting in a further slight increase in accuracy. Compared to the original configuration, the total training time for a given field was thus reduced by about factor 30, while reaching comparable accuracy. The original batch size of 32 and augmentation by random translations of up to 16 pixels were retained, with the nearest pixel values being repeated at the borders. All networks were trained on a Nvidia GTX 1080 Ti 11GB graphics card in the framework PyTorch with a mean squared error loss, the optimizer Adam, and parameters pretrained on ImageNet. Each split required less than 25 minutes of training time.

These design choices were made based on preliminary results for three representative fields: Age, liver fat (field 22402) and visceral adipose tissue volume (VAT) (field 22407). All presented results were achieved with this exact network configuration, without early stopping, hyperparameter tuning, or any other attempt to adapt to individual fields for better performance.

Evaluation. The chosen fields range from volumes to circumferences and simple binary labels, all treated as continuous numerical values. The neural network was trained to predict these values in regression, thereby emulating the reference, and the coefficient of determination R^2 is reported to rate the quality of fit, ranging from 1.0 for a perfect fit to negative values where the non-linear network model performs worse than simply estimating the mean. Additionally, the 95% limits of agreement (LoA) and the mean absolute error (MAE) are provided. In some cases the network output was thresholded to mimic a classification, with a threshold of 0.5 for prediction of sex and 5.5% for fatty liver disease. Without taking the exclusion criteria into account, the reference

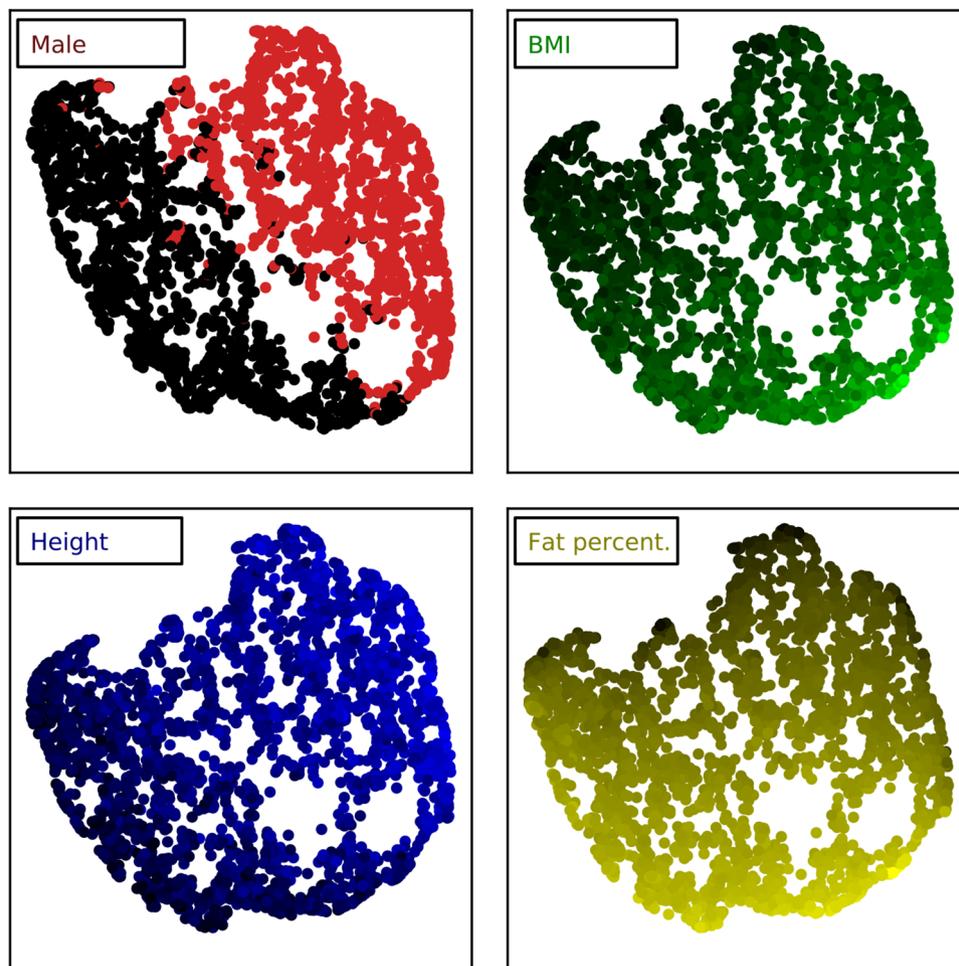


Figure 2. Visualized feature space of body composition, with 64 chosen fields of 3000 subjects. The 2d UMAP projection²¹ minimizes the distance between similar subjects, each represented by one dot. For each subplot, the named variable was removed before dimensionality reduction and used for colouring, with brighter intensity for higher values. The sexes (field 31, 0 for females and 1 for males) distinctly divide in two hemispheres whereas the ranges of BMI (field 21001), height (field 12144) and body fat percentage (field 23281) vary systematically.

liver fat values of 898 of 4219 subjects exceed this threshold. For the prediction, an area under curve (AUC) of the receiver operating characteristic (ROC) curve was calculated.

In some cases, competing measurements of the same property are available from several reference methods, so that their mutual agreement can be compared to the network performance. In the scope of this work, only the atlas-based MRI segmentations²⁴ and measurements from DXA²² are considered in this regard. Both methods examine different regions of interest and therefore show systematic differences. The MRI-based values were therefore first fit to the DXA values by linear regression before reporting their agreement in this analysis. Similarly, many fields describe features specific to the left and right side of the body. Again, the network performance can be put into the context of this inherent bilateral symmetry, but this analysis is abbreviated to report Pearson's coefficient of correlation r only.

In addition to statistical measures, an interpretation of the criteria learned by the network can be attempted with saliency analysis. For each input image, a heat map of relevant image features can be generated using guided gradient-weighted class activation maps^{28,29}. The resulting visualizations were combined by co-registration of subjects³⁰, yielding aggregated saliency maps that describe which image regions on average had the highest impact on the network prediction¹⁷ for an entire cohort of subjects. Each saliency map was generated by the one network that used the corresponding subject as a validation sample in cross-validation. When visualized, the saliency intensities were squared and overlaid as a heat map over the water signal image, without any further post-processing or manual adjustment.

Some properties could be trivial to predict due to strong correlations with simple non-image features such as age and weight. We therefore also provide the results of multiple linear regression based on the age, sex, height and weight as a baseline for comparison with the neural network performance.

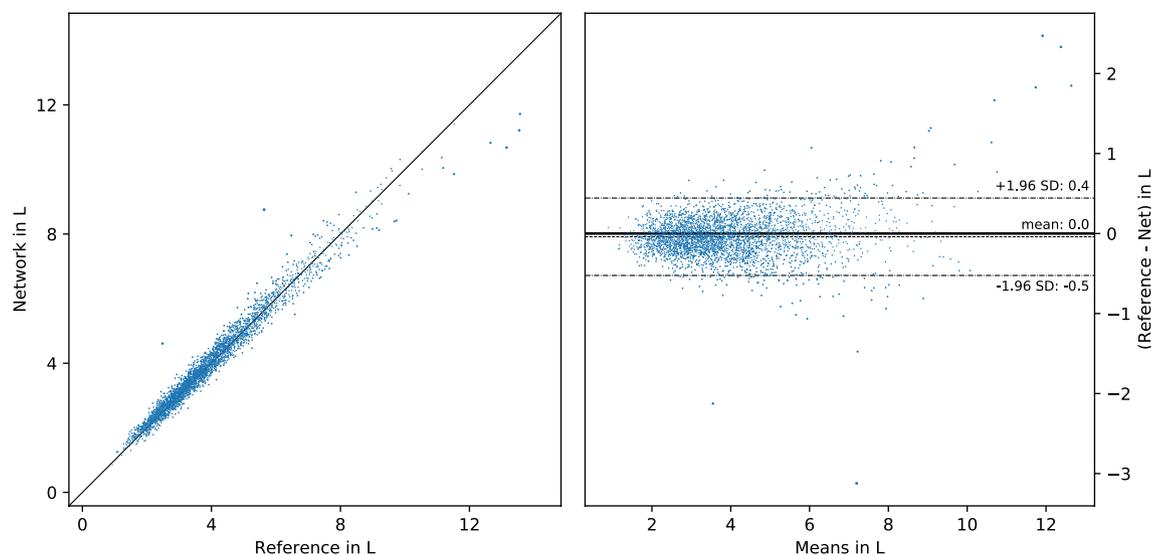


Figure 3. Visualization of the median result with $R^2 = 0.972$, the inference of fat mass in the left leg as measured by dual-energy X-ray absorptiometry (DXA) (field 23266). The diagonal line represents a hypothetical perfect fit, whereas dashed lines represent the 95% limits of agreement (LoA). Point sizes vary for better visibility.

Field	Name	N	Unit	[Min, max]	Mean \pm SD	MAE	LoA	R^2
–	Age	31,172	Years	[44.6, 82.3]	63.9 \pm 7.5	2.46	(– 5.85 to 6.31)	0.829
31	Sex	31,172	(Male = 1)	[0.0, 1.0]	0.5 \pm 0.5	0.00	(– 0.04 to 0.04)	0.999
12144	Height	31,172	cm	[140.0, 204.0]	170.2 \pm 9.4	1.70	(– 4.78 to 4.44)	0.938
21002	Weight	30,382	kg	[39.6, 169.2]	76.5 \pm 15.1	0.78	(– 1.95 to 2.24)	0.995

Table 1. Inference of basic fields. *SD* standard deviation, *MAE* mean absolute error, *LoA* limits of agreement.

Results

A close regression fit is achieved on almost all examined fields. On average, less than 3% of variability in the reference measurements remains unexplained by the network output alone (median $R^2 = 0.972$) and the linear regression baseline was outperformed in all cases. The field with median fit is shown in Fig. 3, and more plots for all fields are available in the Supplementary Material¹⁸. Table 1 lists the basic fields with a MAE of about 2.5 years for age, 0.8kg for body weight and 1.7cm for height. When thresholded, the classification accuracy for the prediction of sex reaches 99.97%, so that only 10 of 31,172 subjects were misclassified.

Some of the most accurate predictions were made for body composition as measured by atlas-based segmentation on MRI (median $R^2 = 0.987$), with a corresponding MAE of 140 mL for visceral adipose tissue (VAT), 220 mL for subcutaneous abdominal adipose (ASAT), and 180 mL for total thigh muscle volume. Additional statistical metrics for these fields and others including those from DXA and liver fat are provided in Supplementary Tables 1, 2, 3, and 4. The lowest performance was achieved on grip strength and pulse rate, where the network nonetheless managed to make a weak, image-based prediction from the MRI. When thresholded at 5.5% to identify subjects with high liver fat, the predictions reached an accuracy of 90%, with a sensitivity of 73%, specificity of 95% and an AUC-ROC of 0.943. Even though the arms are usually not visible in the images, the network succeeded in estimating the grip strength of the right hand with an MAE of about 5kg and furthermore gave a rough estimate of the pulse rate.

Saliency analysis. Examples for saliency maps generated by the network are shown in Fig. 4. The saliency indicates that the network on average correctly targets specific structures on the left or right side of the body. Moreover, the estimate of liver fat appears to be mostly based on image areas with actual liver tissue, whereas the prediction of the pulse rate takes into account features of the heart. The BMI appears to be mostly estimated from the knees and lungs, and the grip strength of the right hand is inferred from features of the corresponding side of the upper body. Complete visualizations of all saliency maps are provided in the Supplementary Material¹⁸.

Agreement between modalities. Measurements from DXA are compared to those derived from atlas-based segmentations of the MRI in Table 2. Each listed comparison yielded lower agreement between these reference methods than achieved by the specific network predictions, evaluated in Supplementary Tables 1 and 2. Although only a one-way fitting of MRI to DXA is shown, this analysis was performed in both directions and

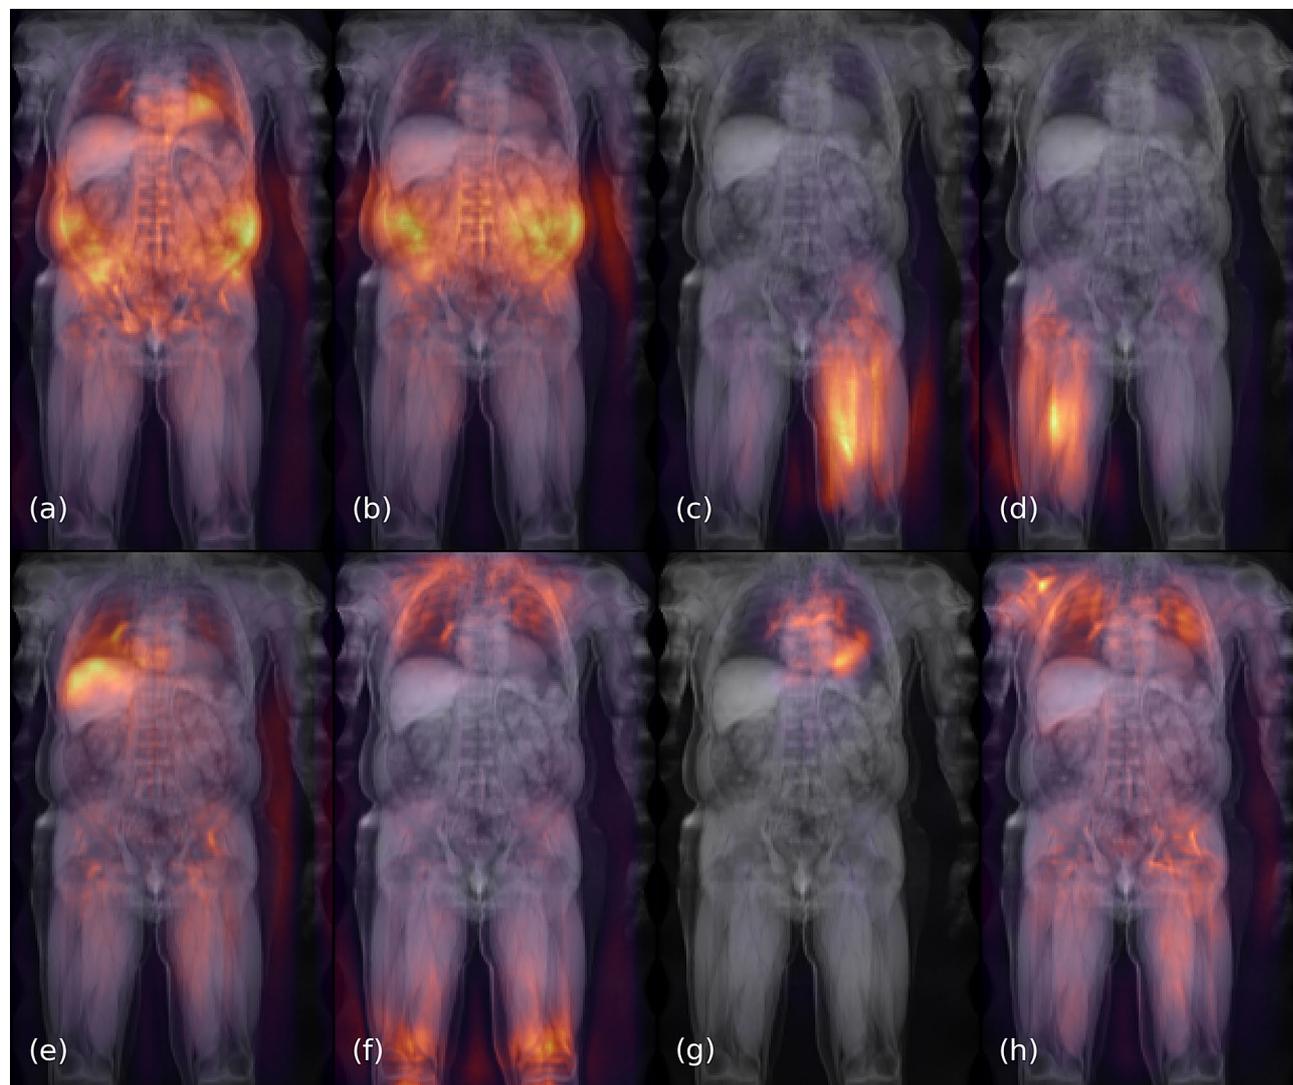


Figure 4. Aggregated saliency of about 3000 subjects for: VAT as derived from atlas-based MRI segmentations (field 22407) (a) or DXA (field 23289) (b), muscle volumes of the anterior left (field 22405) (c) and right thigh (field 22403) (d), liver fat (field 22402) (e), BMI (field 21001) (f), pulse rate (field 102) (g) and grip strength (field 47) (h). The network appears to emulate regions of interest used by different modalities and correctly targets specific limbs and organs.

Property	Field _{MRI}	Field _{DXA}	N	Term	Unit	MAE	LoA	R ²
VAT	22407	23289	4,494	$0.42x - 0.3$	L	0.17	(- 0.45 to 0.45)	0.939
Total trunk fat	22410	23284	4,538	$1.35x + 0.6$	kg	0.75	(- 1.93 to 1.93)	0.971
Total fat tissue	22415	2378	4,326	$1.23x - 0.2$	kg	0.85	(- 2.30 to 2.30)	0.981
Total lean tissue	22416	23280	4,326	$1.88x + 1.9$	kg	1.68	(- 4.46 to 4.46)	0.937

Table 2. Agreement between reference methods. Transforming measurements from atlas-based MRI segmentations of N subjects to their corresponding values from DXA yields the listed linear regression term, mean absolute error (MAE), and limits of agreement (LoA). In all cases, this agreement is exceeded by the accuracy of their individual network predictions.

yielded average LoA between both methods that are 70% wider on average than the LoA between each field and its network predictions.

Bilateral symmetry. In some cases the accuracy of the network predictions also exceeds the inherent, bilateral symmetry of the human body. For a given property, one limb is accordingly more dissimilar to the opposite limb than to its prediction by the network. A field-wise comparison with Pearson r is reported in Supplementary

Table 5. For atlas-based measurements from MRI, the average bilateral correlation for the anterior and posterior thigh muscle volume amounts to $r = 0.979$. The network predictions correlate more strongly with the left- and right-specific measurements for an average $r = 0.989$. For DXA, however, the specific prediction accuracy of the network is lower than the bilateral symmetry, with averages of $r = 0.975$ vs $r = 0.954$ for the arms and 0.987 vs 0.983 for the legs. Although some individuals show strong unilateral atrophy, this effect is not just due to outliers. The fact that the network learned to specifically target either side of the body is also visible in the saliency maps of Fig. 4 and occurs in both the DXA and MRI-based fields.

Discussion

The neural network configuration showed robust performance and closely emulated the chosen measurements by image-based regression on the MRI data, with a median R^2 above 0.97. It not only learned to accurately estimate volumes and circumferences from the simplified, two-dimensional image format, but also to emulate different modalities and make measurements specific to either side of the body. The linear regression baseline was outperformed in all cases and indicates that most of these properties can not be trivially deduced from the basic characteristics of age, sex, height, and weight.

When used to infer metrics related to body composition, the network yielded more faithful approximations of the atlas-based measurements from MRI or DXA than obtained by substituting these two reference methods for each other. This was still the case even after fitting both reference methods to each other with linear regression. The agreement for both modalities on the UK Biobank reported in previous work²⁴ yielded similar error bounds, for a sample with considerable overlap to the subjects examined here. The atlas-based method on MRI has also been previously compared to an alternative method based on T1-weighted images²³, yielding LoA for VAT, ASAT and total trunk fat that are on average more than twice as wide as those relative to the network here. The variability between these two established reference methods can largely be accounted for by differing regions of interest. Whereas the atlas-based method measures VAT up to the thoracic vertebrae Th9²⁰, DXA defines VAT as ranging from the top of the iliac crest up to 20% of the distance to the base of the skull²⁴. This is reflected in the saliency maps of Fig. 4, indicating that the network correctly learned to emulate the different criteria, based on the numerical target label alone.

Many of the most accurate predictions were made for the atlas-based measurements on MRI, where the accuracy of the network also exceeds the inherent similarity in muscle volumes between the left and right leg. There are several possible explanations for this. In contrast to the DXA-based values, these reference measurements were originally performed on the same MRI data that served as a basis for the presented method. The lack of outliers in the reference suggests high quality, closely representing an objective truth that is contained in these images. Furthermore, all images with ground truth values passed the quality control steps applied by the reference. The network was accordingly trained and evaluated on samples that were preselected regarding suitability for body composition analysis. The measurements of the arms and legs from DXA, in contrast, contain outliers and are often based on anatomy that is not entirely contained in the field of view. Including additional imaging stations that would cover the lower legs and head could lead to both more robust inference and better agreement with the DXA measurements, but was rejected during the original study design due to the prohibitive total increase in scan time²⁰. Future studies may benefit from a less targeted acquisition and instead choose to collect less restricted, more comprehensive data, as increasingly powerful tools for automated analysis become available.

Despite being able to use the same MRI data and producing similar measurements, the proposed technique and the atlas-based reference method²⁰ differ substantially in their approach. The network generates no segmentations for manual refinement or quality control. It furthermore requires hundreds or thousands of labelled ground truth images for training and would likely require retraining for different imaging devices and demographics. The atlas-based method relies on just 31 prototype subjects and has been credited for robustness towards different imaging devices and field strengths. In turn, the network can analyse several scans within just seconds instead of minutes and requires no manual intervention or guidance, so that it can easily be scaled to process tens of thousands of subjects. Even though no segmentations are generated, there is also no restriction on using only segmented images as input, but instead arbitrary numerical target labels can be used. This makes it possible to examine more abstract properties, such as grip strength and pulse rate, and to link them to relevant anatomical regions by saliency analysis.

One limitation of this work consists in the lack of an independent test set. This means that it remains unclear whether the already trained networks would reach similar performance on data from other studies and sources. As the used data has been gathered at three different imaging centres, it at least appears that the protocol can be reproduced sufficiently well at different sites to allow for robust performance on future UK Biobank images of the same population. When applied to data from other studies, such as for example the whole-body MRI scans of the German National Cohort³¹, systematic differences in subject demographics, scanning device or protocol are likely to limit the performance however, and retraining of the networks would almost certainly be necessary. The lack of an independent test set might also raise concerns about the network configuration being excessively adapted to the given data. It could be assumed that the repeated runs of the cross-validation during the preliminary experiments may have resulted in design choices that merely represent a coincidental optimum on the cross-validation data itself, with low ability to generalize and possible dependence on confounding factors in the images. However this effect is unlikely to play a significant role since all design choices were based on preliminary experiments on the fields for age, liver fat (field 22402) and VAT (field 22407) only. The resulting configuration is robust without any individual adjustment for a large variety of measurements with tens of thousands of subjects, so that it is exceedingly unlikely that the high performance is coincidental or based on simple confounding effects alone.

Many properties could potentially be predicted with greater accuracy by using customized image formats and more training samples. The resampled, two-dimensional projection effectively compresses the volumetric MRI data by factor 220 and is furthermore encoded to 8bit only. Despite the computational benefits there is no reason to assume that this format is optimal for all examined fields. Among its limitations, the separately normalized water and fat signal only enable an indirect inference of fat fraction values. When inferring these values for certain tissues and organs, the signals are furthermore conflated along the axis of projection. Future work will explore ways to make this information more accessible to the network, which is likely to benefit especially the inference of liver fat. Despite the limitations of the dual-echo Dixon technique for this purpose³², these improvements may ultimately yield higher agreement than observed between other methods such as biopsy and magnetic resonance spectroscopy³³.

When compared to the previous configuration for age estimation¹⁷, the network for age was trained in cross-validation with about 28% more data. The mean absolute error accordingly decreased as expected, from a previous 2.49 years to 2.46 years, roughly following the previously reported relationship between performance and quantity of training data. The ResNet50 performs similar to the VGG16 when using standardization of the target values, but at far higher speed. Its main disadvantage consists in more diffuse saliency maps, possibly due to the final average pooling layer.

The results show that the presented approach can leverage the two-dimensional representation of MRI image data to estimate not only the age but also to emulate a wide range of other measurements for subjects of the UK Biobank. Given only an abstract, numerical target value and the vast amount of images, the regression network learned to identify the correct body region, tissue or limb as used by the reference methods. In its current form the method could be used as a fully automated tool for approximation of missing values for those subjects who have not yet undergone all of the planned examinations. These estimates could then serve for quality control and as a basis for preliminary analyses, months or years before the established gold standard methods have been fully applied. Future work will consist in making the results accessible to the medical community and improving individual measurements with specialized input formats and network configurations, as well as exploring the limits of which other, more abstract properties can be predicted from these scans. Similar approaches could potentially enable the prediction of more variables such as blood biochemistry, disease states, and genetic markers.

Conclusion

The neural network can perform fully automated inference on the UK Biobank MRI data and learned to emulate measurements from DXA, atlas-based segmentations, dedicated liver scans and more in a fast and lightweight, standardized configuration. Saliency and correlation analysis indicate that the network can specifically target the left and right side of the body and identify relevant organs and body regions. Given enough training data for a given demographic and a standardized imaging protocol, further development may ultimately enable fully automated measurements of a wide range of biological metrics from a single 6-minute neck-to-knee body MR image.

Received: 20 April 2020; Accepted: 5 October 2020

Published online: 20 October 2020

References

1. Sudlow, C. *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 Med.* **12**, e1001779. <https://doi.org/10.1371/journal.pmed.1001779> (2015).
2. Cole, J. H. *et al.* Predicting brain age with deep learning from raw imaging data results in a reliable and heritable biomarker. *NeuroImage* **163**, 115–124. <https://doi.org/10.1016/j.neuroimage.2017.07.059> (2017).
3. Ding, Y. *et al.* A deep learning model to predict a diagnosis of Alzheimer disease by using 18F-FDG pet of the brain. *Radiology* **290**, 456–464 (2018).
4. Shahab, S. *et al.* Brain structure, cognition, and brain age in schizophrenia, bipolar disorder, and healthy controls. *Neuropsychopharmacology* **44**, 898 (2019).
5. Xue, W., Islam, A., Bhaduri, M. & Li, S. Direct multitype cardiac indices estimation via joint representation and regression learning. *IEEE Trans. Med. Imaging* **36**, 2057–2067 (2017).
6. Poplin, R. *et al.* Prediction of cardiovascular risk factors from retinal fundus photographs via deep learning. *Nat. Biomed. Eng.* **2**, 158 (2018).
7. Thomas, E. L., Fitzpatrick, J., Malik, S., Taylor-Robinson, S. D. & Bell, J. D. Whole body fat: content and distribution. *Prog. Nucl. Magn. Reson. Spectrosc.* **73**, 56–80 (2013).
8. Prentice, A. M. & Jebb, S. A. Beyond body mass index. *Obes. Rev.* **2**, 141–147 (2001).
9. Neeland, I. J. *et al.* Associations of visceral and abdominal subcutaneous adipose tissue with markers of cardiac and metabolic risk in obese adults. *Obesity* **21**, E439–E447 (2013).
10. Linge, J. *et al.* Body composition profiling in the UK Biobank Imaging Study. *Obesity* **26**, 1785–1795 (2018).
11. Kaul, S. *et al.* Dual-energy X-ray absorptiometry for quantification of visceral fat. *Obesity* **20**, 1313–1318 (2012).
12. Hu, H. H., Chen, J. & Shen, W. Segmentation and quantification of adipose tissue by magnetic resonance imaging. *Magn. Reson. Mater. Phys. Biol. Med.* **29**, 259–276 (2016).
13. Estrada, S. *et al.* Fatsegnet: a fully automated deep learning pipeline for adipose tissue segmentation on abdominal dixon MRI. *Magn. Reson. Med.* <https://doi.org/10.1002/mrm.28022>. <https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.28022>.
14. Langner, T. *et al.* Fully convolutional networks for automated segmentation of abdominal adipose tissue depots in multicenter water-fat MRI. *Magn. Reson. Med.* **81**, 2736–2745 (2019).
15. Weston, A. D. *et al.* Automated abdominal segmentation of CT scans for body composition analysis using deep learning. *Radiology* **290**, 669–679 (2018).
16. Wang, Z. *et al.* An effective CNN method for fully automated segmenting subcutaneous and visceral adipose tissue on CT scans. *Ann. Biomed. Eng.* **48**, 312–328 (2019).
17. Langner, T., Wikstrom, J., Bjerner, T., Ahlstrom, H. & Kullberg, J. Identifying morphological indicators of aging with neural networks on large-scale whole-body MRI. *IEEE Trans. Med. Imaging* <https://doi.org/10.1109/tmi.2019.2950092> (2019).
18. Langner, T. Deep regression for biometry on body MRI. <https://github.com/tarolangner/mri-biometry> (2019).

19. Hu, H. & Kan, H. E. Quantitative proton MR techniques for measuring fat. *NMR Biomed.* **26**, 1609–1629 (2013).
20. West, J. *et al.* Feasibility of MR-based body composition analysis in large scale population studies. *PLoS One* <https://doi.org/10.1371/journal.pone.0163332> (2016).
21. McInnes, L., Healy, J. & Melville, J. Umap: uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426* (2018).
22. Harvey, N., Matthews, P., Collins, R., Cooper, C. *et al.* Osteoporosis epidemiology in UK Biobank: a unique opportunity for international researchers (2013).
23. Borga, M. *et al.* Validation of a fast method for quantification of intra-abdominal and subcutaneous adipose tissue for large-scale human studies. *NMR Biomed.* **28**, 1747–1753 (2015).
24. Borga, M. *et al.* Advanced body composition assessment: from body mass index to body composition profiling. *J. Investig. Med.* **66**, 1–9 (2018).
25. Wilman, H. R. *et al.* Characterisation of liver fat in the UK Biobank cohort. *PLoS One* **12**, e0172921 (2017).
26. Simonyan, K. & Zisserman, A. Very deep convolutional networks for large-scale image recognition. [arXiv:1409.1556](https://arxiv.org/abs/1409.1556) [cs] (2014).
27. He, K., Zhang, X., Ren, S. & Sun, J. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778. <https://doi.org/10.1109/CVPR.2016.90> (2016).
28. Selvaraju, R. R. *et al.* Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *2017 IEEE International Conference on Computer Vision (ICCV)*, 618–626. <https://doi.org/10.1109/ICCV.2017.74> (IEEE, Venice, 2017).
29. Ozbulak, U. Pytorch cnn visualizations. <https://github.com/utkuozbulak/pytorch-cnn-visualizations> (2019).
30. Ekström, S., Malmberg, F., Ahlström, H., Kullberg, J. & Strand, R. Fast graph-cut based optimization for practical dense deformable registration of volume images. [arXiv:1810.08427](https://arxiv.org/abs/1810.08427) [cs] (2018).
31. Bamberg, F. *et al.* Whole-body MR imaging in the German national cohort: rationale, design, and technical background. *Radiology* **277**, 206–220 (2015).
32. Kukuk, G. M. *et al.* Comparison between modified dixon MRI techniques, MR spectroscopic relaxometry, and different histologic quantification methods in the assessment of hepatic steatosis. *Eur. Radiol.* **25**, 2869–2879 (2015).
33. Roldan-Valadez, E. *et al.* In vivo 3T spectroscopic quantification of liver fat content in nonalcoholic fatty liver disease: correlation with biochemical method and morphometry. *J. Hepatol.* **53**, 732–737 (2010).

Acknowledgements

This work was supported by a research grant from the Swedish Heart- Lung Foundation and the Swedish Research Council (2016-01040, 2019-04756) and used the UK Biobank resource under application no. 14237.

Author contributions

T.L. wrote the main manuscript text and conducted the experiments, R.S. revised the manuscript, H.A. and J.K. supervised data access from the UK Biobank and contributed to the statistical evaluation. All authors read and approved the final manuscript.

Funding

Open Access funding provided by Uppsala University.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41598-020-74633-5>.

Correspondence and requests for materials should be addressed to T.L.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020