

Visual Feature Attribution using Wasserstein GANs

Christian F. Baumgartner¹
Ender Konukoglu¹

Lisa M. Koch² Kerem Can Tezcan¹ Jia Xi Ang¹
for the Alzheimer's Disease Neuroimaging Initiative*

¹Computer Vision Lab, ETH Zurich

²Computer Vision and Geometry Group, ETH Zurich

Abstract

Attributing the pixels of an input image to a certain category is an important and well-studied problem in computer vision, with applications ranging from weakly supervised localisation to understanding hidden effects in the data. In recent years, approaches based on interpreting a previously trained neural network classifier have become the de facto state-of-the-art and are commonly used on medical as well as natural image datasets. In this paper, we discuss a limitation of these approaches which may lead to only a subset of the category specific features being detected. To address this problem we develop a novel feature attribution technique based on Wasserstein Generative Adversarial Networks (WGAN), which does not suffer from this limitation. We show that our proposed method performs substantially better than the state-of-the-art for visual attribution on a synthetic dataset and on real 3D neuroimaging data from patients with mild cognitive impairment (MCI) and Alzheimer's disease (AD). For AD patients the method produces compellingly realistic disease effect maps which are very close to the observed effects.

1. Introduction

In this paper we address the problem of visual attribution, which we define as detecting and visualising evidence of a particular category in an image. Pinpointing all evidence of a class is important for a variety of tasks such as weakly supervised localisation or segmentation of structures [41, 43, 65], and better understanding disease effects, and physiological or pathological processes in medical images [67, 18, 12, 19, 13, 27, 54, 30, 31, 63].

Currently, the most frequently used approach to address the visual attribution problem is training a neural network classifier to predict the categories of a set of images and then following one of two strategies: analysing the gradients of

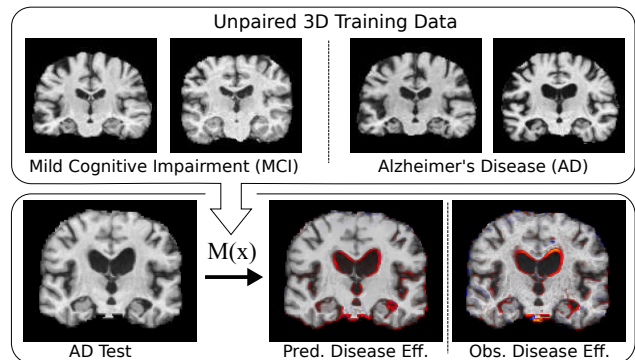


Figure 1. Our proposed method learns a map generating function $M(x)$ from unlabelled training data. Given a test image, this function will generate an image-specific visual attribution map which highlights the features unique to that category. The method is of particular interest for creating medical disease effect maps. We show that on neuroimaging data the method predicts effects in very good agreement with the actual observed effects.

the prediction with respect to an input image [27, 5, 54] or analysing the activations of the feature maps for the image [65, 41, 43] to determine which part of the image was responsible for making the associated prediction.

Visual attribution based directly on neural network classifiers may, under some circumstances, produce undesired results. It is known that such classifiers base their decisions on certain salient regions rather than the whole object of interest. It was recently shown that during training neural networks minimise the mutual information between input and output layers, thereby compressing the input features [50]. These findings suggest that a classifier may ignore features with low discriminative power if stronger features with redundant information about the target are available. In other words, neural network training may be working in opposition to the goal of visual attribution. As a consequence, if there is evidence for a class at multiple locations in the image (such as multiple lesions in medical images) some locations may not influence the classification result and may thus not be detected. We demonstrate this effect on a synthetic dataset in our experiments.

*Data used in preparation of this article were obtained from the Alzheimers Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu).

It would be highly desirable if instead we could visualise evidence of a particular category in a way that captures *all* category-specific effects in an image. Our main contribution is a novel approach towards solving the visual attribution problem which takes a first step in this direction. In contrast to the majority of recent techniques, the method does not rely on a classifier but rather aims at finding a map that, when added to an input image of one category, will make it indistinguishable from images from a baseline category. To this end we propose a generative model in which the additive map is learned as a function of the images. The method is based on Wasserstein generative adversarial networks (WGAN) [2], which have the desirable property that they minimise an approximation of the Wasserstein distance between the distributions of the generated images and the real ones.

We note that our method does not tackle the classification problem but rather assumes that the category labels of the test images have already been determined (e.g. using a separately trained classifier or by an expert). Furthermore, the method requires a baseline category, which is not the case for many benchmark recognition datasets in vision, but is in fact the case for many practical detection applications, especially in medical image analysis.

We demonstrate the method on synthetic 2D data and on large 3D brain MR data, where we aim to predict subject-specific disease effect maps for Alzheimer’s disease (AD).

1.1. Medical motivation

Identifying disease effects at the subject-specific level is of great interest for various medical applications. In clinically oriented research, identifying subject-specific disease effects would be useful for stratification amongst the patient population and to help disentangling diseases such as AD [25] and Schizophrenia [48], that are believed to be composed of multiple sub-types rather than a single disease. Furthermore, for clinicians, subject-specific maps could be helpful in assessing disease status and grading.

In this paper, we chose to study the disease effects of AD with respect to mild cognitive impairment (MCI), which is characterised by a slight decline in cognitive abilities. Patients with MCI are at increased risk of developing AD, but do not always do. We evaluate our method on one of the largest publicly available neuroimaging datasets acquired by the Alzheimer’s Disease Neuroimaging Initiative (ADNI). We used the MCI population as the baseline category and the AD population as the category of interest. Our choice to use MCI as our baseline is motivated by the fact that the ADNI dataset contains a number of MCI subjects who convert AD with imaging data at both stages of the disease. This allowed us to evaluate the predicted disease effects against real observed effects defined as the differences between images at the different stages. Note that

even though using normal controls as the baseline is feasible, it would have been much harder to assess the proposed method due to the small number of control to AD converters in the ADNI dataset.

2. Related work

2.1. Visual attribution

A commonly used approach for weakly supervised localisation or segmentation is to analyse the final feature map of a neural network classifier [41, 43]. The Class Activation Mapping (CAM) method [65] builds on those techniques by reducing the feature maps of the second to last layer using a global average pooling layer, followed by a dense prediction layer. This allows to create class-specific activation maps as a linear combination of the weights in the last layer.

A large amount of works on medical images builds on the CAM technique. Examples include the work of Feng et al. [12] on pulmonary nodule localisation in CT, the work of Ge et al. [18] on skin disease recognition. Other examples are [67], [19]. It is important to note that CAM is restricted in the resolution of its visual attributions by the resolution of the last feature map. Consequently, often post-processing of the predictions is required [12, 13, 43]. In contrast, our proposed method can produce visual attributions at the resolution of the original input images.

Another class of techniques creates saliency maps by backpropagating back to the input image. Examples include Guided Backprop [53], Excitation Backprop [62], Integrated Gradients [54], meaningful perturbations [13].

Similar techniques have been applied in the domain of medical images. Jamaludin et al. [27] use the backprop-based saliency technique proposed by [51] to pinpoint lumbar degradations, and Baumgartner et al. [5, 6] use a variant of [53] to localise fetal anatomy. Gao and Noble [15] apply a similar approach to localise the fetal heart.

2.2. Statistical disease models

Statistical analysis of medical images for identifying disease effects has been an instrumental tool for various diseases and disorders [56, 46, 9] as well as other non-disease related factors [17, 38, 28, 59, 42, 55]. The most common approach is to use regression analysis or machine learning tools to generate population average maps, which highlights features that are salient across the population [3, 32, 61, 16, 39, 45, 14].

Recently, constructing subject-specific maps has received attention. Maumet et al. took a one-versus-all group analysis approach [37, 36], while Konukoglu and Glocker extracted subject-specific maps with predictive models and Markov Random Field restoration [30, 31].

The common drawback in the previous approaches is the need for *registration*. In order to compute disease effect

maps, images of different subjects need to be non-rigidly aligned on a common template where statistical analysis can be performed. The non-rigid registration process brings additional uncertainty to the subject-specific maps. Our work, addresses this shortcoming and generates subject-specific disease effect maps without requiring registration.

2.3. Image generation using GANs

Generative adversarial images conditioned on an input image have been used in diverse applications such as video frame prediction [35], image super-resolution [33], image-translation across domains using paired [26] and unpaired [66] images, and pixel level domain adaptation [7, 49].

In the context of medical images, GANs have been applied to super-resolution in retinal fundus images [34], for semi-supervised cardiac segmentation [64], synthesising computed tomography images from MR images [40, 60] and intraoperative motion modelling [23]. Although some of the above models use 3D data, the examined volumes are usually relatively small [23], or the networks operate in a patch-wise fashion [40]. It is important to note that in the case of brain MR images of Alzheimers disease patients, the diagnostic information is only visible at a high resolution and cannot be determined by considering small local patches only. In this paper, we therefore tackle the challenge of processing large 3D volumes directly.

2.4. Contributions

1. We demonstrate a limitation in current neural network based visual attribution methods using synthetic data.
2. We propose a novel visual attribution technique that can detect class specific regions more completely and at a high resolution.
3. To our knowledge, this is the first application of generative adversarial networks on large structural 3D data.

An implementation of the proposed method is publicly available here: <https://github.com/baumgach/vagan-code>.

3. Visual attribution using WGANs

3.1. Problem Formulation

Our goal is to estimate a map that highlights the areas in an image which are specific to the class the image belongs to. We formulate the problem for two classes $c \in \{0, 1\}$, a baseline class and a class of interest. The formulation however, easily extends to the case of multiple classes of interest. We denote an image with x and the distribution of images coming from class $c = 0$ with $p_d(x|c = 0)$ and images from class $c = 1$ with $p_d(x|c = 1)$. In the case of medical application, $c = 1$ could for example denote the

set of images from a population with a certain disease and $c = 0$ images of control subjects.

We formulate a problem as estimating a map function $M(x)$ that, when added to an image x_i from category $c = 1$, creates an image

$$y_i = x_i + M(x_i), \quad (1)$$

which is indistinguishable from the images sampled from $p_d(x|c = 0)$. Thereby, the map $M(x_i)$ contains all the features which distinguish the input image x_i from the other category. In the case of medical images, M will by definition contain the effects of a disease visible in the images, i.e. a *disease effect map*.

We model the function M using a convolutional neural network, whose parameters we find using a WGAN.

3.2. Wasserstein GANs

In the GAN paradigm a generator function and a discriminator function (both neural networks) compete with each other in a zero-sum game [20]. Given random noise as input, the generator tries to produce realistic images that fool the discriminator, while the discriminator tries to learn the difference between generated and real images.

Arjovski and Bottou pointed out a limitation in this paradigm which precludes a guarantee that the generated images will necessarily converge to the target distribution [1] (although in practice, with appropriate training methods, many impressive results were achieved [44]). Wasserstein GANs are a modification to the classic GAN paradigm where the discriminator is replaced by a critic which does not have an activation function in its final layer and which is constrained to be a K -Lipschitz function. WGANs have better optimisation properties and it can be shown that they minimise a meaningful distance between the generated and real distributions.

3.3. Constrained effect maps using WGANs

In this work we build on WGANs to find the optimal map generation function. In contrast to regular WGANs, we have a map generator function $M(x_i)$, which, during training, takes as input randomly sampled images x_i from category $c = 1$ rather than noise. M tries to generate maps that, when added to x_i , create images y_i appearing to be from category $c = 0$. By trying to distinguish generated images y_i from real images from category $c = 0$, the critic D ensures that the generated maps are constrained to realistic modifications (see Fig. 2 for an overview). In the context of medical images, this means enforcing anatomically realistic modifications to the images.

Building on [2] this leads to the following cost function:

$$\begin{aligned} \mathcal{L}_{GAN}(M, D) = & \mathbb{E}_{x \sim p_d(x|c=0)}[D(x)] \\ & - \mathbb{E}_{x \sim p_d(x|c=1)}[D(x + M(x))]. \end{aligned} \quad (2)$$




Figure 2. Overview of VA-GAN. During training images are sampled from the categories $c \in \{0, 1\}$. Images from $c = 1$ are passed to the map generating function $M(x)$. The map generator aims to create additive maps which produce generated images that the critic $D(x)$ cannot distinguish from images sampled from $p_d(x|c = 0)$. The critic, $D(x)$ tries to assign different values to generated and real images. During testing, $M(x)$ can be used directly to predict a map in a single forward pass.

Optimising Eq. 2 directly could lead to changes in the input image x_i that change the image identity. For instance, the brain anatomy of a subject could be changed to a degree where it does not only capture disease related changes but changes the subject identity. We want to encourage the smallest required map M that still leads to a realistic y_i . Thus add the following data regularisation term to the cost function:

$$\mathcal{L}_{reg}(M) = \|M(x)\|_1, \quad (3)$$

where $\|\cdot\|_1$ is the L1 norm [13].

The final optimisation is then given by

$$M^* = \operatorname{argmin}_M \max_{D \in \mathcal{D}} \mathcal{L}_{GAN}(M, D) + \lambda \mathcal{L}_{reg}(M), \quad (4)$$

where \mathcal{D} is the set of 1-Lipschitz functions.

In order to enforce the Lipschitz constraint we use the optimisation method proposed in [22]. As recommended by [22], we weigh the gradient penalty with a factor of 10 throughout all experiments.

3.4. Network architecture

As we will discuss in more detail in Section 4.3, we design our proposed method with large 3D medical imaging data in mind, which often need to be processed at high resolutions in order to retain diagnostic information. Specifically, in our experiments on neuroimaging data, an input volume size of 128x160x112 voxels is used.

With such large images the limiting factor becomes storing the activations of the networks on GPU memory. With this in mind we design the map generator and the critic networks as follows.

3.4.1 Map generator network

The map generator function should be able to form an internal representation of the visual attributes that characterise the categories. In the case of brain images affected by dementia, it should be able to “understand” the systematic changes involved in the disease. Therefore, a relatively

powerful network is required to adequately model the function M . To this end, we use the 3D U-Net [10] (originally proposed for segmentation), as a starting point. The 3D U-Net has an encoder-decoder structure with a bottle-neck layer in the middle, but additionally introduces skip connections at each resolution level bypassing the bottle-neck. This allows the network to combine high-level semantic information (such as the presence of a structure) with low-level information (such as edges).

In order to reduce GPU memory consumption we reduce the number of feature maps by a factor of 4 in most layers. As in the original 3D U-Net [10] we use batch normalisation for all layers except the final one. The exact architecture is shown in Fig. 2 in the supplementary material.

3.4.2 Critic function

In line with related literature on image generation using GANs [26, 66, 49], we model our critic as a fully convolution network with no dense layers. We loosely base our architecture on the C3D network which achieved impressive results on action recognition tasks in video data by processing them directly in the spatio-temporal 3D space [57]. However, in contrast to that work we only perform 4 pooling steps. After the fourth pooling layer we add another 3x3x3 convolution layer, followed by a 1x1x1 convolution layer which reduces the number of feature maps to one. The final critic prediction is given by a global average pooling operation of that feature map.

It proved important *not* to use batch normalisation for the critic network. Towards the beginning of training generating statistics of a batch with generated and the real images may not produce reasonable estimates, because the images vary considerably from each other. We surmise that this effect prevents the critic from learning when batch normalisation is used. A similar observation was made in [22]. We also experimented with layer normalisation [4], but did not observe improvements.

The exact architecture we used is shown in Fig. 1 in the supplementary material.

3.5. Training

To optimise our networks, we follow [2, 22] and update the parameters of the critic and map generator networks in an alternating fashion. In contrast to the regular GANs [20], WGANs require a critic which is kept close to optimality through-out training. We therefore perform 5 critic updates for every map generator update. Additionally, for the first 25 iterations and every hundredth iteration, we perform 100 critic updates per generator update.

With the above architectures, the maximum batch size that can be used for a single gradient computation on a Nvidia Titan Xp GPU with 12 GB of memory is 2+2 (real+generated). In order to obtain more reliable gradient estimates we aggregate the gradients for a total of 6 mini-batches before performing a training step.

We used the ADAM optimiser [29] to perform the update steps for all experiments. The optimiser parameters were set to $\beta_1 = 0$, $\beta_2 = 0.9$, and we used a learning rate of 10^{-3} . Lastly, we used a weight of $\lambda = 10^2$ for the map regularisation term (see Eq. 4) throughout the paper. Training took approximately 24 hours on an Nvidia Titan Xp.

4. Experiments

We evaluated the proposed method using a synthetically generated dataset and a large number of 3D brain MRI images from the publicly available ADNI dataset.

We compared our proposed visual attribution GAN (VA-GAN) to methods from the literature which have been used for visual attribution both on natural and on medical images. Specifically, we compared against Guided Backpropagation [53], Integrated Gradients [54] and Class Activation Mapping (CAM) [65]. Furthermore, to verify that the WGAN framework is necessary, we also investigated an alternative way of estimating the additive map not based on GANs, which is described in detail in the next section.

All the methods except VA-GAN use classification networks. For simplicity, we used a very similar architecture for these networks as for the critic in VA-GAN, except for two differences: (1) we replaced the last convolution and the global average pooling layer by two dense layers followed by a softmax and (2) we used batch normalisation for all layers, which produced better classification results for the experiments on the ADNI dataset. In addition, for the CAM method we designed the last layer as described by [65] and omitted the last two max pooling layers, which allowed significantly more accurate visual attribution maps due to the higher resolution of the last feature maps.

Lastly, for the experiments on the 2D synthetic data we simply replaced all 3D operations by 2D operations, but left the architectures otherwise unchanged.

4.1. Classifier-based map estimation

In the VA-GAN approach, we generate an additive map which is constrained by the critic to generate a realistic image from the opposite class. To demonstrate that this approach is necessary we also investigated an alternative method of estimating the additive map without a term enforcing realistic maps.

The alternative approach requires training a classifier $p(c = 1) = f(\cdot)$ and then optimising an additive map *image* m that lowers the prediction $p(c = 1)$ as much as possible. That is to say, the image $y_i = x_i + m$ should minimise $f_i(y_i)$. This formulation is almost exactly the same as for the WGAN-based approach (see Eq. 1) except that m is not a function of x_i .

We need to use a regularisation in determining m to avoid trivial solutions, such as imperceptible changes that can fool classifiers [21]. A “well behaved” map can be found by the following minimisation problem:

$$m^* = \underset{m}{\operatorname{argmin}} f(x_i + m) + \omega_1 \|m\|_1 + \omega_2 \sum_u \|\nabla m(u)\|_\beta^\beta. \quad (5)$$

Here u indexes the pixels or voxels of m . The L1 term weighted by ω_1 encourages small maps, while the total variation term weighted by ω_2 encourages smoothness.

We optimise this cost function using the ADAM optimiser using the default internal parameters given in [29] with a learning rate of 10^{-2} and early stopping at 1500 iterations. Furthermore, we set $\beta = 2.0$, $\omega_1 = 10^{-2}$ and $\omega_2 = 10^{-5}$ in all experiments.

This approach is strongly related to the meaningful perturbation masks technique proposed by [13] in which parts of an image are locally deleted by a mask m such that the prediction $f(x)$ is minimised. In preliminary experiments we found that on the medical image problem we studied, visual attribution using destructive masks did not lead to the desired results. Deleting the diagnostic part of an image will not produce an image of the opposite class but rather an image with an undetermined diagnosis. This means such a mask may contain information about the location of diagnostic regions but not about specific disease effects, e.g. enlargement or shrinkage. In contrast, by optimising Eq. 5 we attempt to morph the image into the opposite class, such that diagnostic regions can be changed to have the characteristics of another class. Because of the similarity to [13], we refer to this method as additive perturbation maps.

4.2. Synthetic experiments

Data: In order to quantitatively evaluate the performance of the examined visual attribution methods, we generated a synthetic dataset of 10000 112x112 images with two classes, which model a healthy control group (label 0) and a




Figure 3. Description of synthetic data. We generated noisy observations from ground-truth effect maps. The dataset contained two categories: A baseline category 0 (e.g. healthy images) and category with an effect (e.g. patient images). The images in category 1 contained one of two subtypes, A or B, which is unknown to the algorithms. A: box in the lower right, B: box in the upper left.

patient group (label 1). The images were split evenly across the two categories. We closely followed the synthetic data generation process described in [31] where disease effects were studied in smaller cohorts of registered images.

The control group (label 0) contained images with random iid Gaussian noise convolved with a Gaussian blurring filter. Examples are shown in Fig. 3. The patient images (label 1) also contained the noise, but additionally exhibited one of two disease effects which was generated from a ground-truth effect map: a square in the centre and a square in the lower right (subtype A), or a square in the centre and a square in the upper left (subtype B). Importantly, both disease subtypes shared the same label. The location of the off-centre squares was randomly offset in each direction by a maximum of 5 pixels. This effect was added to make the problem harder, but had no notable effect on the outcome.

Evaluation: We split the data into a 80-20 training and testing set. Moreover, we used 20% of the training set for monitoring the training. Next, we estimated the disease effect maps for all cases from the synthetic patient class using the examined methods.

In order to assess the visual attribution accuracy quantitatively, we calculated the normalised cross correlation (NCC) between the ground-truth label maps and the predicted disease effect maps. The NCC has the advantage that it is not sensitive to the magnitude of the signals. For CAM we used only the positive values to calculate the NCC, while for the backprop-based techniques we used the absolute value, since those techniques do not necessarily predict the correct sign of the changes.

Results: A number of examples of the estimated disease effect maps are shown in Fig. 4. Guided Backpropagation produced similar results to Integrated Gradients. We therefore omitted it in the visual results due to space considera-




Figure 4. Examples of visual attribution on synthetic data obtained using the compared methods.

tions but provide quantitative results.

For the backprop-based methods we consistently observed two behaviours: 1) They tended to focus exclusively on the central square which was always present and was thus the most predictive set of features. This behaviour is consistent with the network compressing away less predictive features discussed earlier [50]. 2) They tended to focus mostly on the edges of the boxes rather than on the whole object. This may have to do with the fact that edges are more salient than other points and, again, are sufficient to predict the presence or absence of the box.

The CAM method managed to capture both squares most of the times, but by design had limited spatial resolution. Note that due to the lower number of max-pooling layers used for the CAM classifier each pixel in the last feature map had a receptive field of only 39x39 pixels. This could mean that many pixels in that feature map could not simultaneously see both of the squares, which may have contributed to the squares being better discerned. However, we did not investigate this further.

Lastly, our proposed VA-GAN method produced the most localised disease effect maps, finding the entire boxes and following the edges closely. It also managed to consistently identify both disease effects.

Table 1. NCC scores for experiments on synthetic data.

Method	mean	std.
Guided Backprop [53]	0.14	0.04
Integrated Gradients [54]	0.36	0.11
CAM [65]	0.48	0.04
Additive Perturbation	0.06	0.03
VA-GAN	0.94	0.07

The quantitative NCC results shown in Table 1 are mostly consistent with our qualitative observations, with VA-GAN obtaining significantly higher NCC than the other methods. The additive perturbation technique achieved a

low score due to its exclusive focus on edges.

4.3. Experiments on real neuroimaging data

In this section, we investigate the methods’ ability to detect the areas of the brain which are involved in the progression from MCI to AD at a subject-specific level. We trained on images from both categories and then generated disease effect maps only for the AD images.

Data: We selected 5778 3D T1-weighted MR images from 1288 subjects with either an MCI (label 0) or AD (label 1) diagnosis from the ADNI cohort. 2839 of the images were acquired using a 1.5T magnet, the remainder using a 3T magnet. The subjects are scanned at regular intervals as part of the ADNI study and a number of subjects converted from MCI to AD over the years. We did not use these correspondences for training, however, we took advantage of it for evaluation as will be described later. An overview of the data is given in the supplemental materials in Section C.

All images were processed using standard operations available in the FSL toolbox [52] in order to reorient and rigidly register the images to MNI space, crop them and correct for field inhomogeneities. We then skull-stripped the images using the ROBEX algorithm [24]. Lastly, we resampled all images to a resolution of 1.3 mm^3 and normalised them to a range from -1 to 1. The final volumes had a size of $128 \times 160 \times 112$ voxels.

Evaluation: We split the data on a subject level into a training, testing and validation set containing 825, 256 and 207 subjects, respectively. We then trained all of the algorithms with both AD and MCI data as described earlier, and generated disease effect maps for the AD subjects from the test set. The validation set was used to monitor the training.

In order to better understand the quality of the generated disease maps we estimated the actual deformations for a number of subjects as follows. We identified all subjects from the test set who were diagnosed with MCI during the baseline examination but progressed to AD in one of the follow-up scans. We then aligned those images rigidly and subtracted them from each other to obtain an *observed disease effect map*. We excluded all subjects which were not acquired with the same field strength, since a large amount of the observed effects could be due to differences in image quality. This left 50 subjects which we evaluated more closely. We note that even for the same field strength there are a number of artefacts due to intensity variations and registration. Furthermore, there are likely to be effects not caused by the disease, such as ageing (which will also be captured by our method), such that the observed disease effect maps could be considered a ground-truth.

Nevertheless, we also evaluated NCC between the observed and the predicted disease effect maps in the same manner as for the synthetic data.




Figure 6. Close-up of the hippocampus region of a subject before (left) and after developing AD (middle). The right panel shows the generated image. The red (hippocampus) and green (ventricles) contours are in the same location in all three images. It can be observed that the map “reverses” some of the atrophy.

Table 2. NCC scores for experiments on neuroimaging data.

Method	mean	std.
Guided Backprop [53]	0.05	0.03
CAM [65]	0.09	0.07
Integrated Gradients [54]	0.13	0.05
Additive Perturbation	0.11	0.05
VA-GAN	0.27	0.15

Results: Fig. 5 shows disease effect maps obtained for a selection of AD subjects (we again omitted Guided Backprop in the figure). The subjects are ordered by increasing progression of the disease as measured by the ADAS13 cognition exam [47]. It can be seen that VA-GAN’s predictions were in very good agreement with the observed effect maps. As is known from the literature [8, 11] the method indicates atrophy in the hippocampi, and general brain atrophy around the ventricles. Furthermore, it is known that in later stages of the disease other brain areas such as the temporal lobe get affected as well [58]. Those effects were also identified by VA-GAN in the last subject in Fig. 5.

The backprop-based methods and additive perturbations were observed to be very noisy and tended to identify *only* the hippocampal areas. We believe that this is in agreement with the findings on the synthetic data. The hippocampus is known to be the most predictive region for AD, however, it is also known that many other regions are involved in the disease. It is likely, that classifiers learned to focus only on the most discriminative set of features ignoring the rest. Lastly, it is hard to interpret the results produced by CAM due to the low resolution. However, the images suggest that this method focuses on similar areas as the other methods.

Quantitative results are given in Table 2. VA-GAN obtained the highest correlation scores, however, it is hard to draw conclusions from these figures due to the noisy nature of the observed effect maps as well as the possible non-disease related effects on the observed effect maps, which are taken to be “ground-truth” in the experiments.

We observed that VA-GAN generally produced very realistic deformations. In Fig. 6 a close-up of the MCI, AD, and generated image is shown for a sample subject. It can




Figure 5. Coronal and sagittal views of generated AD effect maps for three subjects and actual observed effects. Maps are shown as coloured overlay over the input image. The ventricular (arrow A) and hippocampal (arrow B) regions are particularly affected by the disease and are reliably captured by VA-GAN. In later stages also other brain regions such as the temporal lobe (arrow C) are affected. We also report the ADAS13 cognitive exam scores (larger means AD is further progressed) and the ADNI identifier (rid) for each subject.

be seen that our method succeeded in making the generated image more similar to the corresponding MCI image and that the changes were realistic.

5. Limitations and discussion

We have proposed a method for visual feature attribution using Wasserstein GANs. It was shown that, in contrast to backprop-based methods, our technique can capture multiple regions affected by disease, and produces state-of-the-art results for the prediction of disease effect maps in neuroimaging data and on a synthetic dataset.

Currently, the method assumes that the category labels of the test data are known during test-time. In case they are

unknown, the method could be easily combined with classifier which produces this information. We only evaluated the method for the case of two labels. More categories could be addressed by training multiple map generators each mapping to a background class (assuming there is one).

In the future, we plan to model other effects such as ageing or the presence or absence of certain genes on the ADNI data, investigate the method on other datasets and apply it to other problems such as weakly-supervised localisation.

Acknowledgements

We gratefully acknowledge the support of NVIDIA Corporation with the donation of a Titan Xp GPU.

References

- [1] M. Arjovsky and L. Bottou. Towards principled methods for training generative adversarial networks. *arXiv preprint arXiv:1701.04862*, 2017.
- [2] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein GAN. *arXiv preprint arXiv:1701.07875*, 2017.
- [3] J. Ashburner and K. J. Friston. Why voxel-based morphometry should be used. *Neuroimage*, 14(6):1238–1243, 2001.
- [4] J. L. Ba, J. R. Kiros, and G. E. Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [5] C. F. Baumgartner, K. Kamnitsas, J. Matthew, T. P. Fletcher, S. Smith, L. M. Koch, B. Kainz, and D. Rueckert. Sononet: Real-time detection and localisation of fetal standard scan planes in freehand ultrasound. *IEEE Transactions on Medical Imaging*, 2017.
- [6] C. F. Baumgartner, K. Kamnitsas, J. Matthew, S. Smith, B. Kainz, and D. Rueckert. Real-time standard scan plane detection and localisation in fetal ultrasound using fully convolutional neural networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 203–211. Springer, 2016.
- [7] K. Bousmalis, N. Silberman, D. Dohan, D. Erhan, and D. Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. *arXiv preprint arXiv:1612.05424*, 2016.
- [8] H. Braak and E. Braak. Neuropathological staging of Alzheimer-related changes. *Acta neuropathologica*, 82(4):239–259, 1991.
- [9] E. J. Burton et al. Cerebral atrophy in Parkinsons disease with and without dementia: a comparison with Alzheimers disease, dementia with Lewy bodies and controls. *Brain*, 127(4):791–800, 2004.
- [10] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger. 3D U-Net: learning dense volumetric segmentation from sparse annotation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 424–432. Springer, 2016.
- [11] B. C. Dickerson et al. The cortical signature of alzheimer’s disease: regionally specific cortical thinning relates to symptom severity in very mild to mild AD dementia and is detectable in asymptomatic amyloid-positive individuals. *Cerebral cortex*, 19(3):497–510, 2009.
- [12] X. Feng, J. Yang, A. F. Laine, and E. D. Angelini. Discriminative localization in CNNs for weakly-supervised segmentation of pulmonary nodules. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 568–576. Springer, 2017.
- [13] R. Fong and A. Vedaldi. Interpretable explanations of black boxes by meaningful perturbation. *arXiv preprint arXiv:1704.03296*, 2017.
- [14] M. Ganz et al. Relevant feature set estimation with a knock-out strategy and random forests. *NeuroImage*, 122:131–148, 2015.
- [15] Y. Gao and J. A. Noble. Detection and characterization of the fetal heartbeat in free-hand ultrasound sweeps with weakly-supervised two-streams convolutional networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 305–313. Springer, 2017.
- [16] B. Gaonkar and C. Davatzikos. Analytic estimation of statistical significance maps for support vector machine based multi-variate image analysis and classification. *NeuroImage*, 78:270–283, 2013.
- [17] E. Garrido, A. Castello, J. Ventura, A. Capdevila, and F. Rodriguez. Cortical atrophy and other brain magnetic resonance imaging (MRI) changes after extremely high-altitude climbs without oxygen. *International journal of sports medicine*, 14(04):232–234, 1993.
- [18] Z. Ge, S. Demyanov, R. Chakravorty, A. Bowling, and R. Garnavi. Skin disease recognition using deep saliency features and multimodal learning of dermoscopy and clinical images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 250–258. Springer, 2017.
- [19] W. M. Gondal, J. M. Köhler, R. Grzeszick, G. A. Fink, and M. Hirsch. Weakly-supervised localization of diabetic retinopathy lesions in retinal fundus images. *arXiv preprint arXiv:1706.09634*, 2017.
- [20] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [21] I. J. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [22] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville. Improved training of wasserstein GANs. *arXiv preprint arXiv:1704.00028*, 2017.
- [23] Y. Hu, E. Gibson, T. Vercauteren, H. U. Ahmed, M. Emberton, C. M. Moore, J. A. Noble, and D. C. Barratt. Intraoperative organ motion models with an ensemble of conditional generative adversarial networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 368–376. Springer, 2017.
- [24] J. E. Iglesias, C.-Y. Liu, P. M. Thompson, and Z. Tu. Robust brain extraction across datasets and comparison with publicly available methods. *IEEE transactions on medical imaging*, 30(9):1617–1634, 2011.
- [25] K. o. Iqbal. Subgroups of Alzheimer’s disease based on cerebrospinal fluid molecular markers. *Annals of neurology*, 58(5):748–757, 2005.
- [26] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. *arXiv preprint arXiv:1611.07004*, 2016.
- [27] A. Jamaludin, T. Kadir, and A. Zisserman. Spinenet: automatically pinpointing classification evidence in spinal MRIs. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 166–175. Springer, 2016.
- [28] R. Kanai and G. Rees. The structural basis of inter-individual differences in human behaviour and cognition. *Nature Reviews Neuroscience*, 12(4):231–242, 2011.
- [29] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

- [30] E. Konukoglu and B. Glocker. Constructing subject-and disease-specific effect maps: Application to neurodegenerative diseases. In *Medical Computer Vision and Bayesian and Graphical Models for Biomedical Imaging*, pages 3–13. Springer, 2016.
- [31] E. Konukoglu and B. Glocker. SubCMap: subject and condition specific effect maps. *arXiv preprint arXiv:1701.02610*, 2017.
- [32] A. Krishnan et al. Partial least squares (PLS) methods for neuroimaging: a tutorial and review. *Neuroimage*, 56(2):455–475, 2011.
- [33] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. *arXiv preprint arXiv:1609.04802*, 2016.
- [34] D. Mahapatra, B. Bozorgtabar, S. Hewavitharanage, and R. Garnavi. Image super resolution using generative adversarial networks and local saliency maps for retinal image analysis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 382–390. Springer, 2017.
- [35] M. Mathieu, C. Couprie, and Y. LeCun. Deep multi-scale video prediction beyond mean square error. *arXiv preprint arXiv:1511.05440*, 2015.
- [36] C. Maumet, P. Maurel, J.-C. Ferré, and C. Barillot. An a contrario approach for the detection of patient-specific brain perfusion abnormalities with arterial spin labelling. *Neuroimage*, 134:424–433, July 2016.
- [37] C. Maumet, P. Maurel, J.-C. Ferré, B. Carsin, and C. Barillot. Patient-specific detection of perfusion abnormalities combining within-subject and between-subject variances in Arterial Spin Labeling. *Neuroimage*, 81:121–130, Nov. 2013.
- [38] D. Miller and J. OCallaghan. Effects of aging and stress on hippocampal structure and function. *Metabolism*, 52:17–21, 2003.
- [39] B. Mwangi, T. S. Tian, and J. C. Soares. A review of feature reduction techniques in neuroimaging. *Neuroinformatics*, 12(2):229–244, 2014.
- [40] D. Nie, R. Trullo, J. Lian, C. Petitjean, S. Ruan, Q. Wang, and D. Shen. Medical image synthesis with context-aware generative adversarial networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 417–425. Springer, 2017.
- [41] M. Oquab, L. Bottou, I. Laptev, and J. Sivic. Is object localization for free?-weakly-supervised learning with convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 685–694, 2015.
- [42] J. S. Peper, R. M. Brouwer, D. I. Boomsma, R. S. Kahn, H. Pol, and E. Hilleke. Genetic influences on human brain structure: a review of brain imaging studies in twins. *Human brain mapping*, 28(6):464–473, 2007.
- [43] P. O. Pinheiro and R. Collobert. From image-level to pixel-level labeling with convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1713–1721, 2015.
- [44] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- [45] M. Rahim et al. Integrating multimodal priors in predictive models for the functional characterization of Alzheimers disease. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015*, pages 207–214. Springer, 2015.
- [46] H. Rosas et al. Regional and progressive thinning of the cortical ribbon in Huntingtons disease. *Neurology*, 58(5):695–701, 2002.
- [47] W. G. Rosen, R. C. Mohs, and K. L. Davis. A new rating scale for Alzheimer’s disease. *The American journal of psychiatry*, 1984.
- [48] C. A. Ross, R. L. Margolis, S. A. Reading, M. Pletnikov, and J. T. Coyle. Neurobiology of schizophrenia. *Neuron*, 52(1):139–153, 2006.
- [49] A. Shrivastava, T. Pfister, O. Tuzel, J. Susskind, W. Wang, and R. Webb. Learning from simulated and unsupervised images through adversarial training. *arXiv preprint arXiv:1612.07828*, 2016.
- [50] R. Shwartz-Ziv and N. Tishby. Opening the black box of deep neural networks via information. *arXiv preprint arXiv:1703.00810*, 2017.
- [51] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.
- [52] S. M. Smith, M. Jenkinson, M. W. Woolrich, C. F. Beckmann, T. E. Behrens, H. Johansen-Berg, P. R. Bannister, M. De Luca, I. Drobnjak, D. E. Flitney, et al. Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage*, 23:S208–S219, 2004.
- [53] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller. Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*, 2014.
- [54] M. Sundararajan, A. Taly, and Q. Yan. Axiomatic attribution for deep networks. *arXiv preprint arXiv:1703.01365*, 2017.
- [55] P. M. Thompson, T. D. Cannon, K. L. Narr, T. Van Erp, V.-P. Poutanen, M. Huttunen, J. Lönqvist, C.-G. Standertskjöld-Nordenstam, J. Kaprio, M. Khaledy, et al. Genetic influences on brain structure. *Nature neuroscience*, 4(12):1253–1258, 2001.
- [56] P. M. Thompson et al. Cortical change in Alzheimer’s disease detected with a disease-specific population-based brain atlas. *Cerebral Cortex*, 11(1):1–16, 2001.
- [57] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3D convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015.
- [58] V. Villemagne, M. Fodero-Tavoletti, C. Masters, and C. Rowe. Tau imaging: early progress and future directions. *The Lancet Neurology*, 14(1):114–124, 2015.
- [59] K. E. Watkins, F. Vargha-Khadem, J. Ashburner, R. E. Passingham, A. Connelly, K. J. Friston, R. S. Frackowiak, M. Mishkin, and D. G. Gadian. MRI analysis of an inherited speech and language disorder: structural brain abnormalities. *Brain*, 125(3):465–478, 2002.

- [60] J. M. Wolterink, A. M. Dinkla, M. H. Savenije, P. R. Seevinck, C. A. van den Berg, and I. Išgum. Deep MR to CT synthesis using unpaired data. In *International Workshop on Simulation and Synthesis in Medical Imaging*, pages 14–23. Springer, 2017.
- [61] K. J. Worsley et al. Characterizing the response of PET and fMRI data using multivariate linear models. *NeuroImage*, 6(4):305–319, 1997.
- [62] J. Zhang, Z. Lin, J. Brandt, X. Shen, and S. Sclaroff. Top-down neural attention by excitation backprop. In *European Conference on Computer Vision*, pages 543–559. Springer, 2016.
- [63] Q. Zhang, A. Bhalerao, and C. Hutchinson. Weakly-supervised evidence pinpointing and description. In *International Conference on Information Processing in Medical Imaging*, pages 210–222. Springer, 2017.
- [64] Y. Zhang, L. Yang, J. Chen, M. Fredericksen, D. P. Hughes, and D. Z. Chen. Deep adversarial networks for biomedical image segmentation utilizing unannotated images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 408–416. Springer, 2017.
- [65] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2921–2929, 2016.
- [66] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. *arXiv preprint arXiv:1703.10593*, 2017.
- [67] W. Zhu, Q. Lou, Y. S. Vang, and X. Xie. Deep multi-instance networks with sparse label assignment for whole mammogram classification. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 603–611. Springer, 2017.