The Unobserved Anatomy: Negotiating the Plausibility of AI-Based Reconstructions of Missing Brain Structures in Clinical MRI Scans

Paula Muhr

Abstracts (EN, DE) *Vast archives of fragmentary structural brain scans that are routinely acquired in medical clinics for diagnostic purposes have so far been considered to be unusable for neuroscientific research. Yet, recent studies have proposed that by deploying machine learning algorithms to fill in the missing anatomy, clinical scans could, in future, be used by researchers to gain new insights into various brain disorders. This chapter focuses on a study published in 2019, whose authors developed a novel unsupervised machine learning algorithm for synthesising missing anatomy in extremely sparse clinical MRI scans of thousands of stroke patients. By approaching the study from a media-theoretical perspective, I analyse how its authors discursively negotiated the anatomical and operative plausibility of the unobserved anatomy that their black-boxed algorithm reconstructed from the existing sparse data. My analysis foregrounds the processual, relational, context-dependent, and essentially unstable character of the thus established plausibility of the algorithmically synthesised neuroanatomy.*

*Immense Archive fragmentarischer struktureller Gehirnscans, die routinemäßig in medizinischen Kliniken zu Diagnosezwecken erstellt werden, galten bislang als unbrauchbar für die neurowissenschaftliche Forschung. Jüngste Studien haben jedoch vorgeschlagen, dass klinische Scans durch den Einsatz von Machine-Learning-Algorithmen zum Vervollständigen der fehlenden Anatomie in Zukunft genutzt werden könnten, um neue Erkenntnisse über verschiedene Hirnstörungen zu gewinnen. Im Mittelpunkt dieses Kapitels steht eine 2019 veröffentlichte Studie, deren Autor*innen einen neuartigen unüberwachten Algorithmus entwickelten, um fehlende anatomische Informationen in stark unvollständigen klinischen MRT-Scans von Tausenden von Schlaganfallpatient*innen zu synthetisieren. Ich analysiere aus einer medientheoretischen Perspektive, wie die Autor*innen diskursiv die anatomische und operative Plausibilität jener anatomischen Informationen erläuterten, die ihr Black-Box-Algorithmus aus den vorhandenen spärlichen Daten rekonstruierte. Meine Analyse stellt den prozessualen, relationalen, kontextabhängigen und im Wesentlichen instabilen Charakter der so hergestellten Plausibilität der algorithmisch synthetisierten, neuroanatomischen Bilder in den Vordergrund.*

Introduction

Initially introduced in the early 1980s, magnetic resonance imaging (MRI) has been hailed as a breakthrough technology that has revolutionised neuroscience.[1] MRI is a non-invasive imaging technology that facilitates the study of neuroanatomy in living human subjects by generating high-resolution greyscale images of brain structures. Strictly speaking, far from simply "'taking a picture of the brain,' MRI instead measures radio-frequency signals emitted from hydrogen atoms after the application of electromagnetic (radio-frequency) waves, localizing the signal using spatially varying magnetic gradients."[2] In other words, by placing a human subject inside an MRI scanner and exposing their brain to various specifically tailored combinations of static and dynamic magnetic fields, diverse types of structural MRI images can be obtained, each of which provides information about different neuroanatomical features of the same brain. Different MRI imaging modalities include T1-weighted images, T2-weighted images, proton-density images, MR angiography, FLAIR, white-matter nulled images, and diffusion-weighted images such as DTI.[3]

Since they visualise invisible internal structures otherwise inaccessible to human vision, all MRI image modalities are non-mimetic—they do not visually resemble the neuroanatomical properties they display. Despite their non-mimetic character, MRI images have been used for decades in medicine and neuroscience for making expert judgments about human brain structure "in a reliable manner."[4] As discussed elsewhere, this operative capacity of MRI images to deliver reliable information about brain anatomy is grounded in the images' referential relation to actual physical brains.[5] The referential relation, in turn, is a direct consequence of the stan-

1 See, e.g., Jason P. Lerch/André J. Van der Kouwe/Armin Raznahan et al., Studying Neuroanatomy Using MRI, in: *Nature Neuroscience* 20 (3/2017), 314–326, see 314. The conceptual and technological developments of MRI took place in the 1970s. See Scott A. Huettel/Allen W. Song/Gregory McCarthy, *Functional Magnetic Resonance Imaging*, Sunderland 2009, 18–24. But the first commercial MRI scanner for the human body that found immediate applications in hospitals was produced in 1982. See Huettel/Song/McCarthy, *Functional Magnetic Resonance Imaging*, 23.

2 Lerch et al., Studying Neuroanatomy Using MRI, 314.

3 See Huettel/Song/McCarthy, *Functional Magnetic Resonance Imaging*, 122–142; and Mark Jenkinson/Michael Chappell, *Introduction to Neuroimaging Analysis*, Oxford 2018, 27–51.

4 Lerch et al., Studying Neuroanatomy Using MRI, 317.

5 See Paula Muhr, *From Photography to fMRI. Epistemic Functions of Images in Medical Research on Hysteria*, Bielefeld 2022, sections 3.2.1 and 3.2.2. See also Aud Sissel Hoel/Frank Lindseth, Differential Interventions. Images as Operative Tools, in: Kamila Kuc/Joanna Zylinska (eds.), *Photomediations. A Reader*, London 2016, 177–183; and Aurora A. S. Hoel, Images as Active Powers for Reality. A Simondonian Approach to Medical Imaging, in: Emmanuel Alloa/Chiara Cappelletto (eds.), *Dynamics of the Image. Moving Images in a Global World*, Berlin and New York 2020, 287–310.

dardised, reproducible measurement procedures through which these images are generated. In short, it is not their visual properties but their measurement-based character that guarantees MRI images are anatomically plausible and can, therefore, be reliably used by experts for observing, examining, and drawing conclusions about neuroanatomy.

Broadly speaking, there are two major areas of MRI's application. On the one hand, MRI is used in neuroscientific research. Here, MRI is employed to gain insights into the brain's structural organisation and structure-function relations in healthy human subjects, as well as to investigate how pathological changes in brain structure underpin various neurological and psychiatric diseases.[6] On the other hand, MRI is widely utilised in day-to-day clinical practice as "one of the most common diagnostic imaging procedures."[7] Yet, the ways in which MRI is deployed in these two disparate contexts are significantly different.

On the whole, MRI acquisition is typically performed in the form of sequential 2D cross-sectional slices from which a 3D visualisation of the brain can be rendered.[8] In research context, the aim is to generate high-resolution 3D structural MRI images that uniformly and without gaps cover the entire brain in each direction so that there are no missing anatomical structures. Thus typically, images are acquired with isotropic voxels (i.e., with equal sizes in all three directions) whose average size is 1 x 1 x 1 mm or less.[9] The isotropic quality of voxels is crucial in this context as it allows researchers to use automated algorithms for processing and quantitative analysis of the thus obtained MRI images, especially in group studies.[10] However, to obtain generalisable and reproducible findings from their group studies, researchers need to scan many subjects. Since this is both time-consuming and expensive, researchers perennially struggle to acquire sufficient data.

In contrast, millions of MRI scans are routinely generated for diagnostic purposes and stored in hospital archives in clinical settings. But clinical MRI scans are often acquired under time constraints so as not to overstrain patients. Moreover, when used for diagnosing individual patients, MRI scans are not submitted to automated algorithmic analysis. Instead, they are visually inspected by physicians, who are not necessarily interested in examining each patient's entire brain in equal de-

6 See, e.g., Lerch et al., Studying Neuroanatomy Using MRI, 317.

7 Huettel/Song/McCarthy, *Functional Magnetic Resonance Imaging*, 24.

8 Direct 3D imaging is possible but requires longer scanning times, larger data storage, and is more susceptible to artefacts. See, e.g., Huettel/Song/McCarthy, *Functional Magnetic Resonance Imaging*, 117.

9 See Jenkinson/Chappell, *Introduction to Neuroimaging Analysis*, 27–28. The elementary unit of an MRI image is a voxel (i.e., a 3D pixel), whose size determines the spatial resolution of the image.

10 See Jenkinson/Chappell, *Introduction to Neuroimaging Analysis*, 24, 28.

tail.[11] To accommodate the time constraints and facilitate the images' intended uses, MRI scans in clinics are acquired with a high in-plane resolution but with large spacing between consecutive 2D slices. Such scanning procedures result in anisotropic voxels with sizes such as 0.5 x 0.5 x 4 mm.[12] The caveat is that much of the anatomy is not directly measured in such sparsely sampled slices, and, therefore, the anatomy missing from MRI images remains unobserved.

While missing anatomy in sparsely sampled clinical MRI scans is not a problem from the diagnostic perspective, such scans are unusable for research purposes because they cannot be processed with standard image analysis algorithms.[13] Yet from the "big data" perspective,[14] vast clinical repositories of diagnostic MRI scans represent a potentially valuable source of still untapped information on the role of brain anatomy in various diseases. A growing number of recent studies have proposed to resolve this problem and make sparse clinical MRI scans usable for future neuroscientific research by developing specifically tailored machine learning algorithms that can computationally reconstruct the missing anatomy through image imputation.[15] The basic idea is to train artificial intelligence (AI) algorithms to learn to identify underlying anatomical patterns in the sparsely sampled MRI scans and then, based on these patterns, to computationally fill in—i.e., impute—the missing 2D slices that were never directly measured and were thus unobserved (fig. 1).[16]

---

11 See Jenkinson/Chappell, *Introduction to Neuroimaging Analysis*, 27.

12 See Jenkinson/Chappell, *Introduction to Neuroimaging Analysis*, 27.

13 See Adrian V. Dalca/Katherine L. Bouman/William T. Freeman/Natalia S. Rost/Mert R. Sabuncu/Polina Golland, Medical Image Imputation From Image Collections, in: *EEE Transactions on Medical Imaging* 28 (2/2019), 504–514, see 504.

14 John D. Van Horn/Arthur W. Toga, Human Neuroimaging As a 'Big Data' Science, in: *Brain Imaging and Behavior* 8 (2/2013).

15 See, e.g., Yaqiong Chai/Botian Xu/Kangning Zhang et. al., MRI Restoration Using Edge-Guided Adversarial Learning, in: *IEEE Access* 8 (2020), 83858–83870; Dalca et al., Medical Image Imputation; Juan E. Iglesias/Benjamin Billot/Yaël Balbastre et. al., Joint Super-Resolution and Synthesis of 1 mm Isotropic MP-RAGE Volumes from Clinical MRI Exams with Scans of Different Orientation, Resolution and Contrast, in: *NeuroImage* 237 (2021), 118206; Amod Jog/Aaron Carrass/Jerry L. Prince, Self Super-Resolution for Magnetic Resonance Images, in: Sebastien Ourselin/Leo Joskowicz/Mert R. Sabuncu et al.(eds.), *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2016*, Cham 2016, 553–560; Seung K. Kang/Seong A. Shin/Seongho Seo et al., Deep Learning-Based 3D Inpainting of Brain MR Images, in: *Scientific Reports* 11 (6/2021): 1673; and Can Zhao/Blake E. Dewey/Dzung L. Pham et al., SMORE: A Self-Supervised Anti-Aliasing and Super-Resolution Algorithm for MRI Using Deep Learning, in: *IEEE Transactions on Medical Imaging* 40 (3/2021), 805–817.

16 Generally speaking, imputation is not limited to imaging. Instead, in statistics, it refers to a process of completing any type of missing data through statistical evaluation of the existing data. See, e.g., Janus C. Jakobsen/Christian Gluud/Jørn Wetterslev et al., When and How Should Multiple Imputation Be Used for Handling Missing Data in Randomised Clinical Trials—A Practical Guide with Flowcharts, in: *BMC Medical Research Methodology* 17 (1/2017), 162.

Figure 1: Visualisation of sparsely sampled 2D MRI slices stacked into a 3D volume and the missing slices that need to be reconstructed through image imputation. From: Yaqiong Chai/Botian Xu/Kangning Zhang et.al., MRI Restoration Using Edge-Guided Adversarial Learning, in: IEEE Access *8 (2020), 83859, fig. 1a, https://doi.org/10.1109/access.2020.2992204, cc-by-4.0.*

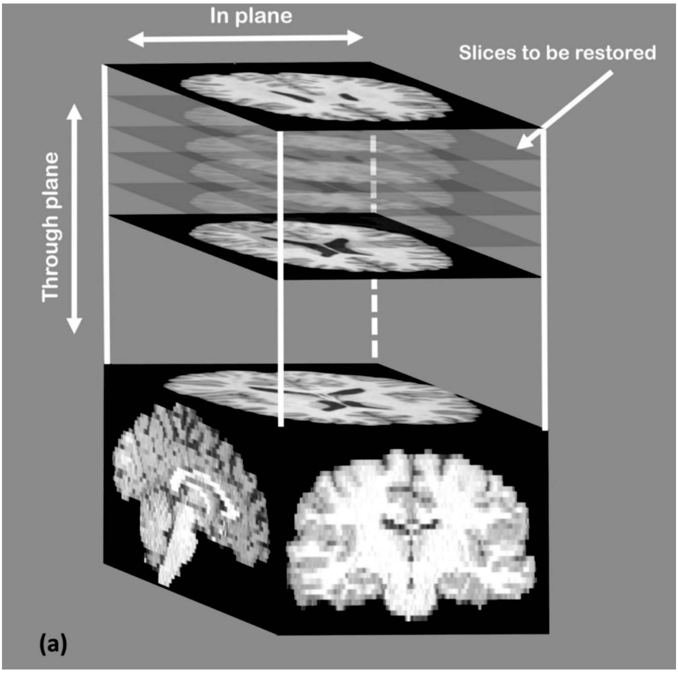

However, this seemingly simple idea is exceedingly challenging to implement in practice. Different studies have proposed using different types of machine learning algorithms for and have worked with different statistical approaches to modelling the missing anatomy in sparse MRI data.[17] The emerging field of MRI image imputation is thus faced with two major questions. First, which AI-based imputation methods are suitable for computing anatomically plausible full 3D MRI images from clinical datasets that comprise sparsely sampled 2D slices? And second, how to prove that the computationally reconstructed brain structures, which were never directly measured, are indeed anatomically plausible and can thus be reliably used to gain insights into the neuroanatomy of actual human brains?

My aim in this chapter is not to offer an overview of this novel research field. Instead, my analysis will focus on a single study, published in 2019, whose authors de-

17 See, e.g., Dalca et al., Medical Image Imputation; and Iglesias et al., Joint Super-Resolution.

veloped a new deep-learning algorithm for synthesising the missing anatomy motivated by the need to analyse clinical MRI scans of thousands of stroke patients.[18] In what follows, I will approach the study by Adrian V. Dalca, Katherine L. Bouman, William T. Freeman, Natalia S. Rost, Mert R. Sabuncu, and Polina Golland from a media-theoretical perspective to examine how its authors addressed the two questions listed above. In doing so, I will outline and discuss how Dalca et al. negotiated the claims of their algorithm's adequacy and the plausibility of the algorithmically reconstructed unmeasured anatomy in sparse clinical MRI scans of stroke patients.

The media-theoretical perspective

My analysis in this chapter is informed by the German media theorist Ludwig Jäger's concept of discursive evidence. According to Jäger, discursive evidence becomes operative when the semantic validity and, thus also, the referential quality of a sign—in our case, MRI images—appears problematic for whatever reason.[19] In such situations, it becomes necessary to establish the legitimacy of the sign's semantic and referential validity through media-specific operations. Such medial operations aim to re-negotiate and thus discursively produce the problematic sign's new meaning. Jäger refers to such medial operations in their entirety as evidential procedures and designates their semantic effects as discursive evidence.

More specifically, Jäger argues that evidential procedures involve the production of a complex network of interactive, reciprocal references not just within a single medium (i.e., intramediality) but also across different media (i.e., intermediality).[20] Moreover, he insists that the production of such intramedial and intermedial references is rhetorically organised, as it is informed by a specific "procedural grammar" of argumentation through which the legitimacy of the problematic sign's validity claim is enacted.[21] The operations entailed in such procedural grammar include paraphrasing, quoting, explaining, demonstrating, comparing, appropriating, commenting, resampling, reinterpreting, translating, and re-addressing.

The crucial aspect, however, is that the discursive evidence generated through such diverse media-based procedures hinges on two things. First, the discursive operations that have enacted the discursive evidence must be fully traceable and avail-

18 See Dalca et al., Image Imputation. In this chapter, I use the terms 'missing,' 'unobserved,' and 'unmeasured' synonymously.

19 See Ludwig Jäger, Schauplätze der Evidenz. Evidenzverfahren und kulturelle Semantik. Eine Skizze, in: Michael Cuntz/Barbara Nitsche/Isabell Otto et al. (eds.), *Die Listen der Evidenz*, Köln 2006, 37–52.

20 See Jäger, Schauplätze der Evidenz, 41.

21 Jäger, Schauplätze der Evidenz, 45.

able for inspection. Or, to quote Jäger, "the procedural visibility of discursive evidence represents an essential moment of its validity."[22] Second, the adequacy and consistency of the discursive operations have to be apparent and beyond doubt in a given context.[23] To emphasise these two aspects, Jäger interchangeably refers to discursive evidence as "evidence as a procedure" or as "procedure-induced" evidence.[24]

Drawing on Jäger's concept of discursive evidence, in this chapter, I will argue that to establish the anatomical plausibility of algorithmically computed missing anatomy in clinical MRI scans, Dalca et al. constructed a complex, rhetorically organised network of intramedial and intermedial references. In line with Jäger, throughout my analysis, plausibility will be treated as a highly dynamic, context-dependent, and multi-layered semantic effect of an ongoing discursive procedure. Thus, instead of defining plausibility in fixed terms, I will explore how plausibility, as an inherently unstable cumulative semantic effect, is processually enacted in my case study.

Depending on which media-specific operations I am analysing, I will examine different procedural aspects of the thus enacted plausibility. For example, I will discuss the methodological plausibility of using a specific algorithm for a particular task but also the referential and operative plausibility of imputed MRI images according to their intended future uses. Additionally, I will elaborate on how the imputed images' visual and quantitative plausibility is constituted through particular medial procedures. In short, for the sake of my analysis, I will pry apart various multifaceted procedural aspects that underpin the establishment of plausibility in my case study. In doing so, I will aim to show that only when all these dynamic aspects are successfully integrated does plausibility emerge as a valid and verifiable semantic effect. But before turning to the analysis of the operational procedure through which Dalca et al. achieved the enactment of plausibility in their study, we first need to examine the epistemic challenges they faced and the solution they provided by developing a novel image imputation method.

Epistemic challenges and their proposed algorithmic solution

As Dalca et al. explicitly stated in the introduction to their study, the development of their novel image imputation method was motivated by a distinctly pragmatic aim. They wanted to make it possible to study "brain MRI scans of thousands of stroke patients acquired within 48 hours of stroke onset" by using standard anal-

22 Jäger, Schauplätze der Evidenz, 46 (my translations).
23 See Jäger, Schauplätze der Evidenz, 46.
24 Jäger, Schauplätze der Evidenz, 46.

ysis algorithms to "quantify white matter disease burden."[25] Due to the patients' acute states, the MRI acquisition had to be performed quickly, resulting in extremely sparsely sampled T_2-FLAIR imaging data with anisotropic voxels sized 0.85 x 0.85 x 6 mm.[26] When 2D slices obtained with such large inter-slice spacing were combined into 3D volumes, much of the patients' anatomy was missing in the through-plane direction of slice acquisition. As mentioned earlier, such significant gaps in data lead to inaccuracies during quantitative algorithmic analysis. But just as problematically, the missing data also derail the performance of basic algorithmic pre-processing steps—such as skull stripping and registration to a common reference space—which are necessary to prepare MRI data for downstream analysis.[27]

Moreover, because "the anatomical structure may change substantially between consecutive slices," making anatomically plausible inferences about the missing data in the slice-acquisition direction is highly challenging.[28] Consequently, various standard interpolation methods used in medical imaging for upsampling low-resolution images to a higher resolution deliver poor results when applied to sparse MRI data.[29] Generally speaking, such methods operate under two assumptions. It is presumed that, first, sparsely measured voxels in the 3D images' through-plane direction can be represented by some generic mathematical functions; and second, that by estimating function parameters, it is possible to infer missing voxel values.[30] To achieve this, so-called single-image methods rely on what is referred to as pattern redundancy or self-similarity—i.e., the fact that "small-scale structures tend to repeat themselves throughout an image."[31] In principle, such repetitive details within a single image can be used "to re-synthesize [missing] high frequency information."[32] Apart from single-image methods, an alternative approach to up-sampling low-resolution medical images is to implement so-called multi-image methods. The latter type of method fuses information across multiple scans of the same subject, provided the scans were acquired from different orientations.[33]

25 Dalca et al., Medical Image Imputation, 504.

26 See Dalca et al., Medical Image Imputation, 504.

27 See Dalca et al., Medical Image Imputation, 504.

28 Dalca et al., Medical Image Imputation, 504.

29 For various standard interpolation methods, see Thomas M. Lehmann/Claudia Gonner/Klaus Spitzer, Survey. Interpolation Methods in Medical Image Processing, in: *IEEE Transactions on Medical Imaging* 18 (11/1999), 1049–1075.

30 See José V. Manjón/Pierrick Coupé/Antonio Buades et al., Non-Local MRI Upsampling, in: *Medical Image Analysis* 14 (6/2010), 784–792.

31 Esben Plenge/Dirk H. Poot/Wiro J. Niessen et al., Super-Resolution Reconstruction Using Cross-Scale Self-Similarity in Multi-Slice MRI, in: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 6* (2013), 123–30, see 123–24.

32 Dalca et al., Medical Image Imputation, 505.

33 See Plenge et al., Super-Resolution,123.

However, both single-image and multi-image upsampling methods struggle when applied to severely undersampled clinical MRI scans. Various single-image methods fail because there is not "enough fine-scale information to provide anatomically plausible reconstruction in the direction of slice acquisition."[34] Conversely, using multi-image upsampling methods is often not feasible in the clinical context, since scans of a single subject's brain from different orientations are rarely available.

In response to these challenges, more recent approaches to medical image upsampling, of which the Dalca et al. study is a pertinent example, have deployed supervised machine learning algorithms. Without being explicitly programmed, such algorithms learn to synthesise missing parts of an image through training on external datasets that typically consist of matching pairs of low-resolution and high-resolution images.[35] By relating an input image to a correct output image, the external training pairs allow the algorithm to learn a mapping function that enables subsequent recovery of high-resolution from low-resolution images. However, the problem Dalca et al. encountered was the lack of an external dataset with non-sparse clinical MRI scans of stroke patients, which they could deploy for training a supervised algorithm. And as Dalca et al. explained, using existing high-resolution MRI scans from healthy subjects as training data did not present a feasible alternative option. In the latter case, the problem was that scans of healthy subjects "may not adequately represent pathology or other properties of clinical populations," which Dalca et al. aimed to study.[36]

To circumvent these problems, Dalca et al. chose to employ the approach called unsupervised machine learning. This type of machine learning utilises algorithms that learn to identify some underlying hidden structure directly in the data of interest without going through a training phase that requires an external dataset.[37] Next, Dalca et al. defined their unsupervised learning task in terms of building a probabilistic generative model that would discover "fine-scale anatomical structure across subjects" in sparse clinical MRI images.[38] Their approach was grounded in the

34 Dalca et al., Medical Image Imputation, 505.

35 See, e.g., Iglesias et al., Joint Super-Resolution, 2. Training data are often obtained by blurring and subsampling high-resolution images to obtain their low-resolution counterparts. See Iglesias et al., Joint Super-Resolution, 2.

36 Dalca et al., Medical Image Imputation, 505.

37 See, e.g., Jason Brownlee, Supervised and Unsupervised Machine Learning Algorithms, in: *Machine Learning Mastery* June 17, 2019, last updated July 19, 2019, URL: https://machinelearningmastery.com/supervised-and-unsupervised-machine-learning-algorithms [last accessed May 2, 2022].

38 Dalca et al., Medical Image Imputation, 504. "Generative modeling is an unsupervised learning task in machine learning that involves automatically discovering and learning the regularities or patterns in input data in such a way that the model can be used to generate or

assumption that, despite individual variations in anatomy across subjects, all sparse 3D MRI images in their clinical collection shared similar local fine-scale anatomical structures. From this perspective, each patient's sparse 3D image "captures some partial aspect" of the underlying anatomical pattern, which is presumed to be common to all images comprising the collection.[39] Put simply, each 3D MRI image is regarded as a variation—or "a noisy manifestation"—of some shared latent pattern.[40] To account for anatomical variability across patients, Dalca et al. decided to mathematically represent the unknown latent anatomical pattern by a statistical model called the Gaussian Mixture Model (GMM).[41] Having defined the type of their generative model, they then had to fit this model to their data. Thus during the learning phase, Dalca et al. deployed the iterative expectation-maximisation (EM) algorithm to estimate the model parameters that best describe the latent anatomical pattern in their collection of clinical MRI scans.

Three aspects of their algorithmic modelling need to be foregrounded for our discussion. First, during parameter estimation, the algorithm learned from a collection of severely undersampled 3D T_2-FLAIR MRI scans of 766 stroke patients. These scans were first registered to "the isotropically sampled common atlas space" and then stacked together.[42] In other words, the algorithm did not learn from a single 3D image but from all mutually aligned 3D images comprising the clinical image collection. But instead of simultaneously operating on the entire brain volumes, the algorithm independently fitted the generative model to separate 3D patches, i.e., predefined smaller image regions. To define the patches, Dalca et al. divided the stacked 3D images into smaller subvolumes, with neighbouring subvolumes mutually overlapping in each direction.[43] Parameter estimation was performed separately at each

---

output new examples that plausibly could have been drawn from the original dataset." Jason Brownlee, A Gentle Introduction to Generative Adversarial Networks (GANs), in: *Machine Learning Mastery* March 16, 2016, last updated August 20, 2020, URL: https://machinelearningmastery.com/what-are-generative-adversarial-networks-gans [last accessed May 2, 2022].

39 Dalca et al., Medical Image Imputation, 504.

40 Dalca et al., Medical Image Imputation, 504.

41 See Douglas Reynolds, Gaussian Mixture Models, in: Stan Z. Li/Anil Jain (eds.), *Encyclopedia of Biometrics*, Boston 2009, 659-63.

42 Dalca et al., Medical Image Imputation, 508. By registering the undersampled scans from subject-specific to the atlas space, Dalca et al. transformed the images' voxels from anisotropic (0.85 x 0.85 x 6 mm) into isotropic (1 x 1 x 1 mm). To enable this transformation, they approximated the resampled voxels in the atlas space as "either observed or missing." Dalca et al., Medical Image Imputation, 508. Working in the atlas instead of the subject-specific space allowed them to avoid "computationally prohibitive updates" during parameter estimation. Dalca et al., Medical Image Imputation, 508.

43 See Dalca et al., Medical Image Imputation, 508. The "subvolumes of 21 x 21 x 21 voxels in the isotropically sampled common atlas space were centered 11 voxels apart in each direction", resulting in cubic patches sized 11 x 11 x 11 voxels. And "instead of selecting just one

overlapping region, which defined a patch. This aspect of Dalca et al.'s approach was conceptually highly significant, enabling them to flexibly capture the differences in local anatomical structures at various locations throughout the 3D brain volume. Moreover, creating overlapping patches within subvolumes had an added benefit. "This aggregation provides more data for each model, which is crucial when working with severely undersampled volumes."[44]

Second, Dalca et al. treated both the observed and the missing voxels in the stacked 3D volumes as constitutive parts of the model. This allowed them to use "simpler algorithms that iteratively fill in missing voxels and then estimate GMM model parameters."[45] Third, Dalca et al. chose not to define in their model the size of the interslice gaps with which the sparse MRI clinical scans were initially acquired. As a result, their model could be applied to "data with varying sparseness patterns, including restoring data in all acquisition directions simultaneously."[46] Taken together, the three modelling aspects aimed to increase the algorithm's precision, speed, and flexibility.

Once their algorithm had learned from the clinical MRI scans the generative model's parameters for each patch separately, Dalca et al. could deploy this model to individually impute each patient's full 3D MRI image in their sparse data collection. By inverting the model, they first reconstructed each 3D image patch separately and then stitched the reconstructed patches together to form a patient's full 3D volume.[47] Notably, to obtain smoother images, Dalca et al. computed not just the missing voxels in each patch but also reconstructed "the entire patch, including the observed voxels."[48] This meant that in the full 3D MRI image, which they had to compute individually for each patient, none of the resulting voxels stemmed directly from the measurement. In short, in the imputed MRI images, the indexical relation to the actual patients' physical brains was broken.

Thus, based on my analysis so far, the question is: How could Dalca et al. convincingly claim that their imputed 3D MRI images were anatomically plausible and can, therefore, be used in neuroscientific research to gain insights into actual physical brains? Informed by Ludwig Jäger's concept of discursive evidence, my following analysis will show that Dalca et al. enacted the anatomical plausibility of their algorithmically imputed images through a discursive procedure that comprised two

patch from each volume at a given location, we collect all overlapping patches within the subvolume centered at that location." Dalca et al., Medical Image Imputation, 508.

44 Dalca et al., Medical Image Imputation, 508.

45 Dalca et al., Medical Image Imputation, 511.

46 Dalca et al., Medical Image Imputation, 511. The computationally more expensive alternative would have been to define missing voxels as latent (i.e., unknown) variables. Dalca et al., Medical Image Imputation, 511.

47 See Dalca et al., Medical Image Imputation, 505.

48 Dalca et al., Medical Image Imputation, 511.

stages. First, they negotiated the plausibility of their AI-based imputation method; and second, they negotiated the plausibility of imaging results obtained through this imputation method. As we are about to see, each stage consisted of multiple steps and required the establishment of a dynamic network of intramedial and intermedial references.

Negotiating the plausibility of the novel imputation method

The initial step Dalca et al. took to demonstrate the scientific plausibility of their novel algorithmic image imputation method entailed embedding their study into a broader context of previous research. As discussed above, their new method was specifically tailored to their research question and the epistemic challenges that arose from it. But in developing their method, Dalca et al. drew on earlier studies that had successfully applied patched-based methods and generative models such as GMM to processing both medical and so-called natural images.[49] By explicitly referencing such studies, Dalca et al. effectively argued that they were adopting and adapting the modelling procedures that had already been proven to deliver plausible results in specialised medical contexts, as well as, more broadly, in computer vision.[50] However, their argument had one caveat—all the previous studies used patched-based methods and GMM to classify or denoise images, which meant that these methods were "typically not designed to handle missing data."[51]

Thus, in the subsequent step, Dalca et al. developed a twofold discursive strategy. First, they focused on describing and justifying how they adapted the patch-based approach to their particular learning task. And second, they elaborated on how they chose to construct a generative model for sparse image patches specifically tailored to their data characteristics. With this aim in mind, Dalca et al. provided a clear-cut explanation of the mathematical underpinnings of their generative model.[52] This was followed by a detailed mathematical description of the algorithm's learning process and the maximum-a-posteriori estimate through which they restored the miss-

49 In computer vision, the term 'natural images' is typically used to denote "everyday [photographic] images taken with regular RGB cameras." Thijs Kooi, Deep Learning: From Natural to Medical Images, in: *Medium* December 19, 2018, URL: https://medium.com/merantix/deep-learning-from-natural-to-medical-images-74827bf51d6b [last accessed May 2, 2022].

50 For the complete list of quoted studies, see Dalca et al., Medical Image Imputation, 505. For example, one of the quoted studies was Daniel Zoran/Yair Weiss, Natural Images, Gaussian Mixtures and Dead Leaves, in: *Proceedings of the 25th International Conference on Neural Information Processing Systems*, Red Hook 2012, 1736–1744.

51 Dalca et al., Medical Image Imputation, 505.

52 See Dalca et al., Medical Image Imputation, 506.

ing data during the imputation.[53] Dalca et al. paid particular attention to clearly delineating each decision they made during the patch selection, modelling, parameter selection, and the subsequent imputation. And just as importantly, for each of these decisions, they discussed possible alternatives, as well as the epistemic consequences of such alternatives.[54] The elaborate description, discussion, and rhetorical justification of each aspect entailed in the development of the new algorithmic method were all the more crucial since, during its implementation, the unsupervised algorithm operated as a black box. Specifically, the exact details of the iterative steps through which the algorithm fitted GMM to the sparse data and later used this model to restore the missing anatomy remained opaque even to Dalca et al.[55]

On the whole, the implicit message generated by Dalca et al. through this "procedural grammar" of argumentation was that no aspect of their newly developed imputation method was arbitrary.[56] Instead, each aspect of the new method was carefully considered, and its potential epistemic consequences were evaluated from multiple perspectives. Through this discursive procedure, the method was established as plausible because each of its aspects was proven to be adequately tailored to the characteristics of the sparse MRI data and to the task of imputing the missing anatomy in such data.

But, in my view, the most persuasive discursive support for the plausibility of the new algorithmic imputation method was delivered in the third step of the authors' argumentation. In this step, Dalca et al. compared the quality of performance of their method to other "state-of-the-art upsampling" algorithmic methods.[57] For the purpose of this comparison, Dalca et al. deployed two alternative methods used in

53 See Dalca et al., Medical Image Imputation, 506–507.

54 For example, Dalca et al. discussed the advantages and disadvantages of defining missing voxels as part of the model instead of designating them as latent variables. They also examined the effects of filling in only the missing but not the observed voxels during the imputation. See Dalca et al., Medical Image Imputation, 508–512.

55 Concerning the opaque, black-boxed character of machine learning algorithms, see, e.g., Malte Ziewitz, Governing Algorithms. Myth, Mess, and Methods, in: *Science, Technology & Human Values* 41 (1/2016), 3–16.

56 Jäger, Schauplätze der Evidenz, 45.

57 Dalca et al., Medical Image Imputation, 504. My analysis in this chapter is informed by the view that instead of being a fixed, static methodological tool, comparing "is a 'relationing' activity" and "a social and historical practice [that] is always bound up with actors and agencies that perform the comparisons and connect them with their purposes and posible outcomes—intended or not." Angelika Epple/Walter Erhart, Practices of Comparing. A New Research Agenda Between Typological and Historical Approaches, in: *Practices of Comparing. Towards a New Understanding of a Fundamental Human Practice*, Bielefeld 2020, 11–38, see 17. Hence, to uncover the epistemic consequences of each comparison and its role in the construction of discursive evidence in our case study, we have to analyse how each comparison was performed and for which particular purposes.

image processing—non-local means and linear interpolation—to impute the missing anatomical structures in their clinical MRI images of stroke patients. They then visually compared the thus obtained results to the images imputed from the same sparse dataset using their new method.[58] The visual comparison was performed at two levels.

On the one hand, by comparing the respective whole-brain views, Dalca et al. assessed how well each method performed when reconstructing large-scale anatomical structures. On the other hand, by utilising multiple close-up panels, the researchers also zoomed in on anatomical substructures to comparatively assess the methods' ability to reconstruct finer details. By constructing these multiple comparative intramedial references across images in which the missing anatomy was computationally synthesised using different methods, Dalca et al. generated convincing discursive evidence for the task-specific plausibility of their imputation method. The direct visual comparison of the imputed images showed that the new method outperformed competing methods by producing "more plausible structure[s], as can be especially seen in the close-up panels focusing on anatomical details."[59]

Yet, despite the undeniable rhetorical persuasiveness of this visual comparison across the methods, one question remained open. So far, Dalca et al. have proven two things through their carefully constructed evidential procedure. First, their new algorithmic approach was methodologically plausible for the intended task; and second, it produced 3D image reconstructions that were relationally more anatomically plausible than those obtained through competing algorithmic methods. However, how could Dalca et al. justifiably claim that their computationally imputed images were sufficiently anatomically plausible in relation to patients' physical brains?

Negotiating the plausibility of imputed images

At this point, Dalca et al. needed to prove that apart from being created through a methodologically adequate algorithmic process, their imputed 3D MRI images were also anatomically plausible in the sense of being able to refer to and thus stand in for actual physical brains. Put simply, Dalca et al. had to generate persuasive discursive evidence for the referential plausibility of the imputed anatomical structures that were never directly measured. What hinged on the quality of such evidence was the operative capability of the algorithmically imputed images to be used in future scientific studies for making judgments about the neuroanatomy of actual brains. In

58 See Dalca et al., Medical Image Imputation, 512, fig. 10.

59 Dalca et al., Medical Image Imputation, 510.

the remainder of the chapter, I will argue that to enact the requisite discursive evidence, Dalca et al. deployed a multistep evidential procedure through which they visually, quantitatively, and operatively negotiated the plausibility of their imputed images.

The problem that Dalca et al. faced was the following. To visually and quantitatively establish the referential plausibility of the algorithmically imputed MRI images, it was necessary to directly compare them to 3D isotropic images of the same patients, which had been acquired by directly measuring the patients' neuroanatomy. However, no such direct comparison was possible because there were no MRI images without the missing anatomy of the stroke patients in Dalca et al.'s clinical dataset. In fact, as discussed earlier, the lack of the stroke patients' full 3D isotropic MRI images motivated the development of the Dalca et al. study in the first place.

Dalca et al. resolved this problem by reverting to an external database of full 3D MRI scans without missing anatomy. The purpose of this external database, which in specialists' terms is called 'ground truth,' was to provide the researchers with a stable external frame of reference for evaluating if their novel imputation method delivered referentially plausible results when reconstructing missing anatomy.[60] Using ground truth datasets for assessing the quality of performance of newly developed machine learning algorithms is standard practice and is considered of central importance for demonstrating a new algorithm's correctness.[61] The decisive aspect in this practice is choosing as ground truth an adequate set of data that have shared salient characteristics with the data to which the algorithm will later be applied. In this respect, it is worth pointing out that Dalca et al. did not choose to employ scans of healthy subjects as their external frame of reference. Instead, without explaining their decision, they chose to use 826 T_1-weighted 3D MRI images from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database.[62] Perhaps Dalca et al. reasoned that a dataset from a neurological patient population had more shared properties with their dataset of stroke patients than MRI scans of healthy subjects.

60 See Dalca et al., Medical Image Imputation, 505.

61 See Florian Jaton, Assessing Biases, Relaxing Moralism: On Ground-Truthing Practices in Machine Learning Design and Application, in: *Big Data & Society* 8, (1/2021), 1–15. In this context, the term 'ground truth' does not imply the existence of some universal or permanent truth. Instead, it is used in a purely relational sense to denote a study-specific benchmark reference for establishing a new algorithm's capability to produce correct results. It is in this sense that I will use the term ground truth in this paper.

62 As succinctly explained by Dalca et al., the "primary goal of ADNI has been to test whether serial magnetic resonance imaging (MRI), positron emission tomography (PET), other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of mild cognitive impairment (MCI) and early Alzheimers disease (AD)." See also https://adni.loni.usc.edu.

However, it should be noted that the MRI images in the ADNI database stemmed from patients with a very different illness from those in Dalca et al.'s clinical collection (i.e., Alzheimer's disease versus stroke).[63] Moreover, the images in the ADNI database were acquired in a different MRI modality. Specifically, the scans from the ADNI database were T_1-weighted images, whereas those of stroke patients were T_2-FLAIR images. Dalca et al. remained tacit about these differences.

After they had defined their external frame of reference, Dalca et al. still had to perform several additional steps to be able to utilise it for the intended purposes. The ADNI database provided Dalca et al. with full 3D scans they could use as ground truths, i.e., benchmark images to which they could visually and quantitatively compare their algorithmically reconstructed images to assess how accurately their novel algorithm imputed the missing anatomy.[64] Yet, to perform such an assessment, Dalca et al. also needed a referential sparse MRI dataset from the same subjects as the non-sparse ADNI images, which they could use as input for their imputation algorithm. To create a referential sparse MRI dataset, Dalca et al. downsampled the original isotropic ADNI images with voxel sizes of 1 x 1 x 1 mm "to slice separation of 6mm (1mm x 1mm in-plane) in the axial direction."[65] They thus obtained a referential input dataset that closely emulated the sparseness pattern in their clinical dataset of stroke patients.

Next, Dalca et al. fed the artificially produced sparse MRI data to their new algorithm and to two alternative upsampling algorithms—non-local means and linear interpolation algorithms. From our discussion in the previous section, we are already familiar with such comparative use of competing methods to impute isotropic 3D images from the same sparse MRI dataset. This time, however, Dalca et al. could directly compare the imputed images with the corresponding ground truth images—i.e., full 3D images obtained by directly measuring the brain anatomy of Alzheimer's patients. The visual comparison showed that the images imputed through the method developed by Dalca et al. were significantly more similar to the ground truth images than the reconstructions produced by the two alternative methods. Specifically, the visual comparison revealed that the new method managed to restore "anatomical structures that are almost entirely missing in the other reconstructions, such as the dura or the sulci of the temporal lobe."[66]

Based on the visual similarity between the images in which the missing anatomy was imputed by their method and the ground truth images obtained by direct measurement of the patients' physical brains, Dalca et al. could now claim that their al-

63 As discussed earlier, different MRI imaging modalities are obtained through different measurement settings and thus visually encode mutually disparate aspects of the brain anatomy.

64 Dalca et al., Medical Image Imputation, 508.

65 Dalca et al., Medical Image Imputation, 508.

66 Dalca et al., Medical Image Imputation, 509.

gorithm was indeed capable of computing anatomically plausible 3D images from sparse MRI data. Moreover, because the novel method imputed anatomically plausible 3D images from the artificially produced sparse ADNI data, by implication, Dalca et al. could argue that it also imputed anatomically plausible 3D images from the sparse clinical data of stroke patients. In effect, the discursive operations of visual comparisons that established multi-layered intramedial references across differently imputed images and the ground truth images enabled Dalca et al. to renegotiate the anatomical plausibility of the algorithmically imputed missing anatomy. Through these operations, Dalca et al. succeeded in supplanting the lack of the imputed images' direct indexical relation to actual physical brains with the iconic relation of visual similarity to the indexical ground truth images.

Yet even at this point, the evidential procedure of enacting the imputed images' anatomical plausibility was not completed. In the following step, Dalca et al. shifted from the visual to the quantitative comparison of similarity between the imputed ADNI images and their corresponding ground truths, i.e., the original isotropic 3D images. This additional level of comparison allowed Dalca et al. to estimate the uncertainty of the imputed images through formal mathematical tools. To facilitate the quantitative comparison, Dalca et al. chose two metrics "commonly used in measuring the quality of reconstruction of compressed or noisy signals"—mean squared error and peak signal-to-noise ratio.[67] Apart from their new imputation method, Dalca et al. also calculated these two metrics for the images reconstructed by three other alternative upsampling methods: nearest neighbour interpolation, non-local means, and linear interpolation. The comparison of the thus obtained metrics across the four methods showed that the full 3D images reconstructed through the new method were relationally more accurate since their error metrics were lower than those of the competing image imputation methods.[68] This, in turn, meant that the images restored through the new method were not just visually but also quantitatively more similar to the ground truth images than those restored through other methods. By calculating the error statistics, Dalca et al. established an additional network of intermedial references. In doing so, they added another discursive aspect to their argumentation, further reinforcing the anatomical plausibility of the images imputed through their method.

But the crowning step of their already elaborate evidential procedure was still to follow. In this final step, Dalca et al. turned to generating discursive evidence that the images imputed by their method were anatomically plausible in operative terms. In other words, they focused on demonstrating that the imputed images were sufficiently anatomically plausible to be used by existing analysis "algorithms that were

67 Dalca et al., Medical Image Imputation, 509.

68 See Dalca et al., Medical Image Imputation, 509.

originally developed for high resolution [MRI] research scans."[69] With this aim in mind, Dalca et al. performed a typical preprocessing step called skull stripping on the original isotropic ADNI images. They then applied the same preprocessing step to 3D scans imputed from the sparse ADNI data by their new method, as well as the by now familiar two alternative methods—non-local means and linear interpolation.

During skull stripping, automated algorithms extract the brain in MRI images by isolating it from the rest of the anatomy. To carry out this operation, the preprocessing algorithms have to identify non-brain voxels, which are then treated as noise and deleted from the image.[70] Dalca et al. argued that automated skull stripping fails not only when applied to sparse MRI data but also when processing full 3D images restored by non-local means or linear interpolation. In contrast, Dalca et al. claimed that when applied to images imputed by the novel algorithm, skull stripping "dramatically improves."[71]

Once again, Dalca et al. rhetorically organised their evidence for this claim by establishing a highly persuasive visual comparison. They showed four skull stripping examples positioned one next to the other in a single composite figure.[72] In the first two examples, skull stripping was performed on brain volumes restored by the alternative upsampling methods. In the third example, skull stripping was performed on a brain volume imputed by the new method and, in the fourth example, on a ground truth isotropic volume. The first two images in the composite figure show the brain encased within the skull, thus visually demonstrating the outright failure of skull stripping on images restored by the upsampling algorithms. In contrast, the other two images within the figure show a brain isolated from the rest of the anatomy. Moreover, the two images of the successfully extracted brain are broadly similar.[73] The latter two images thus demonstrate in clear visual terms that the automated skull stripping algorithm treats the image imputed by the new method similarly to the original isotropic images obtained by directly measuring an actual physical brain. Put simply, the skull stripping algorithm 'recognises' the image imputed by the new method as anatomically plausible.

Thus, the visual comparison of the different skull stripping examples provided the crowning discursive evidence for the anatomical plausibility of the 3D isotropic volumes synthesised from sparse 2D MRI data using the new method developed by

69 Dalca et al., Medical Image Imputation, 504.

70 See Jenkinson/Chappell, *Introduction to Neuroimaging Analysis*, 92–93.

71 Dalca et al., Medical Image Imputation, 511.

72 See Dalca et al., Medical Image Imputation, 511, fig. 9.

73 The image obtained through skull stripping of the image imputed by the new method has less detailed, blurrier borders that the one obtained through skull stripping of the ground truth image. See Dalca et al., Medical Image Imputation, 511, fig. 9.

Dalca et al. This final visual comparison showed that the imputed missing anatomy is not just visually and quantitatively plausible but also operatively plausible from the perspective of standard algorithms for MRI processing. Since Dalca et al. set out to develop an imputation method that would make sparse clinical MRI data analysable with standard image-analysis algorithms, this final visual comparison indicated in a rhetorically effective way that they have successfully achieved their predefined goal.

Conclusion: The relational quality of discursively established plausibility

Overall, my analysis in this chapter has aimed to show that the emerging research field of AI-based imputation of sparse clinical MRI images operates under the assumption that the missing anatomy can be reliably estimated and statistically modelled using the available anatomical information in the sparse images. As demonstrated by the case study discussed above, this field is driven by the development of increasingly sophisticated machine learning methods that learn to identify underlying anatomical patterns in the sparse 2D imaging data and then use the thus identified patterns to infer the anatomical structures that were never directly measured and thus remaind unobserved.

Admittedly, the imputed data are not considered anatomically as accurate as directly measured MRI data of a single subject's brain and should, therefore, "not be used in clinical evaluation" of individual patients.[74] Instead, somewhat more modestly, the aim is to algorithmically reconstruct 3D images that are sufficiently anatomically plausible to be used in "large scale scientific studies" whose goal, in turn, is to produce generalisable knowledge about a particular clinical population's neuroanatomy by averaging MRI data across numerous subjects. In this context, it is crucial for researchers who develop new imputation methods to prove that their black-boxed algorithms can compute sufficiently anatomically plausible reconstructions of the missing anatomy in sparse MRI scans. As I have argued in this chapter, to establish the anatomical plausibility—and thus also the scientific validity of their algorithmically imputed scans—researchers must employ the "procedural grammar" required to enact what Ludwig Jäger has pertinently termed discursive evidence.[75]

Using the example of the Dalca et al. study, I have analysed the targeted medial procedures through which researchers discursively negotiate the methodological plausibility of their new imputation approach and the anatomical plausibility of

74 Dalca et al., Medical Image Imputation, 504.

75 Jäger, Schauplätze der Evidenz, 45.

the thus reconstructed images. On the one hand, the negotiation entails constructing a complex network of dynamic intramedial references across different sparse and isotropic MRI images. On the other hand, it also entails constructing an equally complex network of dynamic intermedial references across various MRI images, different algorithmic imputation methods, and select statistical error metrics. As we have seen, the crucial questions during this process are, first, how to procedurally organise different visual, mathematical, and operative comparisons; second, in which medium to perform a particular comparison; and third, which similarities and differences to emphasise. The thus established anatomical plausibility of the algorithmically reconstructed images is unavoidably relational as it hinges on the adequacy and consistency of the evidential procedure that produced it.

Moreover, such discursively enacted plausibility is inherently unstable because the evidential procedure that produced it is context-dependent and can always be challenged by other researchers in the field. For example, in future studies, the anatomical plausibility of MRI images imputed by Dalca et al.'s algorithm could be disputed by researchers who have established an alternative frame of intramedial and intermedial references to justify the introduction of a different imputation method. Yet it seems to me that, instead of being a drawback, this inherent instability of the discursively enacted semantic effects is, in fact, productive. It drives researchers to progressively increase the relational plausibility of the imputed sparse MRI clinical scans by building on each other's work.

Finally, it is worth emphasising another relational aspect of the discursively negotiated plausibility of Dalca et al.'s new imputation method. As discussed above, Dalca et al. explicitly developed their method for reconstructing sparse clinical MRI data of thousands of stroke patients acquired within 48 hours after stroke onset. We have seen that, during their protracted evidential procedure, the researchers chose to deploy MRI scans of patients with Alzheimer's disease as an external referential frame to test the correctness of their algorithm. I have also emphasised that at no point during the evidential procedure did Dalca et al. test their algorithm on MRI scans of healthy subjects, which they could have taken from one of the standard atlas databases. Instead, they focused exclusively on demonstrating that their new imputation algorithm plausibly operated on different types of brain pathologies.

This relational focus already indicated that the researchers envisaged the application of their algorithm for imputing MRI scans stemming from a wider clinical population beyond stroke patients. Such intentions were implicitly confirmed by Dalca et al. near the end of their paper when they pointed out as a potential limitation of their algorithm that it was not universally applicable to all brain pathologies.[76] As Dalca et al. explained at this point, their algorithm required the preservation of sufficient similarity in anatomical structures across patients. Consequently,

76 Dalca et al., Medical Image Imputation, 511.

their algorithm could not be used for imputing sparse MRI scans of patients with traumatic brain injuries or tumours because, in such clinical cases, MRI scans "may not contain enough consistency to enable the model to learn meaningful covariation."[77]

But despite the authors' explicitly stated intention to develop a machine learning algorithm that successfully imputes missing anatomy in clinical data, I find their apparent disinterest in additionally testing how their algorithm performed on healthy brains somewhat surprising. Perhaps they presumed that if their algorithm performed well on clinical data, by implication, it also performed well on the data of healthy subjects. This, however, does not appear self-evident to me, considering how vastly different the neuroanatomies of healthy individuals can be. Alternatively, maybe Dalca et al. thought there was no sufficient need for imputing MRI scans of healthy subjects, since such scans are typically not acquired under time constraints. However, I suggest that through their strict focus on using clinical data, Dalca et al. discursively, although perhaps inadvertently, inscribed an implicit dichotomy between the 'normal' and the 'pathological' brain into their method.

Furthermore, it seems to me that the potential usefulness of future imputation algorithms, should these be adequately shown to generate anatomically and operatively plausible images from sparse data, is not necessarily limited to clinical MRI datasets. Instead, in principle, such algorithms could also be useful for neuroscientific studies that require MRI data of numerous healthy subjects. In the latter case, it is conceivable that imputation algorithms would allow researchers to collect slightly sparser MRI data of healthy subjects, thus lowering the required scanning time and the costs of performing large-scale studies.

On the whole, with such rich possibilities of application in neuroscientific research, it is safe to assume that new, more sophisticated machine learning methods for MRI imputation will continue to be developed. But, as shown in this chapter, the potential operative usability of each such method will necessarily hinge on the successful negotiation of its plausibility. Therefore, with the development of each new image imputation method, its authors will always be required to establish a new, specifically tailored network of context-dependent intramedial and intermedial references to discursively justify the method's ability to algorithmically synthesise anatomically and operatively plausible MRI images of the unobserved human neuroanatomy. Yet, by its very nature, such relational, dynamically enacted plausibility will always remain semantically unstable and open to further discursive revision.

77 Dalca et al., Medical Image Imputation, 511.

References

BROWNLEE, Jason, A Gentle Introduction to Generative Adversarial Networks (GANs), in: *Machine Learning Mastery* June 17, 2019, last updated July 19, 2019, URL: https://machinelearningmastery.com/what-are-generative-adversarial-networks-gans [last accessed May 2, 2022].

DERS., Supervised and Unsupervised Machine Learning Algorithms, in: *Machine Learning Mastery* March 16, 2016, last updated August 20, 2020, URL: https://machinelearningmastery.com/supervised-and-unsupervised-machine-learning-algorithms [last accessed May 2, 2022].

CHAI, Yaqiong/XU, Botian/ZHANG, Kangning et al., MRI Restoration Using Edge-Guided Adversarial Learning, in: *IEEE Access* 8 (2020), 83858–83870. https://doi.org/10.1109/access.2020.2992204.

DALCA, Adrian V./BOUMAN, Katherine L./FREEMAN, William T. et al., Medical Image Imputation From Image Collections, in: *EEE Transactions on Medical Imaging* 38 (2/2019), 504-514. https://doi.org/10.1109/TMI.2018.2866692.

EPPLE Angelika/ERHART, Walter, Practices of Comparing. A New Research Agenda Between Typological and Historical Approaches, in: *Practices of Comparing. Towards a New Understanding of a Fundamental Human Practice*, Bielefeld 2020, 11–38.

HOEL, Aurora A. S., Images as Active Powers for Reality. A Simondonian Approach to Medical Imaging, in: Alloa, Emmanuel/Cappelletto, Chiara (eds.), *Dynamics of the Image. Moving Images in a Global World*, Berlin/New York 2020, 287–310.

HOEL, Aud Sissel/LINDSETH, Frank, Differential Interventions. Images as Operative Tools, in: Kuc, Kamila/Zylinska, Joanna (eds.), *Photomediations. A Reader,* London 2016, 177–183.

HUETTEL, Scott A./SONG, Allen W./MCCARTHY Gregory, *Functional Magnetic Resonance Imaging*, Sunderland 2009.

IGLESIAS, Juan E./BILLOT, Benjamin/BALBASTRE, Yaël et al., Joint Super-Resolution and Synthesis of 1 mm Isotropic MP-RAGE Volumes from Clinical MRI Exams with Scans of Different Orientation, Resolution and Contrast, in: *NeuroImage* 237 (2021), 118206. https://doi.org/10.1016/j.neuroimage.2021.118206.

JÄGER, Ludwig, Schauplätze der Evidenz. Evidenzverfahren und kulturelle Semantik. Eine Skizze, in: Cuntz, Michael/Nitsche, Barbara/Otto, Isabell et al. (eds.), *Die Listen der Evidenz*, Köln 2006, 37–52.

JAKOBSEN, Janus C./GLUUD, Christian/Wetterslev, Jørn et al., When and How Should Multiple Imputation Be Used for Handling Missing Data in Randomised Clinical Trials—A Practical Guide with Flowcharts, in: *BMC Medical Research Methodology* 17 (1/2017), 162. https://doi.org/10.1186/s12874-017-0442-1.

JATON, Florian, Assessing Biases, Relaxing Moralism. On Ground-Truthing Practices in Machine Learning Design and Application, in: *Big Data & Society* 8, (1/2021), 1–15. https://doi.org/10.1177/20539517211013569.

JENKINSON, Mark/CHAPPELL, Michael, *Introduction to Neuroimaging Analysis*, Oxford 2018.

JOG, Amod/CARRASS, Aaron/PRINCE, Jerry L., Self Super-Resolution for Magnetic Resonance Images, in: Ourselin, Sebastien /Joskowicz, Leo/Sabuncu, Mert R. et al. (eds.), *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2016*, Cham 2016, 553–560.

KANG, Seung K./SHIN, Seong A./SEO, Seongho et al., Deep Learning-Based 3D Inpainting of Brain MR Images, in: *Scientific Reports* 11 (6/2021), 1673. https://doi.org/10.1038/s41598-020-80930-w.

KOOI, Thijs, Deep Learning: From Natural to Medical Images, in: *Medium* December 19, 2018, URL: https://medium.com/merantix/deep-learning-from-natural-to-medical-images-74827bf51d6b [last accessed May 2, 2022].

LEHMANN, Thomas M./GONNER, Claudia/SPITZER, Klaus, Survey: Interpolation Methods in Medical Image Processing, in: *IEEE Transactions on Medical Imaging* 18 (11/1999), 1049–1075.

LERCH, Jason P./VAN DER KOUWE, André J./RAZNAHAN, Armin et al., Studying Neuroanatomy Using MRI, in: *Nature Neuroscience* 20 (3/2017), 314–326. https://doi.org/10.1038/nn.4501.

MANJÓN, José V/COUPÉ, Pierrick/BUADES Antonio et al., Non-Local MRI Upsampling, in: *Medical Image Analysis* 14 (6/2010), 784–792. https://doi.org/10.1016/j.media.2010.05.010.

MUHR, Paula, *From Photography to FMRI. Epistemic Functions of Images in Medical Research on Hysteria*, Bielefeld 2022. https://doi.org/10.14361/9783839461761.

PLENGE, Esben/POOT, Dirk H./NIESSEN, Wiro J. et al., Super-Resolution Reconstruction Using Cross-Scale Self-Similarity in Multi-Slice MRI, in: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 6* (2013), 123–130. https://doi.org/10.1007/978-3-642-40760-4_16.

REYNOLDS, Douglas, Gaussian Mixture Models, in: Li, Stan Z./Jain, Anil (eds.), *Encyclopedia of Biometrics*, Boston 2009.

VAN HORN, John D/TOGA, Arthur W., Human Neuroimaging As a 'Big Data' Science, in: *Brain Imaging and Behavior* 8 (2/2013). https://doi.org/10.1007/s11682-013-9255-y.

ZHAO, Can/DEWEY, Blake E./PHAM, Dzung L. et al., SMORE: A Self-Supervised Anti-Aliasing and Super-Resolution Algorithm for MRI Using Deep Learning, in: *IEEE Transactions on Medical Imaging* 40 (3/2021), 805–817. https://doi.org/10.1109/tmi.2020.3037187.

ZIEWITZ, Malte, Governing Algorithms. Myth, Mess, and Methods, in: *Science, Technology & Human Values* 41 (1/2016), 3–16. https://doi.org/10.1177/0162243915608948.

ZORAN, Daniel/WEISS, Yair, Natural Images, Gaussian Mixtures and Dead Leaves, in: *Proceedings of the 25th International Conference on Neural Information Processing Systems,* Red Hook 2012, 1736–1744.

Figures

Figure 1: Visualisation of sparsely sampled 2D MRI slices stacked into a 3D volume and the missing slices that need to be reconstructed through image imputation. From: Yaqiong Chai/Botian Xu/Kangning Zhang et al., MRI Restoration Using Edge-Guided Adversarial Learning, in: *IEEE Access* 8 (2020), 83859, fig. 1a, https://doi.org/10.1109/access.2020.2992204, cc-by-4.0.