Received: 22 April 2020 | Revised: 16 June 2020 | Accepted: 10 July 2020

DOI: 10.1002/mrm.28460

NOTE

Magnetic Resonance in Medicine

Chemical shift–based prospective k-space anonymization

Hendrik Mattern[1] | **Martin Knoll**[1] | **Falk Lüsebrink**[1,2] | **Oliver Speck**[1,3,4,5]

[1]Biomedical Magnetic Resonance, Otto-von-Guericke-University Magdeburg, Magdeburg, Germany

[2]Medicine and Digitalization, Department of Neurology, Otto-von-Guericke University Magdeburg, Magdeburg, Germany

[3]German Center for Neurodegenerative Disease, Magdeburg, Germany

[4]Center for Behavioral Brain Sciences, Magdeburg, Germany

[5]Leibniz Institute for Neurobiology, Magdeburg, Germany

Correspondence
Hendrik Mattern, Otto-von-Guericke University Magdeburg, Biomedical Magnetic Resonance, Leipziger Str. 44, H65-395, 39120 Magdeburg, Germany.
Email: hendrik.mattern@ovgu.de

Funding information
National Institutes of Health (Grant No. 1R01-DA021146) and the federal state of Saxony-Anhalt (Grant No. I 88)

Purpose: Publicly available data provision is an essential part of open science. However, open data can conflict with data privacy and data protection regulations. Head scans are particularly vulnerable because the subject's face can be reconstructed from the acquired images. Although defacing can impede subject identification in reconstructed images, this approach is not applicable to k-space raw data. To address this challenge and allow defacing of raw data for publication, we present chemical shift–based prospective k-space anonymization (CHARISMA).

Methods: In spin-warp imaging, fat shift occurs along the frequency-encoding direction. By placing an oil-filled mask onto the subject's face, the shifted fat signal can overlap with the face to deface k-space during the acquisition. The CHARISMA approach was tested for gradient-echo sequences in a single subject wearing the oil-filled mask at 7 T. Different fat shifts were compared by varying the readout bandwidth. Furthermore, intensity-based segmentation was used to test whether the images could be unmasked retrospectively.

Results: To impede subject identification after retrospective unmasking, the signal of face and shifted oil should overlap. In this single-subject study, a shift of 3.3 mm to 4.9 mm resulted in the most efficient masking. Independent of CHARISMA, long TEs induce signal decay and dephasing, which impeded unmasking.

Conclusion: To our best knowledge, CHARISMA is the first prospective k-space defacing approach. With proper fat-shift direction and amplitude, this easy-to-build, low-cost solution impaired subject identification in gradient-echo data considerably. Further sequences will be tested with CHARISMA in the future.

KEYWORDS
anonymization, data security, defacing, fat shift

1 | INTRODUCTION

Publicly available provision of data is one keystone of open science.[1,2] Such data enable replication of published results and can be used in the development of new techniques such as image reconstruction with deep learning.[3] However, open data can conflict with data privacy and data protection regulations.[2,4] Brain scans are particularly vulnerable because the subject's face can be reconstructed from the acquired images.[2] Therefore, brain images are commonly de-identified by removing metadata from the data header and cropping out or blurring the face in the images, the so-called defacing.[4,5] Although defacing is possible for reconstructed images, and large, anonymized, open-image repositories exist, such as the human connectome project,[6] defacing is not applicable to k-space raw data. To date, most publicly available raw data are based on knee scans[3,7] or 2D-encoded axial slices of the brain positioned above the eyes.[3] Note that publishing whole-head raw data without defacing is legally possible if the subject waives all rights.[8] However, this raw data handling is an exception, and subjects might not consent to such data-sharing policy.

In this study, we present chemical shift–based prospective k-space anonymization (CHARISMA) and test its masking capabilities for 3D-encoded gradient-echo imaging of the brain at 7 T. The CHARISMA approach leverages the chemical shift, which causes fat protons to precess at a different resonance frequency than water protons.[9,10] In spin-warp imaging, this results in the chemical shift–induced boundary artifact along the frequency-encoding direction (often called "fat shift" in short).[9,10] By placing an oil-filled mask onto the subject's face, the shifted fat signal can overlap with the subject's skin when proper sequence parameters are applied. Hence, the chemical shift, often considered an artifact, is exploited to deface the k-space data during the acquisition. In the following, a proof-of-principle study for CHARISMA is presented. Acquired data, reconstructed images, and the scripts used in this study are publicly available (see Data Availability Statement). Preliminary results of this study have been accepted as an abstract for the 2020 ISMRM.[11]

2 | METHODS

2.1 | Mask design

A commercially available, water-based cooling mask was refilled with colza oil (fat shift approximately 3.5 ppm[10,12]) through a small hole with a syringe/catheter tip. Afterward, the hole was sealed using an adhesive patch (bicycle equipment, no vulcanization). All building materials, the filling process, and final mask are shown in Figure 1. The fabrication of a single mask costs less than 15 Euros.

2.2 | Data acquisition

For a proof of principle, a single subject was scanned with a whole-body MAGNETOM classic 7T scanner (Siemens Healthineers, Erlangen, Germany) equipped with a 32-channel

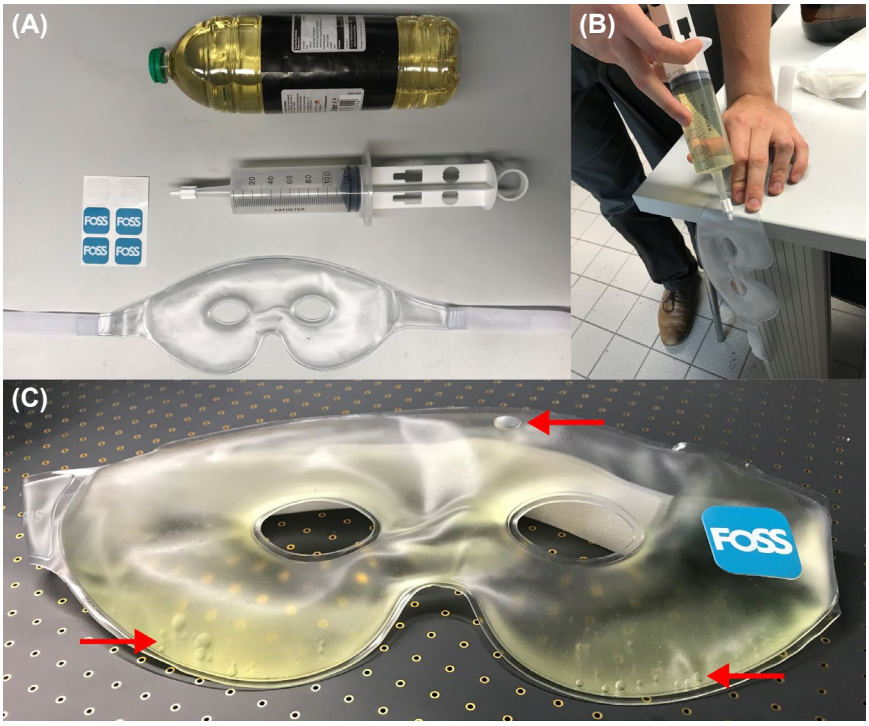

FIGURE 1 A, Required materials include a commercially available cooling mask, colza oil, syringe, and adhesive patch. B, After removing the water-based filling of the mask, the oil is injected through a small hole with a syringe. C, Finally, the hole is sealed with an adhesive patch. Red arrows indicate residual air bubbles after refilling. Fabrication of a single mask costs less than 15 Euros

head coil (Nova Medical, Wilmington, MA). The study was approved by the local ethics committee, and the subject gave written consent before scanning. Furthermore, the subject agreed in written form to the publication of all data acquired, including two volumes acquired without wearing the mask, and consented to the presentation of unmasked 3D face reconstructions in this manuscript. The subject is aware that the data published or the face surface reconstruction presented here may allow identification and does not object to the publication.

To test CHARISMA, a slab-selective gradient-echo sequence was used. At 7 T, this combination should be an ideal fit. Slab-selective gradient echoes are used regularly to acquire T_2^*-weighted data with axial orientation. Hence, the frequency direction is commonly set to anterior–posterior, as required by CHARISMA, and the long TEs used enable lower readout bandwidths, which result in larger fat shifts.[10] Furthermore, the apparent fat shift increases[10] and the T_2^*-contrast mechanism improves with the use of higher magnetic field strengths.[13] Axial slabs with 1-mm isotropic voxel size (256 × 224 × 96 mm FOV) and frequency-encoding direction set to anterior–posterior were acquired, and T_2^* weighting was achieved by setting TE/TR to 9/17 ms and the flip angle to 9°. Flow compensation in all three encoding directions was enabled. To compare the effect of different fat shifts, five protocols with different bandwidths were acquired: 980 Hz/voxel (1.0-mm shift), 390 Hz/voxel (2.5-mm shift), 300 Hz/voxel (3.3-mm shift), 200 Hz/voxel (4.9-mm shift), and 100 Hz/voxel (9.8-mm shift). Additionally, a volume without wearing the mask and a bandwidth of 100 Hz/voxel was acquired.

In addition to the T_2^*-weighting data, proton-density data were acquired by adapting the protocol accordingly (TE/TR = 3.4/8.1 ms and flip angle = 5°). To achieve the shorter TE, flow compensation was turned off. The shorter TE reduced signal decay and minimized susceptibility effects. Hence, the depiction of the subject's face, which is surrounded by intracranial and extracranial air, was improved. Proton density–weighted scans were acquired for two different bandwidths: 980 Hz/voxel (1.0-mm shift) and 200 Hz/voxel (4.9-mm shift). Additionally, a proton density–weighted scan without wearing the mask and a bandwidth of 980 Hz/voxel was acquired.

2.3 | Postprocessing

After bias field correction (using N4[14]), all images were coregistered rigidly (using advanced normalization tools[15]) to the proton density–weighted volume acquired without wearing the mask.

To evaluate the masking efficiency across fat shifts and contrasts, unmasking of all data sets was attempted retrospectively. To that end, ilastik's[16] pixel classification was used to segment the oil content of the mask (if visible), the water content of the mask, air, brain, and face in 3D volumes. The classification is performed semi-automatically. First, the user manually labels a subset of voxels for each class. Based on this sparse training data, ilastik estimates the probability of belonging to each of the classes for every voxel in the 3D data set.[16] This classification is based on the output of image filters (ie, voxel intensity, edgeness, or texture on multiple scales) as features and random forest as a classifier.[16] For unmasking, ilastik's segmentation probability maps were exported and loaded into *MATLAB* 2015b (The MathWorks, Natick, MA). The data were unmasked by subtracting the images weighted by the summed probability of the mask's oil and water content from the images themselves. The labeling, segmentation, and unmasking steps were repeated several times to find iteratively and empirically the best trade-off between oversegmentation and undersegmentation of the mask and face.

The masked and unmasked data were compared using 3D renderings of the face and intensity profiles extracted from 2D slices. For the 3D renderings, data acquired without wearing the mask are also shown to assess recognizability of facial features. For the intensity profiles, the data without the mask were not used, because wearing the mask squeezes the face. This nonrigid deformation impeded the extraction of comparable one-dimensional intensity profiles, but was rather indiscriminating for the human perception of facial features in 3D renderings.

3 | RESULTS

For an efficient masking, the signal of shifted fat and skin should overlap. To that end, CHARISMA was tested in a single subject for different fat shifts by varying the readout bandwidth and for two different contrasts. A retrospective, segmentation-based unmasking was used then to test whether the mask could be removed.

The T_2^*-weighted data sets were acquired for five different fat shifts. As shown in sagittal views and intensity profiles in Figure 2, the best overlap was achieved for 3.3-mm and 4.9-mm shift, with a slight advantage for 4.9 mm (with almost flat intensity profile after unmasking, see Figure 2M,N). After unmasking, the mask and, due to the overlap, face were removed from the images and intensity profiles. If the mask was shifted by 9.8 mm, its apparent location is partially beyond the subject's skin. Although in the one-dimensional profile unmasking removed the mask (see Figure 2O), further headward, unmasking caused signal discontinuities in the face (see Figure 2J). Shifts of 1.0 mm and 2.5 mm were not sufficiently large to overcome the physical distance between mask and skin entirely. Any imperfect overlap improved the unmasking's capability to separate the mask from the face, as shown in the intensity profiles. As a note, the mask's physical

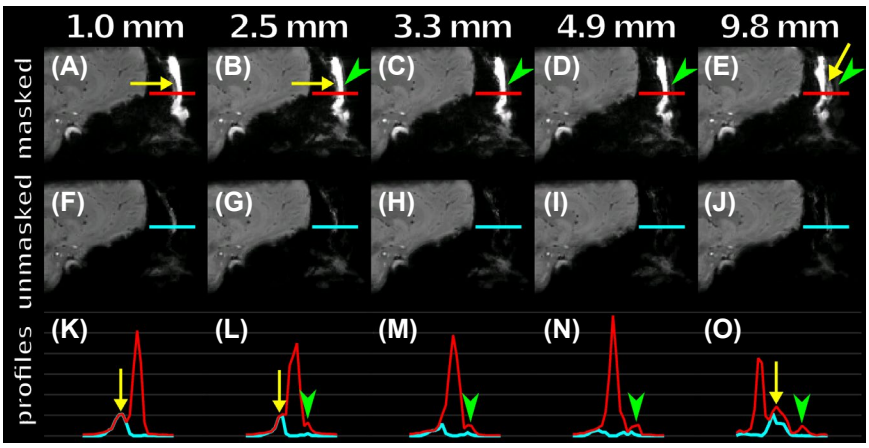

FIGURE 2 Sagittal views of the masked (A-E) and unmasked (F-J) T_2^*-weighted data for five different fat shifts. The corresponding intensity profiles are shown in (K)-(O) in red for prospectively masked, and cyan for retrospectively unmasked conditions, respectively. Red and cyan bars in (A)-(J) show the profiles' location. Yellow arrows show the subject's face, and green arrowheads show the water content of the oil-filled mask. Shifts of 3.3 mm (C,H,M) and 4.9 mm (D,I,N) resulted in the best overlap of mask and face, and therefore the least effective unmasking

location can be seen in the images as a shadow, and its signal originates from the water content of the oil used.

The data were also 3D-rendered, because a human observer would recognize a subject by his or her facial features. The masked and unmasked reconstructions are shown in Figure 3 alongside a volume acquired without wearing the mask. In line with the one-dimensional and 2D analysis, the unmasking efficiency was highest for shifts of 1.0 mm and 2.5 mm, but 3D reconstruction of facial features was overall poor. Unmasking worked best for 1.0-mm shift around the forehead. For this resolution and region, the manual labeling of mask and face, which was required for the segmentation, could be done with high confidence. Around the eyes and nose, this separation was nontrivial. There, the face's surface is more tortuous, causing varying mask–face distances, and therefore overlaps. Furthermore, this varying mask–face distance and small air bubbles inside the masks (see Figure 1C) caused locally varying susceptibility effects. The resulting residual mask fragments after unmasking caused a disruptive perception of the face.

To better delineate the face and mask, short-TE, proton density–weighted protocols with 1.0-mm and 4.9-mm fat shift were acquired (see Figure 4). In line with the T_2^*-weighted data, a shift of 1.0 mm is not sufficiently large to overlap the mask with the face completely. Although for a 4.9-mm shift the mask signal overlaps with the face, in this contrast, the skin's signal can shine through (see pink arrows in Figure 4B). This could impair the masking efficiency. Nevertheless, the naïve intensity-based segmentation was not able to distinguish between oil and face and caused signal dropouts in the face after unmasking (see Figure 4D). For the 1-mm shift, the unmasking removed almost exclusively the mask and not the face. For the 3D rendering in Figure 5, a data set without wearing the mask is shown for comparison. Note that the darker appearance of the mask in Figure 5C (4.9-mm shift) is caused by the residual water content of the oil positioned anterior to the face and oil signal. Because of the insufficient overlap of oil and face for 1.0-mm shift, unmasking could remove the mask to a large extent, enabling subject identification. For 4.9-mm shift, the unmasking was less efficient, and large fragments of the masks are still present after unmasking. Thus, subject identification was complicated by CHARISMA considerably.

4 | DISCUSSION

In this study, we presented CHARISMA, a low-cost solution to deface raw data prospectively. By adjusting the sequence parameters, the fat shift can be used to overlay the mask and face signal, complicating subject identification considerably. As a proof of principle, CHARISMA was tested successfully for 3D gradient echoes with frequency encoding in the anterior–posterior direction. Further sequences need to be tested with CHARISMA to enable reliable, sequence-specific recommendations. Theoretically, CHARISMA should be applicable to any sequence meeting the following three requirements, but there are challenges that limit its practical application.

First, the fat shift direction has to be in the anterior–posterior direction to shift the oil onto the subject's face. This is the case for spin warp imaging sequences such as gradient and (turbo) spin echoes with the frequency encoding set to anterior–posterior or EPI with anterior–posterior-oriented phase encoding. Note that this theoretical range of application requires thorough testing before reliable recommendations can be given. Hence, CHARISMA is not applicable to protocols such as 3D-encoded, whole-brain acquisitions or 2D-encoded, sagittal multislice imaging with frequency encoding in the head–foot direction.

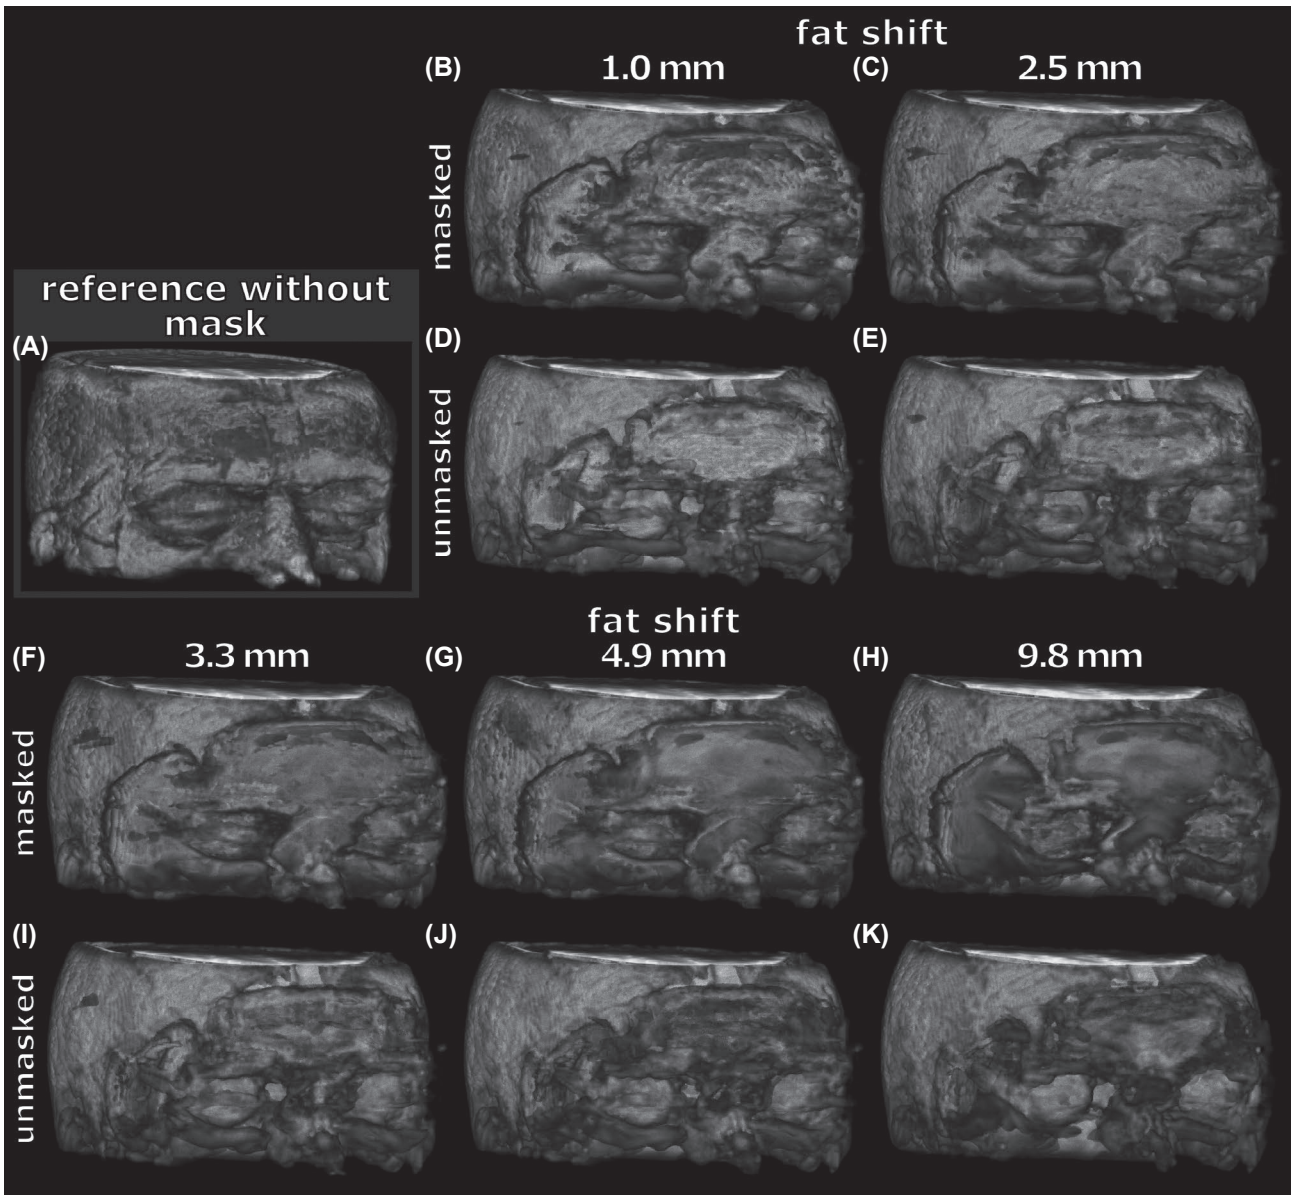

FIGURE 3 Three-dimensional rendering of the masked (B,C,F-H) and unmasked (D,E,I-K) T_2^*-weighted data for five different fat shifts. A, For comparison, a volume without wearing the mask is shown. The naïve, retrospective, intensity-based unmasking removed the mask only partially, and subject identification remained impeded by the residual mask

Second, the shifted oil needs to overlap with the subject's face. The fat shift amplitude increases with the magnetic field strength. Therefore, the integration of CHARISMA should be easier at ultrahigh field compared with lower field strengths. As the fat shift amplitude is a function of the imaging resolution and readout bandwidth used, the protocol needs to be adapted to CHARISMA. Commonly, sequence protocols are optimized for a specific goal. These optimized protocols might prevent the application of CHARISMA at all or would require sequence adaptation. Based on this study, gradient-echo protocols with long TEs or low readout bandwidth should be compatible with CHARISMA (as discussed later). For other sequences or weightings, protocols could be optimized within the constraints given by CHARISMA. This could reduce the sequence's contrast or image quality compared to the case without CHARISMA applied, but would enable prospective k-space defacing, and thus impair subject identification. Alternatively, the mask filling could be changed to a material with another resonance frequency to adapt CHARISMA rather than the protocol for a given imaging scenario.

Third, no fat suppression or fat–water imaging[12,17] is applied in the sequence, as it would jeopardize CHARISMA's masking.

Hence, like the fat shift itself, CHARISMA is not bound to a single contrast mechanism, but its requirements limit the range of possible use cases in practice. Regarding the axial gradient-echo sequences tested in this study, CHARISMA

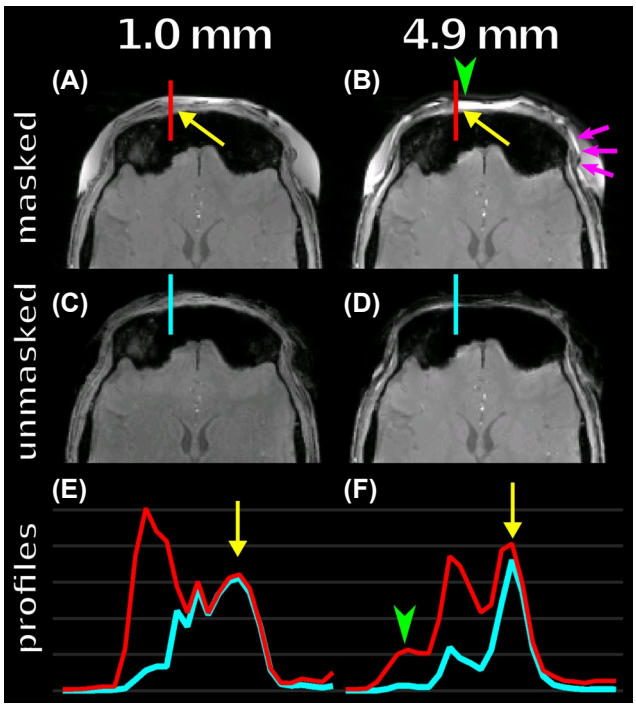

FIGURE 4 Axial views of the masked (A,B) and unmasked (C,D) proton density–weighted data for two different fat shifts. The corresponding intensity profiles are shown in (E) and (F) in red for prospectively masked, and cyan for retrospectively unmasked, conditions, respectively. The red and cyan bars in (A)-(D) show the profiles' location. Yellow arrows indicate the subject's face, and green arrowheads indicate the water content of the oil-filled mask. In contrast to the 1.0-mm shift, for the 4.9-mm mask and face overlap, the face partially shines through (pink arrows)

was applicable to the T_2^*-weighted protocols without any sequence adaptation. For the proton density–weighted protocols, no flow compensation was applied to enable short TEs, even if low readout bandwidths, and thus large fat shifts, were used.

Regardless of the contrast used, in this single-subject study, a fat shift amplitude of 3.3 mm to 4.9 mm achieved the best overlap of mask and face signal. For these shifts, the face's surface is fully covered by the thickness of the mask, but further tests with multiple subjects are required because the optimal shift amplitude might depend on the subject, mask geometry, mask fixation, and mask-filling level.

Although the fat shift is not a function of the image weighting, in this study, considerable differences in retrospective unmasking were observed between T_2^*-weighted and proton density–weighted data. Long TEs cause signal decay and increase susceptibility weighting, which complicated labeling and separation of face and mask in the segmentation-based unmasking. More advanced unmasking approaches might compensate for this T_2^*-weighting associated challenge. Furthermore, such techniques could improve unmasking in general, but the development of such a deanonymization technique requires considerable effort and was beyond the scope of this study. Moreover, gradient-echo sequences with long TEs could be defaced further by leveraging susceptibility-induced signal loss. In this study, signal inhomogeneities induced by air bubbles and varying mask–face distance prevented reliable unmasking incidentally. By adding paramagnetic particles to the mask's skin

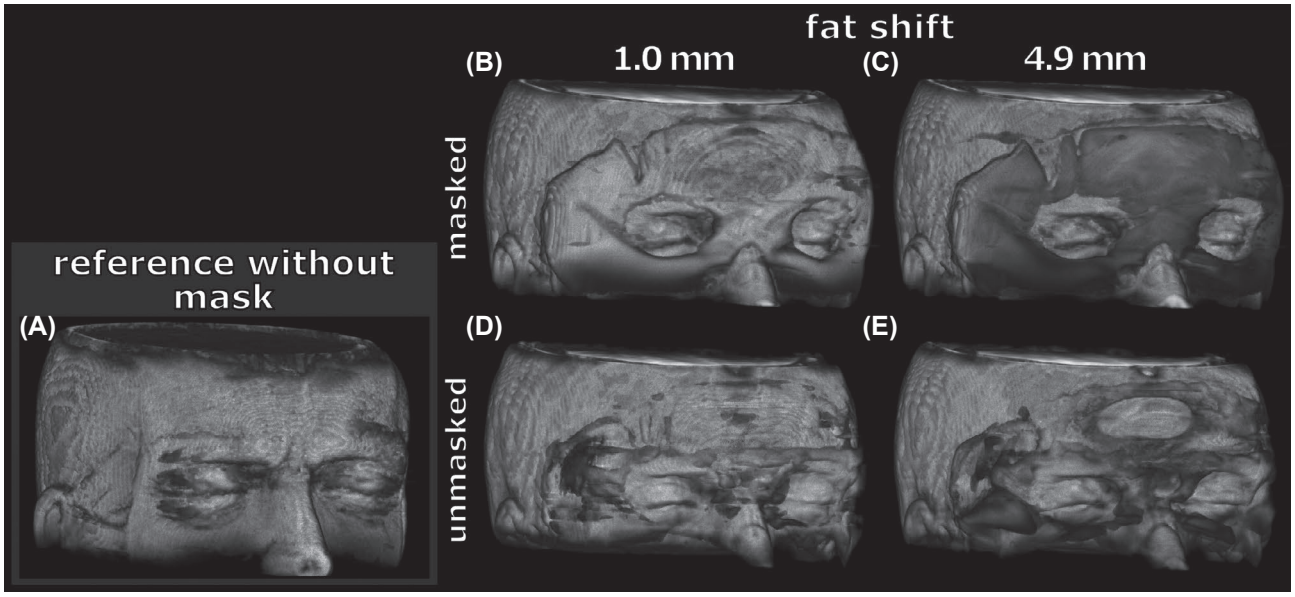

FIGURE 5 Three-dimensional rendering of the proton density–weighted data (A) without wearing the mask (B,C), with the mask, and after unmasking the data retrospectively (D,E). Due to the missing overlap of mask and face for the 1.0-mm fat shift, the segmentation-based unmasking could remove most of the mask. For the 4.9-mm shift, the face and mask overlap resulted in less effective unmasking

side deliberately, susceptibility-induced signal loss could be used as a complementary defacing strategy to the fat shift presented here.

Like any security mechanism, CHARISMA provides no absolute security. Although the mask could be extended to conceal also the eyes, nose and mouth, with the current mask, these structures are not covered to allow the use of visual stimuli and improve subject comfort. Nevertheless, eyes, nose, and mouth are facial features that are particularly relevant for subject identification.[4] Therefore, mask design should trade off subject comfort and data security, and should be done in consultation with the local ethics committee. Regardless of whether CHARISMA is used, subjects should be informed about the risks associated with publicly available data. Future technological advances are a potential threat to the subject's data privacy. Deep learning–based inpainting, which could undo retrospective image defacing to some extent,[18] could be applied to CHARISMA. Furthermore, access to privileged data repositories could uncover the subject's identity by matching brain images or dental records.

Despite these challenges, CHARISMA is, to our best knowledge, the first prospective anonymization strategy for k-space raw data. Future alternatives could be based on spatially selective electro-magnetic fields to dephase or saturate the magnetization, prevent excitation, or not to detect the signal in the facial region. Potentially, technologies such as parallel transmission, parallel imaging, nonlinear gradients, or saturation pulses can be applied for these purposes. Although these techniques may be more effective and provide general solutions, they are associated with considerable costs and efforts compared with the low-cost, do-it-yourself CHARISMA approach, and remain topics for future research. Potentially, CHARISMA's partial solution to the k-space anonymization challenge could stimulate the development of more general defacing techniques to address the demand by more and more funding bodies, to publish all data acquired openly in accordance with data security regulations.

5 | CONCLUSION

To our best knowledge, CHARISMA is the first approach to anonymize k-space raw data prospectively. With proper fat shift direction and amplitude, this easy-to-build, low-cost solution complicates subject identification considerably. At 7 T, CHARISMA and T_2^*-weighted imaging are an ideal fit.

Eventually, this work could foster the discussion and development of prospective k-space anonymization to enable open (raw) data for brain examinations. Publicly available data repositories could then be used for the development of advanced reconstruction methods or to improve reproducibility in MRI research.

ACKNOWLEDGMENT

Open access funding enabled and organized by Projekt DEAL.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are openly available in GitLab at https://gitlab.com/hmattern/tmi/-/tree/master/CHARISMA (commit number 5f935846; SHA-1 hash 5f935846f45e1b38ba778765c1822339ee1d323a). For the validation of the presented mask-based defacing, a comparison to the ground truth is required. Therefore, the data acquired without wearing the mask were included in this data repository. Although this procedure does not follow the recommendations given in the MRM author guidelines (version February 2020), it enables us to reproduce the results of, to our best knowledge, the first k-space defacing approach. Furthermore, the subject consented in written form to this publication.

ORCID

Hendrik Mattern https://orcid.org/0000-0001-5740-4522

REFERENCES

1. Stikov N, Trzasko JD, Bernstein MA. Reproducibility and the future of MRI research. *Magn Reson Med*. 2019;82:1981-1983.
2. Poldrack RA, Gorgolewski KJ. Making big data open: data sharing in neuroimaging. *Nat Neurosci*. 2014;17:1510-1517.
3. Zbontar J, Knoll F, Sriram A, et al. fastMRI: an open dataset and benchmarks for accelerated MRI. 2018. ArXiv:1811.08839.
4. Bischoff-Grethe A, Ozyurt IB, Busa E, et al. A technique for the deidentification of structural brain MR images. *Hum Brain Mapp*. 2007;28:892-903.
5. Milchenko M, Marcus D. Obscuring surface anatomy in volumetric imaging data. *Neuroinform*. 2013;11:65-75.
6. van Essen DC, Smith SM, Barch DM.Behrens TEJ, Yacoub E, Ugurbil K. The WU-Minn Human Connectome Project: an overview. *Neuroimage*. 2013;80:62-79.
7. Ong F, Amin S, Vasanawala S, Lustig M. mridata.org: an open archive for sharing MRI raw data. In: *Proceedings of the Joint Annual Meeting of ISMRM-ESMRMB*, Paris, France, 2018. Abstract #3425.
8. Lüsebrink F, Sciarra A, Mattern H, Yakupov R, Speck O. T1-weighted in vivo human whole brain MRI dataset with an ultrahigh isotropic resolution of 250 μm. *Sci Data*. 2017;4:170032.
9. Babcock EE, Brateman L, Weinreb JC, Horner SD, Nunnally RL. Edge artifacts in MR images: chemical shift effect. *J Comput Assist Tomogr*. 1985;9:252-257.
10. Parizel PM, van Hasselt BA, van den Hauwe L, van Goethem JW, de Schepper AM. Understanding chemical shift induced boundary artefacts as a function of field strength: influence of imaging parameters (bandwidth, field-of-view, and matrix size). *Eur J Radiol*. 1994;18:158-164.
11. Mattern H, Knoll M, Lüsebrink F, Speck O. Chemical sHift bAsed pRospectIve k-Space anonyMizAtion (CHARISMA). In: *Proceedings of the Virtual Meeting of ISMRM & SMRT*, 2020. Abstract number #1049
12. Bley TA, Wieben O, François CJ, Brittain JH, Reeder SB. Fat and water magnetic resonance imaging. *J Magn Reson Imaging*. 2010;31:4-18.

13. Ladd ME, Bachert P, Meyerspeer M, et al. Pros and cons of ultra-high-field MRI/MRS for human application. *Prog Nucl Magn Reson Spectrosc*. 2018;109:1-50.
14. Tustison NJ, Avants BB, Cook PA, et al. N4ITK: improved N3 bias correction. *IEEE Trans Med Imaging*. 2010;29:1310-1320.
15. Avants B, Epstein C, Grossman M, Gee J. Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Med Image Anal*. 2008;12:26-41.
16. Berg S, Kutra D, Kroeger T, et al. ilastik: interactive machine learning for (bio)image analysis. *Nat Methods*. 2019;16:1226-1232.
17. Berglund J, Rydén H, Avventi E, Norbeck O, Sprenger T, Skare S. Fat/water separation in k-space with real-valued estimates and its combination with POCS. *Magn Reson Med*. 2020;83:653-661.
18. Abramian D, Eklund A. Refacing reconstructing anonymized facial features using GANS. In: *Proceedings of the 2019 IEEE International Symposium on Biomedical Imaging*, Venice, Italy, 2019. pp. 1104-1108.

How to cite this article: Mattern H, Knoll M, Lüsebrink F, Speck O. Chemical shift–based prospective k-space anonymization. *Magn Reson Med*. 2021;85:962–969. https://doi.org/10.1002/mrm.28460