



King's Research Portal

DOI:

[10.1109/TPAMI.2023.3346330](https://doi.org/10.1109/TPAMI.2023.3346330)

Document Version

Early version, also known as pre-print

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Fidon, L., Aertsen, M., Kofler, F., Bink, A., David, A. L., Deprest, T., Emam, D., Guffens, F., Jakab, A., Kasprian, G., Kienast, P., Melbourne, A., Menze, B., Mufti, N., Pogledic, I., Prayer, D., Stuempflen, M., Elslander, E. V., Ourselin, S., ... Vercauteren, T. (2024). A Dempster-Shafer Approach to Trustworthy AI With Application to Fetal Brain MRI Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(5), 3784-3795. Article 10388220. <https://doi.org/10.1109/TPAMI.2023.3346330>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

A Dempster-Shafer approach to trustworthy AI with application to fetal brain MRI segmentation

Lucas Fidon¹, Michael Aertsen², Florian Kofler^{3,9,10}, Andrea Bink⁴,
Anna L. David⁵, Thomas Deprest², Doaa Emam^{2,6}, Frédéric
Guffens², András Jakab⁷, Gregor Kasprian⁸, Patric Kienast⁸,
Andrew Melbourne¹, Bjoern Menze^{11,3}, Nada Mufti¹, Ivana
Pogledic⁸, Daniela Prayer⁸, Marlene Stuempflen⁸, Esther Van
Elslander², Sébastien Ourselin¹, Jan Deprest², and Tom
Vercauteren¹

¹School of Biomedical Engineering & Imaging Sciences, King's
College London, London, UK

²Department of Radiology, University Hospitals Leuven, Leuven,
Belgium

³Department of Informatics, Technical University Munich, Munich,
Germany

⁴ Department of Neuroradiology and Clinical Neuroscience Center,
University Hospital Zurich and University of Zurich, Zurich,
Switzerland

⁵Institute for Women's Health, University College London,
London, UK

⁶Department of Gynecology and Obstetrics, University Hospitals
Tanta, Tanta, Egypt

⁷Center for MR Research, University Children's Hospital Zurich,
University of Zurich, Zurich, Switzerland

⁸Department of Biomedical Imaging and Image-guided Therapy,
Medical University of Vienna, Vienna, Austria

⁹ Department of Diagnostic and Interventional Neuroradiology,
School of Medicine, Klinikum rechts der Isar, Technical University
of Munich, Germany

¹⁰TranslaTUM - Central Institute for Translational Cancer
Research, Technical University of Munich, Germany

¹¹Department of Quantitative Biomedicine, University of Zurich,
Zurich, Switzerland

Abstract

Deep learning models for medical image segmentation can fail unexpectedly and spectacularly for pathological cases and for images acquired at different centers than those used for training, with labeling errors that violate expert knowledge about the anatomy and the intensity distribution of the regions to be segmented. Such errors undermine the trustworthiness of deep learning models developed for medical image segmentation. Mechanisms with a fallback method for detecting and correcting such failures are essential for safely translating this technology into clinics and are likely to be a requirement of future regulations on artificial intelligence (AI). Here, we propose a principled trustworthy AI theoretical framework and a practical system that can augment any backbone AI system using a fallback method and a fail-safe mechanism based on Dempster-Shafer theory. Our approach relies on an actionable definition of trustworthy AI. Our method automatically discards the voxel-level labeling predicted by the backbone AI that are likely to violate expert knowledge and relies on a fallback atlas-based segmentation method for those voxels. We demonstrate the effectiveness of the proposed trustworthy AI approach on the largest reported annotated dataset of fetal T2w MRI consisting of 540 manually annotated fetal brain 3D MRIs with neurotypical or abnormal brain development and acquired from 13 sources of data across 6 countries. We show that our trustworthy AI method improves the robustness of a state-of-the-art backbone AI for fetal brain MRI segmentation on MRIs acquired across various centers and for fetuses with various brain abnormalities.

Automatic segmentation of medical images is needed for personalized medicine and to study anatomical development in healthy populations as well as populations with a pathology. Artificial Intelligence (AI) algorithms for medical image segmentation can reach super-human accuracy on average [28] and yet most radiologists do not trust them [3, 5]. This is partly because, for some cases, AI algorithms fail spectacularly with errors that violate expert knowledge about the segmentation task when the AI was applied across imaging protocol and anatomical pathologies [3, 17, 22] (Fig.1b). This sense of distrust is exacerbated by the current lack of clear fit-for-purpose regulatory requirements for AI-based medical image software [34].

The legal framework for the deployment in clinics of AI tools for medical segmentation is likely to soon become more stringent once the European Union has proposed its Artificial Intelligence Act to regulate AI and AI trust is at the core of this proposal [12, 13]. Trust is an attitude of the human user and, in particular, AI trust is context-dependent. In contrast, AI trustworthiness is an attribute of the AI system [29, 25]. It is worth noting that a human user can trust an AI system that is not trustworthy and conversely. Guidelines for trustworthy AI claim that AI trustworthiness must precede trust in the deployment of AI systems to avoid miscalibration of the human trust with respect to the trustworthiness of an AI system [12]. In Psychology, trust of humans in AI can be defined as the belief of the human user that the AI system will satisfy the criteria of a set of contracts of trust. This contract-based definition of human-AI trust reflects the plurality and the context-dependency of human-AI trust. In particular, the user may trust an AI system for one population or one type of scanner but not trust it for others. An AI system is trustworthy to a contract of trust if it can maintain this contract in all situations within the contract scope.

The EU ethics guidelines for trustworthy AI, that upheld the AI Act, advocate that “AI systems should have safeguards that enable a *fallback* plan in case of problems” [12]. We argue that those safeguards should implement a *fail-safe* mechanism in relation with a collection of contracts of trust so as to improve the trustworthiness of the overall system.

In this article, we propose the first trustworthy AI framework with a fail-safe and a fallback for medical image segmentation (see A Dempster-Shafer approach to Trustworthy AI for medical image segmentation). The proposed framework consists of three main components: first, a backbone AI algorithm, that can be any AI algorithm for the task at hand, second, a fallback segmentation algorithm, that can be any safe segmentation algorithm but potentially less precise than the backbone AI algorithm, and third, a fail-safe method that automatically detects local conflicts between the backbone AI algorithm prediction and the contracts of trust and switches to the fallback algorithm in case of conflicts. This is illustrated for fetal brain 3D MRI segmentation in Fig. 1. The proposed principled fail-safe method is based on Dempster-Shafer (DS) theory. DS theory allows to model partial information about the task, which is typically the case for expert knowledge. For example, in human brain anatomy, one piece of expert knowledge is that the cerebellum is located in the lower back part of the brain. This gives us information only about the segmentation of the cerebellum while a segmentation algorithms will typically compute segmentation for many other tissue types in addition to the cerebellum. The Dempster’s rule of combination of DS theory is an efficient mathematical tool to combine independent sources of information that discards contradictions among the sources (see Background on Dempster-Shafer theory). In our framework, the AI-based segmentation and each expert knowledge are treated as independent sources of information and the Dempster’s rule of combination is employed to act as the fail-safe (see Fail-safe mechanism: switching between the AI algorithm and the fallback).

To demonstrate the applicability of the developed trustworthy AI framework, we propose one implementation for fetal brain segmentation in MRI. For the backbone AI model, we used the state-of-the-art deep learning-based segmentation pipeline nnU-Net [28] (see nnU-Net as backbone AI segmentation algorithm). For the fallback model we used a registration-based segmentation method inspired by the state-of-the-art multi-atlas method GIF [6] (see Multi-atlas segmentation as fallback segmentation algorithm). We also show that our fail-safe formulation is flexible enough to model both spatial location-based (see Dempster-Shafer anatomical contract of trust for fetal brain segmentation) and intensity-based (see Dempster-Shafer intensity-based contract of trust for fetal brain segmentation) contracts of trust about the regions of interest to be segmented (see General case with multiple contracts of trust:). Spatial location-based fail-safes are implemented using the masks of the regions of interest computed by the registration-based fallback algorithm. However, in the fail-safe, the masks are interpreted differently. In this case, the mask of a region R is used only to exclude labeling voxels outside of the mask as belonging to R. This is illustrated in Fig. 1c. Inspired by the margins used for safety in radiation therapy planning to account for spatial registration errors [41], we first add spatial margins to the fallback masks before excluding the labels seen as anatomically unlikely according to the dilated fallback masks. While allowing the masks to overlap, this helps preventing miscoverage of the regions of interest that is the

only source of error in this formulation of the fail-safe.

We evaluated the proposed trustworthy AI method on fetal brain segmentation into eight tissue types using 3D T2w MRI. The segmentation of fetal brain MRI is essential to study normal and abnormal fetal brain development. In the future, reliable analysis and evaluation of fetal brain structures could also support diagnosis of central nervous system pathology, patient selection for fetal surgery, evaluation and prediction of outcome, hence also parental counselling. In particular, fetal brain 3D T2w MRI segmentation presents multiple challenges for trustworthy AI [17]. There are variations in T2w MRI protocols used across clinical centers and there is a spectacular variation of the fetal brain anatomy across gestational ages and across normal and abnormal fetal brain anatomy.

Results

Evaluation on a large multi-center fetal brain T2w MRI dataset. To effectively evaluate the performance of the proposed trustworthy AI framework as a suitable method to improve the trustworthiness of a backbone AI model using a fallback model, we have selected the task of fetal brain segmentation in 3D MRI. This task is clinically relevant and is characterized by large image protocol variations and large anatomical variations.

Deep learning-based AI methods for fetal brain MRI segmentation have recently defined state-of-the-art segmentation performance [15, 16, 18, 26, 31, 35, 43, 55], gradually replacing image registration-based segmentation methods [36] in the literature. Most previous work on deep learning for fetal brain MRI segmentation trained and evaluated their models using only MRIs of healthy fetuses or only MRIs acquired at one center. However, the segmentation performance of deep learning methods typically degrades when images from a different center or a different scanner vendor as the one used for training are used or when evaluating the segmentation performance on abnormal anatomy [2, 30, 37, 42, 44, 47]. One study has reported such issues for fetal brain MRI segmentation [17].

Thus, we have compared the proposed trustworthy AI algorithm to the backbone AI algorithm only [28] and the fallback algorithm only, consisting of a registration-based segmentation method [6]. We have used a large multi-centric fetal brain 3D MRI dataset that consists of a total of 540 fetal brain 3D MRIs with neurotypical or abnormal brain development, with gestational ages ranging from 19 weeks to 40 weeks, and with MRIs acquired at 13 hospitals across 6 countries. The task consists of segmenting automatically a fetal brain 3D MRI into eight clinically relevant tissue types: the corpus callosum, the white matter, the cortical gray matter, the deep gray matter, the cerebellum, the brainstem, the intra-axial cerebrospinal fluid (CSF), and the extra-axial CSF.

Stratified evaluation across brain conditions and acquisition centers.

The evaluation of AI-based segmentation algorithms has shown that the performance of deep learning models can vary widely across clinically relevant populations and across data acquisition protocols [17] (Fig. 1b, Fig. 2).

Therefore, we performed a stratified comparison of the backbone AI algorithm, the fallback algorithm, and the trustworthy AI algorithm across two

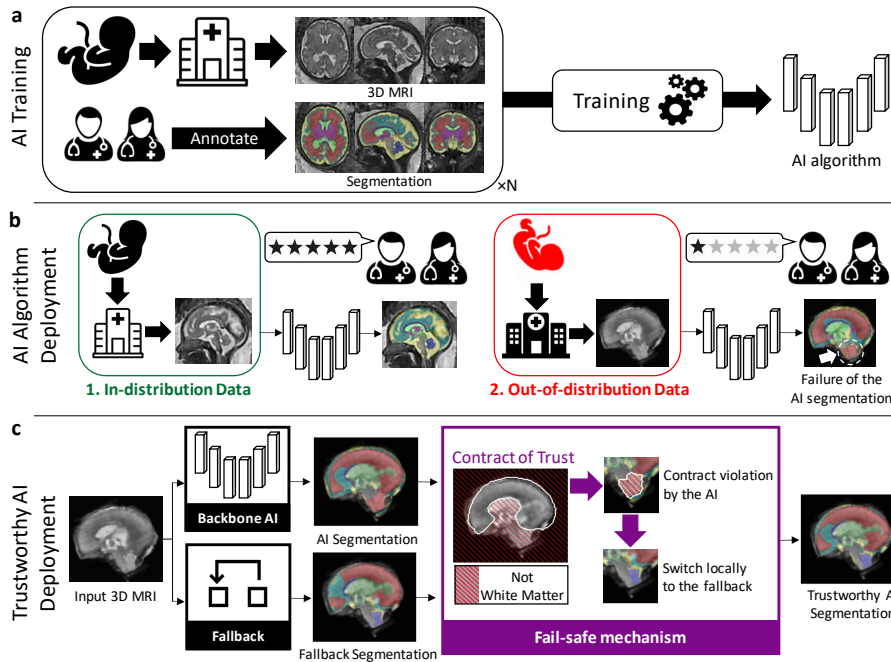


Figure 1: **Schematics of our principled method for trustworthy AI applied to medical image segmentation.** **a.** Deep neural networks for medical image segmentation (AI algorithm) are typically trained on images from a limited number acquisition centers. This is usually not sufficient to cover all the anatomical variability. **b.** When the trained AI algorithm is deployed, it will typically give satisfactory accuracy for images acquired with the same protocol as training images and with a health condition represented in the training dataset (left). However, the AI algorithm might fail with errors that are not anatomically plausible, for images acquired with a slightly different protocol as training images and/or representing anatomy underrepresented in the training dataset (right). **c.** Schematic of the proposed trustworthy AI algorithm. The backbone AI segmentation algorithm is coupled with a *fallback* segmentation algorithm. Experts knowledge about the anatomy is modelled using Dempster-Shafer theory. When part of the AI segmentation is found to contradict expert knowledge for a voxel, our trustworthy AI algorithm automatically switches to the fallback segmentation for this voxel.




Figure 2: Illustration of the improved robustness of the proposed trustworthy AI method (TW-AI) as compared to a state-of-the-art AI method. (Top) 3D MRI of a neurotypical fetus at 28 weeks of gestation acquired at the same center as the training data for the AI. (Bottom) 3D MRI of a fetus with a high-flow dural sinus malformation at 28 weeks of gestation acquired at a different center as the training data for the AI. Severe violations of the anatomy by the AI, but not the TW-AI, are highlighted in white.

groups of acquisition centers and three groups of brain conditions. The composition of the dataset for each group is summarized in Fig. 4 and detailed in Fetal brain MRI dataset. The acquisition centers were split into two groups, that we called *in-scanner distribution* and *out-of-scanner distribution*, depending if 3D MRIs acquired at a given center were present in the training dataset or not. Four out of thirteen data sources were used for the training of the backbone AI algorithm. In addition, the 3D MRIs were also separated based on the underlying brain condition of the fetus. The first group, *neurotypical*, contains the fetuses diagnosed by radiologists with a normal brain development using ultrasound and MRI. The second group, *spina bifida*, contains the fetuses with a condition called spina bifida aperta. We use the term *spina bifida* for short in this work. Cases of spina bifida aperta are typically accompanied by severe anatomical brain abnormalities [20, 45] with a type II Chiari malformation and an enlargement of the ventricles being most prevalent. The Chiari malformation type II is characterized by a small posterior fossa and hindbrain herniation in which the medulla, cerebellum, and fourth ventricle are displaced caudally into the direction of the spinal canal [40]. The third group, *other pathologies*, contains fetuses with various pathologies other than spina bifida and causing an abnormal brain development, such as corpus callosum agenesis and dysgenesis, intracranial hemorrhage and cyst, aqueductal stenosis, and Dandy-Walker malformation. Those other pathologies were not present in the training dataset of the backbone AI algorithm and spatio-temporal atlases are not available for the fallback and the fail-safe algorithm. Hence, the testing 3D MRIs classified as other pathologies allow us to measure the segmentation performance of

the trustworthy AI algorithm outside of the domain covered by the anatomical contracts of trust.

Figures 3a,3b show the results of the overall stratified evaluation in terms of Dice score and Hausdorff distances at 95% percentile. The detailed results per region of interest can be found in the appendix (Fig. 7,8).

Scoring of trustworthiness by a panel of radiologists. The Dice score and the Hausdorff distance are the two most standard metrics used for measuring the quality of automatic segmentations [4]. However, those two metrics do not directly measure the trustworthiness of segmentation algorithms [33]. Therefore, we have also conducted an evaluation of the trustworthiness of the automatic segmentations as perceived by radiologists. We have asked a panel of 8 experts from 4 different hospitals to score the trustworthiness of automatic segmentations from 0 (totally unacceptable) to 5 (perfect fit) for each region of interest and for the three segmentation algorithms. Independent scoring were performed by raters at different hospitals. The scoring protocol and details about the panel of experts can be found in section Human expert scoring method for evaluating the trustworthiness of fetal brain 3D MRI segmentation. The scoring was performed for the same 3D MRIs by all radiologists. We have used a subset of 50 3D MRIs from the out-of-distribution group of the testing dataset consisting of 20 neurotypical fetuses, 20 spina bifida fetuses, and 10 fetuses with other abnormalities. Those cases were selected per condition at random among the 3D MRIs of the publicly available FeTA dataset [43]. We have chosen to restrict this analysis to fetal brain 3D MRIs from the FeTA dataset, a subset of the *out-of-scanner distribution* group, because those data are publicly available which facilitated collaborations across centers and because rating of the quality of the 3D MRIs had been performed in previous work [43]. The *out-of-scanner distribution* group is the most relevant group for the evaluation of trustworthiness because this corresponds to the situation in which AI algorithms generalization is the most challenging and clinically relevant. In addition, this allows us to share our segmentations and scores publicly, thereby improving the reproducibility of our results. The overall scoring results can be found in Fig. 3e and the detailed results per region of interest can be found in the appendix (Fig. 11).


Expert raters reported that most of the 3D MRIs were of good quality. However, they noticed that the algorithms were very dependent on the quality of the 3D MRIs they were based on for the *spina bifida* group. We found a positive correlation between the mean-class trustworthiness scores and the quality of the 3D MRI for the *spina bifida* group (Pearson $r = 0.43$). There was no correlation between scores and 3D MRI quality for the *neurotypical* group (Pearson $r = -0.1$) and the 3D MRIs of the *other pathologies* were all of high quality. In addition, the more structurally abnormal the brains were due to the pathologies, the more difficult it was to compare the algorithms. In the case of the Chiari malformations, this applies in particular to the cerebellum and brainstem.

Stratified evaluation across gestational ages. The anatomy and the size of the fetal brain change significantly from 19 weeks of gestation until term for both neurotypical fetuses [21] and fetuses with spina bifida [20]. This age-related

variability is a challenge for segmentation algorithms for fetal brain MRI [17].


We analysed the performance of the proposed trustworthy AI algorithm for fetal brain segmentation as a function of the gestational age and compared it to the backbone AI algorithm based on deep learning [28] and the fallback algorithm based on image registration [6]. We grouped the fetuses with neurotypical or spina bifida condition with the same gestational age rounded to the closest week. The mean and the confidence intervals at 95% for the overall performance in terms of Dice score (resp. Hausdorff distance) across regions of interest can be found in Fig. 3c (resp. Fig. 3d). The detailed results per region of interest can be found in the appendix (Fig. 9,10).

Our results show how the backbone-AI algorithm, based on deep learning, and the fallback algorithm, based on image registration, are complementary. Overall, the backbone-AI algorithm achieves higher Dice scores than the fallback algorithm, while the fallback achieved lower Hausdorff distances than the backbone-AI method (Fig. 3). Our proposed trustworthy AI algorithm successfully combines the backbone AI and the fallback algorithms. It achieves higher or similar segmentation performance than those two algorithms in terms of established segmentation quality metrics such as the Dice score and the Hausdorff distance across all gestational ages for *neurotypical* and *spina bifida*. Using a Wilcoxon signed-rank test, we found that the trustworthy AI method significantly outperforms the backbone AI in terms of both Dice score and Hausdorff distance for every group ($p < 0.05$), except for the Dice score of the in-scanner distribution *other pathologies* group for which there is no statistical difference between the trustworthy AI and AI methods. The trustworthy AI method also significantly outperforms the fallback for the two metrics for all groups, except for the Hausdorff distances of the out-of-scanner-distribution *spina bifida* group for which there is no statistical difference between the trustworthy AI and fallback methods. This result is further supported by the trustworthiness scores of expert paediatric radiologists specialized in fetal brain anatomy. The trustworthy AI method outperformed the AI and performed on a par with the fallback for the *neurotypical* and *spina bifida* groups and the trustworthy AI method outperformed the fallback and performed on a par with the AI method for the *other pathologies* group.




(a) Mean-ROI Dice Score (in %)

(b) Mean-ROI HD95 (in mm)



(c) Mean and 95% CI of Dice Scores

(d) Mean and 95% CI of HD95



(e) Mean-ROI Trustworthiness Scores for out-of-scanner distribution 3D MRIs

Figure 3: **Comparison of the backbone AI, fallback, and trustworthy AI segmentation algorithms.** In (a) (resp. (b)), we report for each algorithm the distributions of Dice scores (resp. Hausdorff distances at 95% percentile (HD95)) for in-scanner distribution ($n = 182$) and out-of-scanner distribution ($n = 167$) images and for neurotypical ($n = 141$), spina bifida ($n = 157$), and other pathologies ($n = 51$). In (c) (resp. (d)), we report mean and 95% confidence intervals for the Dice scores (resp. HD95) across gestational ages, for neurotypical and spina bifida cases. In (e), we report 4 scoring by a panel of 8 experts of the trustworthiness of the automatic segmentations for a subset of the out-of-scanner distribution testing 3D MRIs ($n = 50$). Each expert was asked to score from 0 (totally unacceptable) to 5 (perfect fit) the trustworthiness of each ROI. The scores displayed here are averaged across ROIs. Results per ROI can be found in the appendix (Fig. 7, 8, 9, 10, 11). Box limits are the first quartiles and third quartiles. The central ticks are the median values. The whiskers extend the boxes to show the rest of the distribution, except for points that are determined to be outliers. Outliers are data points outside the range median $\pm 1.5 \times$ interquartile range.

Discussion

A principled and practical trustworthy AI method. We have mathematically formalized a method for trustworthy AI with a fallback based on Dempster-Shafer theory. For application to fetal brain MRI segmentation, we have shown that our trustworthy AI method can be implemented using anatomy-based and intensity-based priors. We have proposed to interpret those priors as contracts of trust in Human-AI trust theory. Altogether, we showed that our principled trustworthy AI method improves the robustness and the trustworthiness of a state-of-the-art AI algorithm for fetal brain 3D MRI segmentation.

AI-based and registration-based algorithms are complementary. AI-based algorithms and registration-based algorithms have different error patterns. In several situations we have found that the registration-based method tends to achieve better segmentation performance in terms of Hausdorff distance as compared to the AI-based method while the AI-based method achieved better segmentation performance in terms of Dice score. We have found that the segmentation performance of the fallback algorithm decreases less than for the backbone AI algorithm, when comparing out-of-scanner distribution to in-scanner distribution for neurotypical and spina bifida fetal brain 3D MRIs. In our scoring of trustworthiness on out-of-scanner distribution data, we have also found that the fallback algorithm outperformed the backbone AI algorithm for neurotypical and spina bifida cases (Fig. 3e). We think this is because the anatomical prior used by registration-based segmentation methods prevents mislabelling voxels far from the real anatomy. In contrast, AI-based method are unconstrained and such errors can occur. This is what we observe for the out-of-distribution cases displayed in Fig. 1c, 2. Our proposed fail-safe method uses the registration-based segmentation with added margins with the aim to automatically detect and discard such errors that were found to occur more often for AI-based approach than for registration-based approach.

The contracts of trust hold for sub-populations covered by brain atlases. Our implementation of trustworthy AI for fetal brain segmentation depends on the availability of spatio-temporal segmentation atlases of the fetal brain in 3D MRI. While such atlases currently exist for neurotypical fetal brain [21, 53] and fetuses with spina bifida [20], it is not the case for other fetal brain pathologies. Therefore, our contracts of trust are not expected to hold for the group *other pathologies*. This illustrates how AI trustworthiness is context-dependent. We found that the *other pathologies* group is the only one for which radiologists associated the fallback method, based solely on the atlases, with a lower trustworthy scores than the backbone AI algorithm (Fig. 3e). Surprisingly, we found that the trustworthy AI algorithm still performs better or on a par with the backbone AI algorithm for the *other pathologies* group. We associate this with the use of our margins and to the proposed voxel intensity prior for the cerebrospinal fluid that are specific to the trustworthy AI algorithm. For the *other pathologies* group, we used the margin values estimated for spina bifida. Our group *other pathologies* gathers diverse rare developmental diseases associated with different variations of the fetal brain anatomy. However, due to the low number of examinations available per pathology, grouping them was

necessary for evaluation purposes. This introduces biases when comparing the segmentation performance for the *other pathologies* groups associated with in-scanner and out-of-scanner distribution. In particular, some of the fetuses with *other pathologies* in the in-scanner distribution had very severe brain anatomical abnormalities, such as aqueductal stenosis with large supratentorial ventricles and caudal displacement of the cerebellum or intracranial hemorrhage with parenchymal destruction and ventriculomegaly. In contrast, the one in the out-of-scanner distribution have milder brain abnormalities, such as moderate ventriculomegaly, and there were no cases with parenchymal destruction. This explains why, for this condition, we observe more outliers with low Dice scores and high Hausdorff distances for the backbone AI algorithm for in-scanner-distribution as compared to out-of-distribution 3D MRIs (Fig. 3a,3b). This is also the only group for which the trustworthy AI method does not significantly outperform the AI method in terms of Dice score (Fig. 7a).

The two histograms of gestational ages for the training spina bifida 3D MRIs and the in-scanner-distribution testing spina bifida 3D MRIs are not uniform and have the same shape (see Fig. 4). In contrast the histogram of gestational ages for the out-of-scanner-distribution testing spina bifida is more uniform. This might partly explain the degradation of Dice scores and Hausdorff distances, with the appearance of outliers, between in-scanner-distribution and out-of-scanner-distribution for the backbone AI algorithm (Fig. 3a,3b). Training and in-scanner-distribution testing spina bifida MRIs were mostly clinical data acquired at UHL. In this center, MRI of spina bifida are typically performed a few days before and after the surgery that is performed prior to 26 weeks of gestation. In addition, a follow-up MRI is sometimes performed one month after the surgery. This explains the two modes observed in the histograms for those two groups. In the training data, the use of the spina bifida atlas [20], that has a uniform gestational age distribution, makes the second mode less visible. Our results suggest the trustworthy AI algorithm is more robust than the AI algorithm to the gestational ages distributional shift between training and testing.

For gestational ages lower than 27 weeks, the Dice scores and Hausdorff distances degrade for all the algorithms (Fig. 3c,3d). For the backbone AI algorithm this is surprising given that more MRIs acquired at gestational ages lower than 27 weeks than higher were present in the training dataset (Fig. 4). Poorer MRI quality, which is typical for younger fetuses, might explain this degradation. In addition, the ratio of spina bifida over neurotypical examinations is higher for gestational ages lower than 27 weeks in our dataset. The abnormal brain anatomy of spina bifida cases leads to more difficult segmentation compared to neurotypical cases. This is particularly the case for several classes: the cerebellum, the extra-axial cerebrospinal fluid (CSF), the cortical gray matter, and the brainstem (Fig. 7,8,9,10,11). The cerebellum is more difficult to detect using MRI before surgery as compared to early or late after surgery [1, 8]. This has already been found to affect the segmentation performance of AI-based algorithms in previous work [17]. For neurotypical fetuses, the extra-axial CSF is present all around the cortex. However, for fetal brain MRI of spina bifida fetuses with gestational ages of 27 weeks or less this is often not the case and the extra-axial CSF might be reduced to several small connected components that do not embrace the entire cortex anymore [20]. We hypothesize that the spina bifida atlas does not cover well this variability of the extra-axial CSF [20]. Due

to the explicit spatial regularization, medical image registration cannot tackle such differences of topology. Therefore, using the atlas currently available, the contract of trust for extra-axial CSF does not apply for this group of spina bifida cases. It can also influence nearby regions, such as the cortical gray matter in this case. For the fallback algorithm and the trustworthy AI algorithm, a further degradation of the segmentation performance for gestational ages lower than 21 weeks was expected because the fetal brain atlases used start at 21 weeks. For gestational ages of 21 weeks or higher, the trustworthy AI outperforms either the backbone AI-algorithm or the fallback algorithm and performs better or on a par with the best other algorithms for all regions of interest in terms of Dice score and Hausdorff distance (Fig. 9,10). The confidence intervals are also similar or narrower for the trustworthy AI algorithm than for the other algorithms for gestational ages higher or equal to 21 weeks. This illustrates that our contracts of trust improve the robustness of the proposed trustworthy AI algorithm for spina bifida for the range of gestational ages covered by the atlas used [20].

Future work. For this work we have created the largest manually segmented fetal brain MRI dataset to date that consists of 540 fetal brain 3D MRIs from 13 acquisition centers. A recent trend in medical image processing using AI is to gather even larger multi-institutional datasets using methods such as federated learning [48]. One can hypothesize that, with enough data, the AI algorithm would get more accurate even in the worst case until eventually reaching the same accuracy as the trustworthy AI algorithm in all cases. However, results of our stratified evaluation suggest that this will require manually annotated 3D MRIs for every scanner acquisition protocol, for every condition, and for every gestational age. To give an order of magnitude of the required dataset size, if we consider that 10 3D MRIs are required for each gestational age from 19 weeks to 38 weeks, for each of 10 conditions and each of 5 hospitals, we would already need 10,000 3D MRIs for both training and testing. Given the low prevalence of some conditions [17] and the cost of obtaining fully-segmented data, classical supervised learning approaches might not be sufficient. This rough estimation does not even include important confounding factors such as ethnicity and gender. Altogether, this suggests that gathering more training data to improve the AI algorithm prior to deployment might not be sufficient to make the AI algorithm alone trustworthy.

The proposed fail-safe mechanism, that is part of our trustworthy AI method, could be used to help improving the backbone AI continuously after its deployment. An AI incident could be declared when a large part of the AI algorithm prediction was discarded by the *fail-safe mechanism*. This would allow automatic detection of images to correct and include in priority in the training set to update the backbone AI algorithm. In addition, reporting such incidents could help to further improve the trust of the user. In the context of trust, it is important to report such issues even when the incidents were handled correctly using the fallback segmentation algorithm. In addition, as part of the European Union Medical Device Regulations (EU MDR) Article 87 on “Reporting of serious incidents and field safety corrective actions” [14], it is a requirement for medical device manufacturers to report device-related incidents. Previous methods for global segmentation failures detection, i.e. at the image-level, were

proposed [32, 49]. In contrast, our fail-safe mechanism approaches the problem locally, i.e. at the voxel-level, by using an atlas-based anatomical prior and an intensity-based prior.

The margins used in our trustworthy AI segmentation algorithm for fetal brain MRI could also support interactive segmentation. Instead of providing voxel-level corrections or scribbles, the annotator could interact with the automatic segmentation by manually adapting the margins for its annotation. After manual adjustment, the voxels outside the margins are automatically marked as correctly labelled while for the voxel inside the margins will be assigned a set of possible labels. This yields partial annotations that can be exploited to improve the backbone AI method using partially-supervised learning methods [16]. This use of margins is similar, in terms of user interaction, to the safety margins that are used in clinics for radiation therapy planning [41].

The expert raters also emphasized that some of the most frequent major violations in the cortex layer could be quickly removed manually and that they would have given higher scores to the segmentations if they could interact with them. This echoes previous work on computational-aided decision making that found that users are more satisfied with imperfect algorithms if they can interact with them [9]. Our findings suggest that allowing interactions would also increase the trust of human users in AI algorithms for medical image segmentation.

Histograms of Number of 3D MRIs per Gestational Age (in weeks)

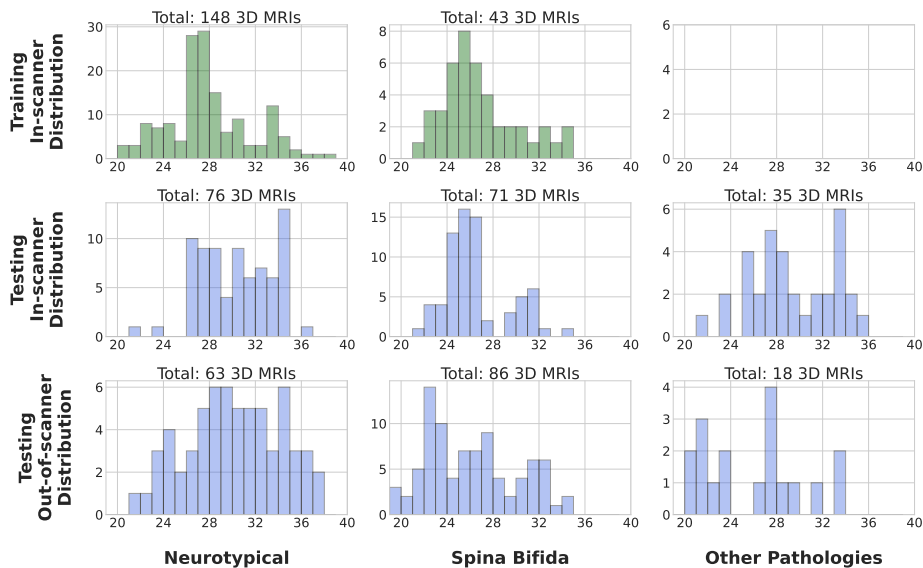


Figure 4: **Composition of the training and testing datasets (total: 540 3D MRIs)** *In-scanner distribution* designates the 3D MRIs acquired at the same center as the training data. *Out-of-scanner distribution* designates the 3D MRIs acquired at different centers than the training data. This is the largest fetal brain MRI dataset reported to date.

Methods

Fetal brain MRI dataset

We have collected a dataset with a total of 540 fetal brain 3D MRIs with neurotypical or abnormal brain development and from 13 sources of data across 6 countries.

The dataset consists of 326 3D MRIs acquired at University Hospital Leuven (UHL), 88 3D MRIs from the FeTA dataset [43] (data release 1 and 2), 11 3D MRIs acquired at Medical University of Vienna (MUV), 29 3D MRIs acquired at King’s College London (KCL), 27 3D MRIs acquired at University College London Hospital (UCLH), 4 3D MRIs acquired at Manchester (MCT), 4 3D MRIs acquired at Belfast (BFT), 2 3D MRIs acquired at Cork (CRK), 1 3D MRIs acquired at Newcastle (NCS), 1 3D MRIs acquired at Liverpool (LVP), and 47 3D MRIs from three fetal brain brain atlases. The three open-access fetal brain spatio-temporal atlases consist of 18 population-averaged 3D MRIs computed from fetal neurotypical brain MRIs acquired at Boston Children’s Hospital, USA [21], 14 population-averaged 3D MRIs computed from fetal neurotypical brain MRIs acquired in China [53], and 15 population-averaged 3D MRIs computed from fetal spina bifida brain MRIs acquired at UHL and UCLH [20].

Data from UHL includes 192 3D MRIs of neurotypical fetuses, 99 3D MRIs of fetuses with spina bifida aperta, and 35 3D MRIs of fetuses with an abnormal brain anatomy due to a condition other than spina bifida. The majority of the

neurotypical fetuses was scanned for a suspected abnormality somewhere else than in the brain, while a minority was scanned for screening of brain abnormality but was proven neurotypical after MRI. The 35 3D MRIs of fetuses with other abnormalities consisted of: 3 examinations of a case with an enlarged sub-arachnoid space, 3 cases of intraventricular hemorrhage, 1 cases of intracranial hemorrhage, 1 case with a partial rombencephalosynapsis, 1 case with a closed lip Schizencephaly, 4 cases with Dandy-Walker malformation, 1 case with an unilateral ventriculomegaly due to a hemorrhage, 1 case with choroid plexus papilloma, 1 case with high-flow dural sinus malformation, 7 cases with corpus callosum agnesis, 1 case with corpus callosum agnesis with interhemispheric cyst, temporal cysts and delayed gyration, 2 cases with tuberous sclerosis, 1 case with a Blake’s pouch cyst, 2 cases with aqueductal stenosis, 1 case with an idiopathic dilatation of the lateral ventricles, 2 cases with cytomegalovirus encephalitis, and 1 case with parenchyma loss due to an ischemic insult.

Data from the publicly available FeTA dataset [43] includes 34 3D MRIs of fetuses with a normal brain development, 36 3D MRIs of fetuses with spina bifida aperta, and 18 3D MRIs of fetuses with an abnormal brain anatomy due to conditions other than spina bifida. Those 18 3D MRIs of fetuses with other abnormalities consisted of: 3 cases with heterotopia, 8 cases with ventriculomegaly without spina bifida, 2 cases with aqueductal stenosis, 2 cases with interhemispheric cyst, 1 case with cerebellar hemorrhage, 1 case with a high-flow dural sinus malformation, and 1 case with bilateral subependymal cysts and temporal cysts.

Data from KCL consists exclusively of brain 3D MRIs of fetuses with a normal brain development. Data from MUV, UCLH, MCT, BFT, CRK, NCS, and LVP consist only of 3D MRIs of fetuses diagnosed with spina bifida aperta.

The composition of the training and testing datasets is summarized in Fig.4. The training dataset consists of the 47 volumes from the three fetal brain atlases, 144 neurotypical cases from UHL and 28 spina bifida cases from UHL. The rest of the data is used for testing ($n = 349$). In the testing dataset, 3D MRIs that were acquired at UHL are designated as *in-scanner-distribution* while the data from other acquisition centers are designated as *out-of-scanner-distribution* data.

The 3D MRIs with an isotropic image resolution of 0.8mm have been reconstructed from the stacks of 2D MRI slices acquired at UHL, MUV, UCLH, MCT, BFT, CRK, NCS, and LVP using the state-of-the-art and publicly available software NiftyMIC [11]. The original 2D MRI slices were also corrected for bias field in the NiftyMIC pipeline version 0.8 using a N4 bias field correction step as implemented in SimpleITK version 1.2.4. Brain masks for those MRIs were computed using MONAI**fbs** [46], an automatic method for fetal brain extraction in 2D fetal MRIs. The 3D brain masks are reconstructed using the automatic 2D brain masks along with the 3D MRIs in NiftyMIC.

The 3D MRI reconstructions for the FeTA dataset is described in [43]. Two volumes of spina bifida cases were excluded from the total FeTA dataset because the poor quality of the 3D MRI reconstruction (`sub-feta007` and `sub-feta009`) did not allow to manually segment them reliably for the seven tissue types. The brain masks for those 3D MRIs were computed directly using the 3D MRIs and an atlas-based method as described in our previous work [18].

Background on Human-AI trust

In this section, we describe the definition of human-AI trust and trustworthiness that we use in this paper.

Artificial intelligence is defined as any automation perceived by the individual using it as having an intent [29].

Human-AI trust is multi-dimensional. For example, the user can trust an AI medical image segmentation algorithm for a given tissue type, for images coming from a given type of scanner, or for a given population and not another. This observation that trust has several facets and is context-dependent [25] has motivated the introduction of *contractual trust* [29]. A **contract of trust** is an attribute of the AI algorithm which, if not fulfilled, causes a risk in using the AI algorithm. A contract of trust is not necessarily related to the accuracy of the AI algorithm for the task at hand.

For the automatic segmentation of the heart on chest CT images, one contract of trust could be: "The heart labels are always on the left side of the body." This contract is not directly related to the accuracy of the AI segmentation algorithm. The AI algorithm can fulfil this contract and yet compute an inaccurate segmentation of the heart. This contract can also be restricted to CT images of sufficient quality. This allows to model that contract may hold in the context of CT images of sufficiently high quality but not otherwise. Context can be added to the contract and several contracts can be derived from the contract above for different contexts. In addition, this contract does not apply to individuals with dextrocardia (their heart is located on the right).

An AI algorithm is defined as **trustworthy** with respect to a contract of trust if it provides guarantees that it will abide by the obligations of said contract [29].

The requirements proposed in the EU guidelines for trustworthy AI [12] are examples of contract or set of contracts of trust. One important requirement of trustworthy AI is the technical robustness and safety [12] which is the focus of our work. The EU guidelines propose to achieve trustworthiness in practice using a *fallback plan*. However, no technical means of implementing such plan have been provided or published. Problems with the backbone AI algorithm should be detected using a *fail-safe* algorithm and a *fallback* algorithm should be available.

In the remaining of this section we will present a theoretical framework for the implementation of a trustworthy AI system leveraging Dempster-Shafer theory to implement a failsafe fallback plan. And we will show how our framework can be used to maintain several concrete contracts of trust for fetal brain MRI segmentation.

Mathematical notations

- \mathbf{C} : the set of all classes to be segmented
- $2^{\mathbf{C}}$: the set of all subsets of \mathbf{C}
- \mathbf{x} : a voxel or a pixel
- Ω : the set of all voxels or pixels (image domain)
- p : a probability vector

- m : a basic probability assignment (BPA) in Dempster-Shafer theory
- \oplus : Dempster's rule of combination

Background on Dempster-Shafer theory

In this section, we give a brief background on Dempster-Shafer theory where we introduce solely the mathematical concepts that will be used for our method: the basic probability assignment (BPA) and Dempster's rule of combination.

In Dempster-Shafer theory [50], *basic probability assignments* are a generalization of probabilities that allow to model partial information and to combine different sources of information using Dempster's rule.

Let \mathbf{C} be the set of all classes and $2^{\mathbf{C}}$ the set of all subsets of \mathbf{C} . A **basic probability assignment (BPA)** on \mathbf{C} is a function $m : 2^{\mathbf{C}} \mapsto [0, 1]$ that satisfies

$$\begin{cases} m(\emptyset) = 0 \\ \sum_{A \subset \mathbf{C}} m(A) = 1 \end{cases} \quad (1)$$

Probabilities on \mathbf{C} are functions $p : \mathbf{C} \mapsto [0, 1]$ that satisfy $\sum_{c \in \mathbf{C}} p(c) = 1$. Probabilities are equivalent to the BPAs that assign non-zeros weights only to the singletons, i.e. the sets $A = \{c\}$ for $c \in \mathbf{C}$. Indeed, given a probability p the BPA $m^{(p)}$ associated with p is defined as: $\forall c \in \mathbf{C}, m^{(p)}(\{c\}) = p(c)$ and $\forall A \subset \mathbf{C}$ with $|A| \neq 1, m^{(p)}(A) = 0$. Basic probability assignments on \mathbf{C} are therefore more general than probabilities on \mathbf{C} .

For $A \subset \mathbf{C}$, $m(A)$ is the probability that our knowledge about the true label is exactly and only: "the true class is one of the classes in A ". In particular, it does not imply that $m(B) > 0$ for any set B such that $B \subsetneq A$ or $A \subsetneq B$. This is in contrast to probabilities that can weight only the individual classes. BPAs allow to represent more precisely than probabilities what we know (and don't know) about the true class of a voxel. For example, the extreme case where we know nothing about the true class can be represented by the BPA m such that $m(\mathbf{C}) = 1$. The best one can do to try representing this case with probabilities is to define a probability p such that $\forall c \in \mathbf{C}, p(c) = \frac{1}{|\mathbf{C}|}$. However, this choice of p corresponds to the knowledge that the class distribution is uniformly random which is different from knowing nothing about the true class.

We now introduce the concept of *complete contradiction* between BPAs which we need for Dempster's rule of combination. Two BPAs on \mathbf{C} , m_1 and m_2 , are said to be **completely contradictory** if

$$\sum_{E, F \subset \mathbf{C} | E \cap F = \emptyset} m_1(E)m_2(F) = 1 \quad (2)$$

Following (1), m_1 and m_2 are completely contradictory if and only if one cannot form a pair of overlapping sets of classes (A, B) such that m_1 commits some belief to A , i.e. $m_1(A) > 0$, and m_2 commits some belief to B , i.e. $m_2(B) > 0$.

Dempster's rule of combination

Dempster's rule of combination allows to combine any pair (m_1, m_2) of BPAs on \mathbf{C} that are not completely contradictory using the formula

$$\forall A \subset \mathbf{C}, \quad (m_1 \oplus m_2)(A) = \begin{cases} \frac{\sum_{E, F \subset \mathbf{C} | E \cap F = A} m_1(E)m_2(F)}{1 - \sum_{E, F \subset \mathbf{C} | E \cap F = \emptyset} m_1(E)m_2(F)} & \text{if } A \neq \emptyset \\ 0 & \text{if } A = \emptyset \end{cases} \quad (3)$$

For any pairs (m_1, m_2) of BPAs on \mathbf{C} that are not completely contradictory, $m_1 \oplus m_2$ is also a BPA on \mathbf{C} . In addition, the relation \oplus is symmetrical and associative.

Special case of combining a probability with a BPA

One particular case that will be useful for our method is the combination of a probability p on \mathbf{C} , which we consider as a BPA, with a generic BPA m on \mathbf{C} using Dempster's rule of combination.

Since p is a probability, for $A \subset \mathbf{C}$, $p(A)$ can be non-zeros only if A is a singleton, i.e. if it exists a class $c \in \mathbf{C}$ such that $A = \{c\}$. For simplicity, we will therefore use the abusive notation when considering p as a BPA: $p(c) := p(A) = p(\{c\})$.

The relation of complete contradiction (2) between p and m can be simplified

$$\sum_{E, F \subset \mathbf{C} | E \cap F = \emptyset} p(E)m(F) = \sum_{c \in \mathbf{C}} \sum_{F \subset (\mathbf{C} \setminus \{c\})} p(c)m(F) = 1 \quad (4)$$

Similarly, if p and m are not completely contradictory, the Dempster's rule between p and m can be simplified. Let $A \subset \mathbf{C}$, $A \neq \emptyset$, using (3) we have

$$\begin{aligned} (p \oplus m)(A) &= \frac{\sum_{E, F \subset \mathbf{C} | E \cap F = A} p(E)m(F)}{1 - \sum_{E, F \subset \mathbf{C} | E \cap F = \emptyset} p(E)m(F)} \\ &= \frac{\sum_{c \in \mathbf{C}, F \subset \mathbf{C} | \{c\} \cap F = A} p(c)m(F)}{1 - \sum_{c \in \mathbf{C}} \sum_{F \subset (\mathbf{C} \setminus \{c\})} p(c)m(F)} \end{aligned} \quad (5)$$

We remark that $p \oplus m$ is also a probability on \mathbf{C} . Indeed, let $c \in \mathbf{C}$, it can exist $F \subset \mathbf{C}$ such that $\{c\} \cap F = A$ only if A is the singleton $A = \{c\}$ (we have assumed $A \neq \emptyset$). Therefore if A is not a singleton, i.e. $|A| > 1$, the sum on the numerator is empty and equal to 0. As a result, we can write for any class $c \in \mathbf{C}$

$$\begin{aligned} (p \oplus m)(c) &= \frac{\sum_{c' \in \mathbf{C}, F \subset \mathbf{C} | \{c'\} \cap F = \{c\}} p(c')m(F)}{1 - \sum_{c' \in \mathbf{C}} \sum_{F \subset (\mathbf{C} \setminus \{c'\})} p(c')m(F)} \\ &= \frac{p(c) \sum_{F \subset \mathbf{C} | c \in F} m(F)}{1 - \sum_{c' \in \mathbf{C}} \sum_{F \subset (\mathbf{C} \setminus \{c'\})} p(c')m(F)} \end{aligned} \quad (6)$$

A Dempster-Shafer approach to Trustworthy AI for medical image segmentation

In this section, we present the formulation of our method for trustworthy AI-based medical image segmentation.

Our trustworthy AI segmentation method consists of three main components: 1) a backbone AI segmentation algorithm; 2) a fallback segmentation algorithm; and 3) and a fail-safe method that detects area of conflict between the AI algorithm segmentation and the contracts of trust and switches to the fallback algorithm for those regions. An illustration for the task of fetal brain MRI segmentation is given in Fig. 1.

The AI segmentation algorithm is a high-accuracy segmentor that can be, for example, a state-of-the-art convolutional neural network for medical image segmentation. The fallback segmentation algorithm is a segmentor that might achieve lower accuracy than the AI, but is superior to the AI for other desirable properties such as robustness. The AI and fallback segmentation algorithms take as input an image to be segmented and compute for each voxel of the image a probabilities vector with one probability for each class to be segmented.

The fail-safe mechanism aims at detecting erroneous predictions of the AI segmentation algorithm that contradict one of the contracts of trust. The contracts of trust embed domain knowledge such as "there can't be white matter in this part of the brain" or "hyperintense voxels on T2 fetal brain MRI are always cerebrospinal fluid". It is worth noting that, in general, those contracts are not sufficient to segment an entire image. Most contract will not enforce a specific segmentation but rather impose that the automatic segmentation meets certain constraints. The fail-safe mechanism is not a segmentation algorithm. In the context of image segmentation, contract of trusts can only reduces the set of possible classes and reweights the class probabilities of the segmentation of a pixel or voxel. To implement the fail-safe mechanism, we propose to use a basic probability assignment (BPA) that acts on the backbone AI and the fallback class probabilities using Dempster's rule of combination (6). In addition, we assume that the fallback class probabilities never completely contradicts the BPA representing the contracts of trust. As a result, Dempster's rule of combination can be used to switch automatically between the backbone AI algorithm and the fallback algorithm when the AI class probabilities completely contradict the BPA.

Formally, the trustworthy segmentation prediction is defined for an input image I and for all voxel position \mathbf{x} as

$$p_{I,\mathbf{x}}^{\text{TWAI}} = ((1 - \epsilon)p_{I,\mathbf{x}}^{\text{AI}} + \epsilon p_{I,\mathbf{x}}^{\text{fallback}}) \oplus m_{I,\mathbf{x}}^{\text{fail-safe}} \quad (7)$$

where \oplus is the Dempster's combination rule (3), $p_{I,\mathbf{x}}^{\text{AI}}$ is the class probability prediction of the AI segmentation algorithm for voxel \mathbf{x} of image I , $p_{I,\mathbf{x}}^{\text{fallback}}$ is the class probability prediction of the fallback segmentation for voxel \mathbf{x} of image I , and $m_{I,\mathbf{x}}^{\text{fail-safe}}$ is the BPA of the fail-safe mechanism for voxel \mathbf{x} of image I . The parameter ϵ is a constant between 0 and 1.

Toy example: trustworthy traffic lights.

One contract of trust for a trustworthy traffic light system at a crossing is that green should not be shown in all directions of the crossing at the same time. To maintain this contract, traffic light controllers may use a fail-safe conflict monitor unit to detect conflicting signals and switch to a fallback light protocol. One possible fallback is to display flashing warning signal for all traffic lights.

Formally, for the example of two traffic lights at a single-lane passage, the

set of classes is all the pairs of color the two traffic light can display at the same time $\mathbf{C} = \{(c_1, c_2) \mid c_1, c_2 \in \{\text{green, orange, red, flash}\}\}$.

Let p^{backbone} be the probability of the pair of traffic lights for the default light algorithm. The probability of the fallback algorithm p^{fallback} is then defined such as $p^{\text{fallback}}(\text{flash, flash}) = 1$. The contract of trust is m defined as $m^{\text{not-all-green}}(\mathbf{C} \setminus \{(\text{green, green})\}) = 1$.

Let $\epsilon \in]0, 1]$, the trustworthy light algorithm is given by

$$p^{\text{TWAI}} = ((1 - \epsilon)p^{\text{backbone}} + \epsilon p^{\text{fallback}}) \oplus m^{\text{not-all-green}} \quad (8)$$

Using (6), we obtain

$$\begin{aligned} p^{\text{TWAI}}(\text{green, green}) &= 0 \\ p^{\text{TWAI}}(\text{flash, flash}) &= \frac{\epsilon p^{\text{fallback}}(\text{flash, flash})}{1 - (1 - \epsilon)p^{\text{backbone}}(\text{green, green})} \end{aligned} \quad (9)$$

where the amount of conflict (4) between p^{backbone} and $m^{\text{not-all-green}}$ is equal to $p^{\text{backbone}}(\text{green, green})$ and the amount of conflict between p^{fallback} and $m^{\text{not-all-green}}$ is equal to 0. In the case of complete contradiction between the default algorithm and the contract of trust, i.e. $p^{\text{backbone}}(\text{green, green}) = 1$, the trustworthy algorithm switch completely to the fallback algorithm with $p^{\text{TWAI}}(\text{flash, flash}) = p^{\text{fallback}}(\text{flash, flash}) = 1$.

Fail-safe mechanism: switching between the AI algorithm and the fallback

In this section, we describe how the the fail-safe mechanism $m^{\text{fail-safe}}$ allows to switch between the backbone AI segmentation algorithm and the fallback segmentation algorithm.

In our framework, we assume that the fallback segmentation algorithm always produces segmentation probabilities that do not contradict entirely the BPA of the contracts of trust. A trivial example of such fallback, is the uniform segmentation algorithm that assigns an equal probability to all the classes to be segmented and for all voxels. In contrast, we do not make such compatibility assumption for the AI segmentation algorithm. Not only does this make our approach applicable with any AI segmentation algorithm, but our method also relies on the incompatibility between the AI segmentation algorithm prediction and the contracts of trust to detect failure of the AI segmentation algorithm and to switch to the fallback segmentation algorithm. Formally, however small ϵ is, as long as it remains strictly positive ($\epsilon > 0$), when $p_{I,x}^{\text{AI}}$ is completely contradictory with $m_{I,x}^{\text{fail-safe}}$, we obtain that $p_{I,x}^{\text{TWAI}}$ depends only on $p_{I,x}^{\text{fallback}}$ and not on $p_{I,x}^{\text{AI}}$. On the contrary, when $p_{I,x}^{\text{AI}}$ is not contradictory with $m_{I,x}^{\text{fail-safe}}$, we obtain that $p_{I,x}^{\text{TWAI}}$ depends mainly on $p_{I,x}^{\text{AI}}$ and the contribution of $p_{I,x}^{\text{fallback}}$ is negligible for ϵ small enough.

This is due to a property of Dempster's rule of combination that can be illustrated by the following example of pairs of BPAs [54]. Let $\epsilon_1, \epsilon_2 \in]0, 1]$ and let m_1 and m_2 two BPAs on $\{a, b, c\}$ defined by: $m_1(\{a\}) = 1 - \epsilon_1$, $m_1(\{b\}) = \epsilon_1$, $m_1(\{c\}) = 0$, $m_2(\{a\}) = 0$, $m_2(\{b\}) = \epsilon_2$, and $m_2(\{c\}) = 1 - \epsilon_2$. We obtain $(m_1 \oplus m_2)(\{b\}) = 1$, however small $\epsilon_1, \epsilon_2 > 0$ are.

This characteristic of Dempster’s rule of combination has sometimes been pointed out as counter-intuitive and as a motivation against the use of this operation [54]. In contrast, we believe this to be an interesting characteristic and we use this property at our advantage to implement the fail-safe mechanism to switch between the AI algorithm and the fallback algorithm. We recommend to choose a value of ϵ small compared to 1. As a result, the fallback class with the highest probability becomes the trustworthy AI prediction highest probability only when the AI highest probabilities prediction have been eliminated by the contracts of trust BPAs under the Dempster’s combination rule. In this sense, the fallback probabilities act as a discounting factor in Dempster Shafer theory [50, 51]. Rolf Haenni [23] summarized the motivation for the use of such discounting coefficient by Sherlock Holmes’ statement [10]: “When you have eliminated the impossible, whatever remains, however improbable, must be the truth”.

Here, we consider the case in which the AI algorithm predicted probability $p_{I,\mathbf{x}}^{\text{AI}}$ is completely contradictory with $m_{I,\mathbf{x}}^{\text{fail-safe}}$ for a voxel \mathbf{x} . Using (4), this implies

$$\left\{ \begin{array}{l} \sum_{c' \in \mathbf{C}} \left(\sum_{\mathbf{C}' \subset (\mathbf{C} \setminus \{c'\})} p_{I,\mathbf{x}}^{\text{AI}}(c') m_{I,\mathbf{x}}^{\text{fail-safe}}(\mathbf{C}') \right) = 1 \\ \forall c' \in \mathbf{C}, \forall \mathbf{C}' \subset \mathbf{C} \mid c' \in \mathbf{C}', \quad p_{I,\mathbf{x}}^{\text{AI}}(c') m_{I,\mathbf{x}}^{\text{fail-safe}}(\mathbf{C}') = 0 \end{array} \right.$$

Since by construction of our Trustworthy AI approach, the fallback algorithm probability $p_{I,\mathbf{x}}^{\text{fallback}}$ is *not* completely contradictory with $m_{I,\mathbf{x}}^{\text{fail-safe}}$, we have that $((1 - \epsilon)p_{I,\mathbf{x}}^{\text{AI}} + \epsilon p_{I,\mathbf{x}}^{\text{fallback}})$ is not completely contradictory with $m_{I,\mathbf{x}}^{\text{fail-safe}}$. Using Dempster’s rule of combination (6) we obtain, for every class $c \in \mathbf{C}$

$$p_{I,\mathbf{x}}^{\text{TWAI}}(c) = \frac{\sum_{\mathbf{C}' \subset \mathbf{C} \mid c \in \mathbf{C}'} ((1 - \epsilon)p_{I,\mathbf{x}}^{\text{AI}}(c) + \epsilon p_{I,\mathbf{x}}^{\text{fallback}}(c)) m_{I,\mathbf{x}}^{\text{fail-safe}}(\mathbf{C}')}{1 - \sum_{c' \in \mathbf{C}} \sum_{\mathbf{C}' \subset (\mathbf{C} \setminus \{c'\})} ((1 - \epsilon)p_{I,\mathbf{x}}^{\text{AI}}(c') + \epsilon p_{I,\mathbf{x}}^{\text{fallback}}(c')) m_{I,\mathbf{x}}^{\text{fail-safe}}(\mathbf{C}')}$$

Since, in this illustrative case, we assume that $p_{I,\mathbf{x}}^{\text{AI}}$ is completely contradictory with $m_{I,\mathbf{x}}^{\text{fail-safe}}$ and $\epsilon > 0$, $p_{I,\mathbf{x}}^{\text{TWAI}}(c)$ becomes

$$\begin{aligned} p_{I,\mathbf{x}}^{\text{TWAI}}(c) &= \frac{\sum_{\mathbf{C}' \subset \mathbf{C} \mid c \in \mathbf{C}'} \epsilon p_{I,\mathbf{x}}^{\text{fallback}}(c) m_{I,\mathbf{x}}^{\text{fail-safe}}(\mathbf{C}')}{1 - (1 - \epsilon) - \sum_{c' \in \mathbf{C}} \sum_{\mathbf{C}' \subset (\mathbf{C} \setminus \{c'\})} \epsilon p_{I,\mathbf{x}}^{\text{fallback}}(c') m_{I,\mathbf{x}}^{\text{fail-safe}}(\mathbf{C}')} \\ &= \frac{\sum_{\mathbf{C}' \subset \mathbf{C} \mid c \in \mathbf{C}'} p_{I,\mathbf{x}}^{\text{fallback}}(c) m_{I,\mathbf{x}}^{\text{fail-safe}}(\mathbf{C}')}{1 - \sum_{c' \in \mathbf{C}} \sum_{\mathbf{C}' \subset (\mathbf{C} \setminus \{c'\})} p_{I,\mathbf{x}}^{\text{fallback}}(c') m_{I,\mathbf{x}}^{\text{fail-safe}}(\mathbf{C}')} \\ &= (p_{I,\mathbf{x}}^{\text{fallback}} \oplus m_{I,\mathbf{x}}^{\text{fail-safe}})(c) \end{aligned}$$

In this last equality, one can observe that, however small $\epsilon > 0$ can be, the trustworthy AI prediction for voxel \mathbf{x} does not depend anymore on the AI algorithm probability but only on the fallback algorithm probability. In other words, we have switched totally from the backbone AI algorithm to the fallback algorithm.

General case with multiple contracts of trust:

In general, $m_{I,\mathbf{x}}^{\text{fail-safe}}$ is a sum of contracts of trust BPAs that are not completely contradictory and $m_{I,\mathbf{x}}^{\text{fail-safe}}$ can be written as

$$m_{I,\mathbf{x}}^{\text{fail-safe}} = \bigoplus_{k=1}^K m_{I,\mathbf{x}}^{(k)} \quad (10)$$

where each $m_{I,\mathbf{x}}^{(k)}$ is a basic probability assignment (BPA), K is the number of BPAs, and $\bigoplus_{k=1}^K$ is the Dempster’s rule of combination (3) of K BPAs computed in any order. $\bigoplus_{k=1}^K$ is well defined because Dempster’s rule of combination \oplus is associative and because the $m_{I,\mathbf{x}}^{(k)}$ are not completely contradictory. The $m_{I,\mathbf{x}}^{(k)}$ represent the contracts of trust in our framework.

Specifically, for medical image segmentation we propose the following trustworthy AI model:

$$p_{I,\mathbf{x}}^{\text{TWAI}} = ((1 - \epsilon)p_{I,\mathbf{x}}^{\text{AI}} + \epsilon p_{I,\mathbf{x}}^{\text{fallback}}) \oplus m_{I,\mathbf{x}}^{\text{anatomy}} \oplus m_{I,\mathbf{x}}^{\text{intensity}} \quad (11)$$

where $m_{I,\mathbf{x}}^{\text{anatomy}}$ is the anatomical contract of trust BPA for voxel \mathbf{x} of image I , and $m_{I,\mathbf{x}}^{\text{intensity}}$ is the intensity contract of trust BPA for voxel \mathbf{x} of image I . The definitions and properties of $m_{I,\mathbf{x}}^{\text{anatomy}}$ and $m_{I,\mathbf{x}}^{\text{intensity}}$ will be derived in following two subsections of this section Dempster-Shafer anatomical contract of trust for fetal brain segmentation and Dempster-Shafer intensity-based contract of trust for fetal brain segmentation.

Dempster-Shafer anatomical contract of trust for fetal brain segmentation

In this section, we describe our proposed anatomical prior basic probability assignment (BPA) m^{anatomy} that is used in our trustworthy AI method (11).

Our anatomical prior is computed using the segmentations computed using the multi-atlas segmentation algorithm [27]. Atlas-based segmentation algorithms are anatomically-constrained due to the spatial smoothness that is imposed to the spatial transformation used to compute the segmentation. In practice, this is achieved thanks to the parameterization of the spatial transformation and the regularization loss in the registration optimization problem [6, 19, 39]. Therefore, if implemented correctly, atlas-based automatic segmentations can inherit from the anatomical prior represented by segmentation atlases.

In terms of contract of trust, every binary segmentation mask corresponding to a specific region of interest in a fetal brain atlas is associated with an anatomical contract of trust. Indeed, each of those binary segmentation masks represents the anatomy of a given tissue type, for a given gestational age and a given population of fetuses. The anatomical contracts derived from atlas-based segmentation are therefore context-dependent since the segmentation masks are specific to a class, to a gestational age, and to the population of fetuses that was used to compute the atlas [20, 21, 53]. Since only neurotypical fetal brain atlases [21, 53] and a spina bifida fetal brain atlas [20] are available in our work, our anatomical contract of trust will hold only for neurotypical and spina bifida fetuses.




Figure 5: **Implementation of the anatomical BPA for the white matter of a fetus with spina bifida.** The white matter basic probability assignment (BPA) is computed by dilating the white matter mask of the fallback algorithm $M^{\text{white matter}}$ using the margin $\eta^{\text{white matter}}$. The margin aims at eliminating the false negative for the mask $M^{\text{white matter}}$ and are estimated using the training dataset. The white matter BPA imposes that no white matter can be predicted outside the dilated mask. No constraint is imposed inside the dilated mask. The same approach is applied to all the regions of interest to be segmented.

Due to the spatial smoothness imposed to the spatial transformation, atlas-based automatic segmentations will usually be correct up to a spatial margin. Therefore, we propose to add spatial margins to the atlas-based segmentation to compute the BPAs of our anatomical contract of trust. This approach is inspired by the safety margins used in radiotherapy to account for errors including spatial registration errors [41] and is illustrated in Fig. 5 for the white matter of a fetal brain 3D MRI.

Formally, let M^c a 3D (binary) mask from an atlas for class $c \in \mathbf{C}$. We propose to define the BPA map $m^{(c)} = \left(m_{\mathbf{x}}^{(c)} \right)_{\mathbf{x} \in \Omega}$ associated with M^c as

$$\forall \mathbf{x}, \quad \begin{cases} m_{\mathbf{x}}^{(c)}(\mathbf{C} \setminus \{c\}) = 1 - \phi(d(\mathbf{x}, M^c)) \\ m_{\mathbf{x}}^{(c)}(\mathbf{C}) = \phi(d(\mathbf{x}, M^c)) \end{cases} \quad (12)$$

where $d(\mathbf{x}, M^c)$ the Euclidean distance from any voxel position vector \mathbf{x} to M^c , and $\phi : \mathbb{R}_+ \rightarrow [0, 1]$ with $\phi(0) = 1$ and ϕ non-increasing. In the following, we call **thresholding function** any function ϕ that satisfies the properties above. We give below two examples of thresholding functions. The first one is based on hard margins, while the second one is based on exponential soft margins.

$$\begin{aligned} 1) \quad \forall d \geq 0, \quad \phi(d) &= \begin{cases} 1 & \text{if } d \leq \eta \\ 0 & \text{otherwise} \end{cases} \\ 2) \quad \forall d \geq 0, \quad \phi(d) &= \exp\left(-\frac{d}{\eta}\right) \end{aligned} \quad (13)$$

In both case, $\eta > 0$ is a hyper-parameter homogeneous to a distance and can be interpreted as a safety margin for the anatomical prior. The BPA for the first

thresholding function (hard margins) can be implemented efficiently without computing explicitly the distance between every voxel \mathbf{x} and the mask M^c . An illustration is given in Fig. 5 for the white matter. We describe a method to tune the margins at training time for each class at the end of this section.

With the definition of the BPA $m_{\mathbf{x}}^{(c)}$ in (12), we formalize the following belief: far enough from the mask M^c we know for sure that the true class is not c , i.e. $m_{\mathbf{x}}^{(c)}(\mathbf{C} \setminus \{c\}) = 1$, otherwise we do not know anything for sure regarding class c , i.e. $m_{\mathbf{x}}^{(c)}(\mathbf{C}) > 0$.

No contradiction between the anatomical contracts of trust: Following the assumption of Dempster’s rule of combination in (3), we need to make sure that the BPAs m_c defined as in (12) are nowhere completely contradictory.

Proof: We show that there is always at least one class that is compatible with the set of anatomical prior. For this we need to show that for all voxel \mathbf{x} , there exists $c \in \mathbf{C}$ such that $m_{\mathbf{x}}^{(c)}(\mathbf{C} \setminus \{c\}) < 1$ and $m_{\mathbf{x}}^{(c)}(\mathbf{C}) > 0$. This holds in our case because the masks $\{M^c\}_{c \in \mathbf{C}}$ form a partition of the set of all the voxels and because, following (12), for the voxels \mathbf{x} inside mask M^c , $m_{\mathbf{x}}^{(c)}(\mathbf{C}) = 1$ and $m_{\mathbf{x}}^{(c)}(\mathbf{C} \setminus \{c\}) = 0$.

Formula for the anatomical prior: We can now define the anatomical prior BPA used in (11) for image I and voxel \mathbf{x} as

$$m_{I,\mathbf{x}}^{\text{anatomy}} = \bigoplus_{c \in \mathbf{C}} m_{\mathbf{x}}^{(c)} \quad (14)$$

where m_c is the BPA associated to the mask M_I^c for a given thresholding function, and M_I^c is the mask for class c of the segmentation obtained using the multi-atlas fallback segmentation algorithm.

We prove that for all voxel \mathbf{x} and for all subset of classes $\mathbf{C}' \subset \mathbf{C}$, the anatomical BPA mass that the true label of \mathbf{x} is not in \mathbf{C}' is equal to

$$m_{I,\mathbf{x}}^{\text{anatomy}}(\mathbf{C} \setminus \mathbf{C}') = \prod_{c \in \mathbf{C}} \left(\delta_c(\mathbf{C}') m_{\mathbf{x}}^{(c)}(\mathbf{C} \setminus \{c\}) + (1 - \delta_c(\mathbf{C}')) m_{\mathbf{x}}^{(c)}(\mathbf{C}) \right) \quad (15)$$

where for all $c \in \mathbf{C}$, δ_c is the Dirac measure associated with c defined as

$$\forall \mathbf{C}' \subset \mathbf{C}, \quad \delta_c(\mathbf{C}') = \begin{cases} 1 & \text{if } c \in \mathbf{C}' \\ 0 & \text{if } c \notin \mathbf{C}' \end{cases} \quad (16)$$

The proof of (15) can be found in the Appendix.

In practice, we are particularly interested in summing the anatomical prior BPA with probabilities using the particular case of Dempster’s rule in (6). Let \mathbf{x} a voxel and $p_{I,\mathbf{x}}$ a probability on \mathbf{C} for voxel \mathbf{x} of image I that is not completely contradictory with $m_{I,\mathbf{x}}^{\text{anatomy}}$. For all $c \in \mathbf{C}$, we can show that

$$\left(p_{I,\mathbf{x}} \oplus m_{I,\mathbf{x}}^{\text{anatomy}} \right) (c) = \frac{p_{I,\mathbf{x}}(c) m_{\mathbf{x}}^{(c)}(\mathbf{C})}{\sum_{c' \in \mathbf{C}} p_{I,\mathbf{x}}(c') m_{\mathbf{x}}^{(c')}(\mathbf{C})} \quad (17)$$

A proof of this equality can be found in the Appendix. It is worth noting that, due to the specific form of $m_{I,\mathbf{x}}^{\text{anatomy}}$ and because $p_{I,\mathbf{x}}$ is a probability, the computational cost of $p_{I,\mathbf{x}} \oplus m_{I,\mathbf{x}}^{\text{anatomy}}$ is $\mathcal{O}(\mathbf{C})$ even though there are $2^{|\mathbf{C}|}$ elements in $2^{\mathbf{C}}$.




Figure 6: **Illustration of the margin distance.** The margin distance is the minimal dilation radius to apply to the predicted binary mask so that it covers entirely the ground-truth binary mask. We have proposed to use the margin distance to define our margins used in our definition of the anatomical BPAs.

Another important remark is that when $p_{I,\mathbf{x}}$ is completely contradictory with $m_{I,\mathbf{x}}^{\text{anatomy}}$, we have $\sum_{c' \in \mathbf{C}} p_{I,\mathbf{x}}(c') m_{\mathbf{x}}^{(c')}(\mathbf{C}) = 0$.

No contradiction between the anatomical prior and the fallback: For our trustworthy AI model (11) to be valid, we need to show that the probability $((1 - \epsilon)p_{I,\mathbf{x}}^{\text{AI}} + \epsilon p_{I,\mathbf{x}}^{\text{fallback}})$, is not completely contradictory with the anatomical prior BPA.

Proof: Since $\epsilon > 0$, it is sufficient to show that $p_{I,\mathbf{x}}^{\text{fallback}}$ is not completely contradictory with $m_{I,\mathbf{x}}^{\text{anatomy}}$ for every voxel \mathbf{x} . In (12) and (14) we have defined the BPA maps $m_{I,c}$ based on the multi-atlas segmentation which is equal to p_I^{fallback} . Therefore, for any voxel \mathbf{x} , let $c \in \mathbf{C}$ the class such that $\mathbf{x} \in M^c$. We have $d(\mathbf{x}, M^c) = 0$ and therefore $m_{I,\mathbf{x}}(\mathbf{C}) = \phi(d(\mathbf{x}, M^c)) = 1$. And we have $p_{I,\mathbf{x}}^{\text{fallback}}(c) > 0$. This shows that $p_{I,\mathbf{x}}^{\text{fallback}}$ and $m_{I,\mathbf{x}}^{\text{anatomy}}$ are not completely contradictory.

Tuning the margins: The margins η were tuned for each class and each condition independently using the 3D MRIs of the fold 0 of the training dataset.

In our definition of the anatomical prior BPAs (12), the margins account for false negatives in the multi-atlas segmentation. The anatomical prior BPAs for a class c will impose to the probabilities of class c to be zeros for every voxel outside of the mask after adding the margins using dilation. Therefore, we chose the margin for a given class c to be the minimal dilation radius for the dilated mask to cover entirely the true region of class c even if it creates overlaps with other regions.

For this purpose, we propose to use a modified Hausdorff distance, called *margin distance*, that considers only the false negatives. An illustration is given in Fig. 6. Let $\text{HD}_{95}(M_{\text{pred}}, M_{\text{gt}})$ denotes the Hausdorff distance at 95% of percentile between a predicted binary mask M_{pred} and the ground-truth mask

M_{gt} . The margin distance of interest between M_{pred} and M_{gt} is defined as

$$\text{HD}_{95}^{FN}(M_{pred}, M_{gt}) = \text{HD}_{95}(M_{pred}, M_{pred} \cup M_{gt}) \quad (18)$$

The margin $\eta_{c,cond}$ for class $c \in \mathbf{C}$ and condition $cond$ (neurotypical or spina bifida) is chosen as the 95% percentile value of HD_{95}^{FN} on the fold 0 of the training dataset for the given class and condition.

For fetuses with a condition other than neurotypical or spina bifida, we chose $\eta_{c,other\ pathologies} = \max\{\eta_{c,neurotypical}, \eta_{c,spina\ bifida}\}$.

Dempster-Shafer intensity-based contract of trust for fetal brain segmentation

In this section, we describe our proposed intensity prior basic probability assignment (BPA) $m^{\text{intensity}}$ that is used in our trustworthy AI method for fetal brain 3D MRI segmentation (11).

In T2-weighted MRI, which is the MRI modality used in this work, it is known that the hyper-intense voxels inside the brain are highly likely to be part of the cerebrospinal fluid (CSF). We therefore propose to model this intensity prior about high intensities as a contract of trust.

Regarding hypo-intense voxels, it is unclear how to derive similar prior because even the CSF classes contain hypo-intense voxels, such as the choroid plexus for the intra-axial CSF class and the vein of Galena and straight sinus for the extra-axial CSF class [43]. It is also worth noting that voxels outside the brain (*background* class) can also be hyper-intense.

Formula for the intensity BPA: Let $\mathbf{C}_{high} \subset \mathbf{C}$ be the subset of classes that contain all the classes that partition the entire CSF (intra-axial CSF and extra-axial CSF) and the background. Let $I = \{I_{\mathbf{x}}\}_{\mathbf{x} \in \Omega}$ be the volume and Ω the volume domain of a fetal brain 3D MRI. We propose to fit a Gaussian mixture model (GMM) with two components to the image intensity distribution of I . The two components of parameters $(\mu_{high}, \sigma_{high})$ and $(\mu_{low}, \sigma_{low})$ are associated to high and low intensities. We propose to define the intensity prior BPA for all voxels $\forall \mathbf{x}$, up to a normalization factor, as

$$\begin{cases} m_{I,\mathbf{x}}^{\text{intensity}}(\mathbf{C}_{high}) \propto \frac{1}{\sigma_{high}} \exp\left(\frac{1}{2} \left(\frac{I_{\mathbf{x}} - \mu_{high}}{\sigma_{high}}\right)^2\right) \\ m_{I,\mathbf{x}}^{\text{intensity}}(\mathbf{C}) \propto \frac{1}{\sigma_{low}} \exp\left(\frac{1}{2} \left(\frac{I_{\mathbf{x}} - \mu_{low}}{\sigma_{low}}\right)^2\right) \end{cases} \quad (19)$$

It is worth noting that $m_{I,\mathbf{x}}^{\text{intensity}}(\mathbf{C}) > 0$. Therefore, no probability will be set to 0 using the Dempster's rule of combination with $m^{\text{intensity}}$. In other words, $m^{\text{intensity}}$ does not forbid any assignment. This is in contrast with the anatomical BPAs defined in Dempster-Shafer anatomical contract of trust for fetal brain segmentation.

Let \mathbf{x} a voxel and $p_{I,\mathbf{x}}$ a probability on \mathbf{C} for voxel \mathbf{x} of image I . Since $m_{I,\mathbf{x}}^{\text{intensity}}(\mathbf{C}) > 0$, $p_{I,\mathbf{x}}$ is not completely contradictory with $m_{I,\mathbf{x}}^{\text{intensity}}$. Using

Dempster’s rule, we have, for all class $c \in \mathbf{C}$

$$\left(p_{I,\mathbf{x}} \oplus m_{I,\mathbf{x}}^{\text{intensity}}\right)(c) \propto \begin{cases} \left(m_{I,\mathbf{x}}^{\text{intensity}}(\mathbf{C}_{high}) + m_{I,\mathbf{x}}^{\text{intensity}}(\mathbf{C})\right) p_{I,\mathbf{x}}(c) & \text{if } c \in \mathbf{C}_{high} \\ m_{I,\mathbf{x}}^{\text{intensity}}(\mathbf{C}) p_{I,\mathbf{x}}(c) & \text{otherwise} \end{cases} \quad (20)$$

This can be interpreted as a soft-thresholding operation

$$\left(p_{I,\mathbf{x}} \oplus m_{I,\mathbf{x}}^{\text{intensity}}\right)(c) \propto \begin{cases} \left(1 + \frac{m_{I,\mathbf{x}}^{\text{intensity}}(\mathbf{C}_{high})}{m_{I,\mathbf{x}}^{\text{intensity}}(\mathbf{C})}\right) p_{I,\mathbf{x}}(c) & \text{if } c \in \mathbf{C}_{high} \\ p_{I,\mathbf{x}}(c) & \text{otherwise} \end{cases} \quad (21)$$

Thus, only the probabilities for the background and CSF classes in \mathbf{C}_{high} are increased in the case of a voxel \mathbf{x} with relatively high intensity. In particular, the probabilities remain approximately unchanged for a voxel \mathbf{x} with relatively low or medium intensity. This reflects the fact that the background and CSF classes also contain hypo-intense voxels. The hyper-intense voxels must be in \mathbf{C}_{high} while we can not say anything about hypo-intense voxels in general. There are hypo-intense voxels in every class.

nnU-Net as backbone AI segmentation algorithm

The AI segmentation algorithm used is based on nnU-Net [28] which is a state-of-the-art deep learning-based method for medical image segmentation. We have chosen the nnU-Net deep learning pipeline because it has led to state-of-the-art results on several segmentation challenges, including the FeTA challenge 2021 for automatic fetal brain 3D MRI segmentation [18, 43]. We have used the code available at <https://github.com/MIC-DKFZ/nnUNet> without modification for our backbone AI.

Deep learning pipeline: The nnU-Net pipeline is based on a set of heuristics to automatically select the deep neural network architecture and other training hyper-parameters such as the patch size. In this work, a 3D U-Net [7] was selected with one input block, 4 down-sampling blocks, one bottleneck block, 5 upsampling blocks, 32 features in the first level, instance normalization [52], and the leaky-ReLU activation function with slope 0.01. This 3D U-Net has a total of 31.2M trainable parameters. The patch size selected is $96 \times 112 \times 96$ voxels.

Preprocessing: We have used the same pre-processing as in our previous work [18]. The 3D MRIs are skull-stripped using the brain mask after applying a dilation operation (3 iterations using a structuring element with a square connectivity equal to one) and setting the values outside the dilated brain mask to 0. The brain masks are all computed automatically either during the 3D reconstruction for the Leuven data using NiftyMIC [11, 46], or for the other data using a multi-atlas segmentation method based on affine registration [38] and three fetal brain atlases [20, 21, 53]. The intensity values inside the dilated brain mask are clipped to the percentile values at 0.5% and 99.5%, and after clipping the intensity values inside the dilated brain mask are normalized to zero mean and unit variance.

Training: the training dataset is split at random into 5 folds. In total, five 3D U-Nets are trained with one for each possible combination of 4 folds for

training and 1 fold for validation. The AI segmentation algorithm consists of the ensemble of those five 3D U-Nets. Each 3D U-Net is initialized at random using He initialization [24]. The loss function consists of the sum of the Dice loss and the cross entropy loss. Stochastic gradient descent with Nesterov momentum is used to minimize the empirical mean loss on the training dataset, with batch size 4, weight decay 3×10^{-5} , initial learning rate 0.01, deep supervision on 4 levels, and polynomial learning rate decay with power 0.9 for a total of 250,000 training iterations. The data augmentation methods used are: random cropping of a patch, random zoom, gamma intensity augmentation, multiplicative brightness, random rotations, random mirroring along all axes, contrast augmentation, additive Gaussian noise, Gaussian blurring, and simulation of low resolution. For more implementation details, we refer the interested reader to [28] and the nnU-Net GitHub page.

Inference: The probabilistic segmentation prediction of the AI segmentation algorithm is the average of the five probabilistic segmentation prediction of the five 3D U-Nets after training. In addition, for each 3D U-Net, test-time data augmentation with flip around the 3 spatial axis is performed.

Multi-atlas segmentation as fallback segmentation algorithm

The fallback segmentation algorithm that we propose to use is based on a multi-atlas segmentation approach. Multi-atlas segmentation [27] is one of the most trustworthy approaches for medical image segmentation in terms of anatomical plausibility. The multi-atlas segmentation that we use is inspired by the Geodesic Information Flows method (GIF) [6], which is a state-of-the-art multi-atlas segmentation algorithm.

In this section, we give details about the three main steps of the multi-atlas segmentation algorithm used. First, the selection of the atlas volumes to use to compute the automatic segmentation. Second, the non-linear registration algorithm to propagate each atlas segmentation to the 3D MRI to be segmented. And third, the fusion method used to combine the propagated segmentations from the atlas volumes.

Atlas volumes selection: We used the volumes from two neurotypical fetal brain 3D MRI atlases [21, 53] and one spina bifida fetal brain 3D MRI atlas [20]. Let GA be the gestational age rounded to the closest number of weeks of the 3D MRI to be segmented. We select all the atlas volumes with a gestation age in the interval $[GA - \Delta GA, GA + \Delta GA]$ with $\Delta GA = 1$ week for the neurotypical fetuses and $\Delta GA = 3$ for spina bifida fetuses. This way approximately the same number of atlas volumes are used for neurotypical and spina bifida fetuses.

Non-linear registration: Our image registration step aims at spatially aligning the selected atlas volumes with the 3D MRI to be segmented. We used NiftyReg [39] to compute the non-linear image registrations. The non-linear image registration optimization problem is the following

$$\begin{cases} \min_{\Theta} \mathcal{L}(I_{subject}, I_{atlas}, \phi(\Theta)) + R(\Theta) \\ R(\Theta) = \alpha_{BE} BE(\phi(\Theta)) + \alpha_{LE} LE(\phi(\Theta)) \end{cases} \quad (22)$$

where I_{atlas} is the segmented atlas volume to be registered to the 3D reconstructed MRI $I_{subject}$ that we aim to segment, $\phi(\Theta)$ is a spatial transformation

parameterized by cubic B-splines of parameters Θ with a grid size of 4 mm. The data term \mathcal{L} is the local normalized cross correlation (LNCC) with the standard deviation of the Gaussian kernel of the LNCC was set to 6 mm. The regularization term R is a linear combination of the bending energy (BE) and the linear energy (LE) regularization functions applied to $\phi(\Theta)$ with $\alpha_{BE} = 0.1$ and $\alpha_{LE} = 0.3$.

Prior to the non-linear registration, the brain mask of $I_{subject}$ was used to mask the voxels outside the brain and I_{atlas} was registered to $I_{subject}$ using an affine transformation. The affine transformation was computed using a symmetric block-matching approach [38] based on image intensities and the brain masks. The optimization is performed using conjugate gradient descent and a pyramidal approach with 3 levels [39]. The hyper-parameters for the non-linear registration were chosen to be the same as in a recent registration pipeline to compute a fetal brain atlas [20].

Segmentations fusion: Once all the atlas volumes $\{I_k\}_{k=1}^K$ and their probabilistic segmentations $\{S_k\}_{k=1}^K$ have been registered to the 3D MRI to be segmented $I_{subject}$ using the transformations $\{\phi_k\}_{k=1}^K$, we need to fuse the aligned segmentations $\{S_k \circ \phi_k\}_{k=1}^K$ into one segmentation. This fusion is computed via a voxel-wise weighted average using heat kernels [6].

The heat map for atlas k at voxel \mathbf{x} is defined as [6]

$$w_k(\mathbf{x}) = \exp(-(D(k, \mathbf{x}))^2) \quad (23)$$

where $D(k, \mathbf{x})$ is a surrogate of the morphological similarity between $I_{subject}$ and $I_k \circ \phi_k$ at voxel \mathbf{x} . The distance $D(k, \mathbf{x})$ is the sum of two components

$$D(k, \mathbf{x}) = \alpha L(I_{subject}, I_k \circ \phi_k)(\mathbf{x}) + (1 - \alpha) F(\phi_k)(\mathbf{x}) \quad (24)$$

with $\alpha = 0.5$, $L(I_{subject}, I_k \circ \phi_k) = B * (I_{subject} - (I_k \circ \phi_k))^2$ the local sum of squared differences convoluted (convolution operator $*$) by a B-spline kernel B of order 3, and $F(\phi_k)(\mathbf{x})$ the Euclidean norm of the displacement field at voxel \mathbf{x} (in mm) after removing the low spatial frequencies of ϕ_k using a Gaussian kernel with a standard deviation of 20 mm. The hyper-parameters chosen are the same as in GIF [6]. Before computing D , the intensity values of the images are normalized to zero mean and unit variance inside the brain mask.

The multi-atlas segmentation at voxel \mathbf{x} is computed using the heat kernels as

$$S_{multi-atlas}(\mathbf{x}) = \frac{\sum_{k=1}^K w_k(\mathbf{x})(S_k \circ \phi_k)(\mathbf{x})}{\sum_{k=1}^K w_k(\mathbf{x})} \quad (25)$$

Our implementation is available at trustworthy-ai-fetal-brain-segmentation.

Hyper-parameters tuning: The hyper-parameters that we tuned are Δ GA, the selection strategy for the atlas volumes, and the fusion strategy for combining the probabilistic segmentation of the atlas volumes after non-linear registration. For Δ GA we tried the values $\{0, 1, 2, 3, 4\}$. For the selection strategy we compared the condition-specific strategy described above to using all the atlases irrespective of the condition of the fetus. And for the fusion strategy we compared the GIF-like fusion strategy described in above to a simple average.

The data used for the selection of the hyper-parameters are the training data of the first fold that was used for the training of the AI segmentation algorithm. The mean Dice score across all the segmentation classes and the

number of volumes to register were used as our selection criteria to find a trade-off between segmentation accuracy and computational time. The results can be found in the appendix. The approach selected consists of GIF-like atlas segmentations fusions, condition specific atlas selection and $\Delta GA = 1$ for the neurotypical condition and $\Delta GA = 3$ for the spina bifida condition.

Human expert scoring method for evaluating the trustworthiness of fetal brain 3D MRI segmentation

The trustworthiness scoring is done for each of the tissue types: white matter, intra-axial CSF, cerebellum, extra-axial CSF, cortical gray matter, deep gray matter, and brainstem.

The evaluation is performed using a Likert scale ranging from 0 star to 5 stars to answer the question *"Is the automatic segmentation of the tissue type X trustworthy?"*:

☆☆☆☆☆ Strongly disagree / there are several severe violations of the anatomy that are totally unacceptable

★☆☆☆☆ Disagree / there is one severe violation of the anatomy that is totally unacceptable

★★☆☆☆ Moderately disagree / there are violations of the anatomy that make the acceptability of the segmentation questionable

★★★☆☆ Moderately agree / there are many minor violations of the anatomy that are acceptable

★★★★☆ Agree / there are a few minor violations of the anatomy that are acceptable

★★★★★ Strongly agree / perfect fit of the anatomy

This evaluation is performed on 50 3D MRIs from the FeTA dataset. We selected at random 20 neurotypical cases, 20 spina bifida cases, and 10 cases with other brain pathologies.

The scoring was performed independently by 4 individuals or groups of expert raters: MA, paediatric radiologist at University Hospital Leuven with several years of experience in manual segmentation of fetal brain MRI; AJ, MD and group leader at University Children's Hospital Zurich; AB, professor of neuro-radiology at University Hospital Zurich; and by MS and PK jointly, two MDs at Medical University of Vienna (MUV) with more than 400 hours of experience in manual segmentation of fetal brain MRI, under the supervision of 3 experts: DP, professor of radiology at MUV, GK, professor of paediatric radiology at MUV, and IP, neuroradiologist at MUV.

The human expert raters were given access to the 3D MRIs, the backbone AI segmentation, the fallback segmentation, the trustworthy AI segmentation, and the ground-truth manual segmentation for each case. The segmentation algorithms are anonymized for the raters by assigning to each of them a number from 1 to 3. The assignment of numbers to the segmentation methods was performed randomly for each case independently, i.e. we used a different random assignment for each case.

The segmentation protocol that was used for the manual segmentations can be found at <http://neuroimaging.ch/feta>. An empty csv file `ranking.csv` is made available to each rater to complete.

Data availability

Readers and reviewers can email the corresponding author (lucas.fidon@kcl.ac.uk) to request access to the private fetal brain MRI data. Access to the data from UHL will require approval by the ethics committee at UHL. Access to the data acquired at UCLH, MCT, BFT, CRK, NCS, and LVP will require approval by the Caldicott guardian at UCLH. Access to the data from MUV can be requested using the contact details available at <https://www.cir.meduniwien.ac.at/research/fetal/>. Access to the data from KCL will require approval by the ethics committee at KCL.

The FeTA dataset [43] is publicly available on Synapse: <https://doi.org/10.7303/syn23747212>. Access requires registration to Synapse and agreement to the terms of use. The manual segmentations for the fetal brain MRI of FeTA dataset, that we have contributed in this work and our previous work [17, 18], are publicly available on Zenodo: <https://doi.org/10.5281/zenodo.5148612> under the term of the Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported license (CC BY-NC-ND 3.0). Access to the data is restricted. Readers and reviewers can apply for access to the data by filling in a form. The only requirement is to acknowledge that the applicant will not use those data for commercial purposes.

The first spatio-temporal atlas of the normal developing fetal brain [21] that we have used is publicly available at http://crl.med.harvard.edu/research/fetal_brain_atlas/. Access requires readers to fill in an access form. Alternatively, one can download the fetal brain atlas directly from the NiftyMIC [11] GitHub repository. The second spatio-temporal atlas of the normal developing fetal brain [53], that was based on a chinese population, is publicly available at <https://github.com/DeepBMI/FBA-Chinese>. The spina bifida spatio-temporal atlas [20] is publicly available on Synapse under the terms of the Creative Commons Zero "No rights reserved" data waiver (CC0 1.0 Public domain dedication): <https://www.synapse.org/#!Synapse:syn25887675/wiki/611424>. Access requires registration to Synapse and agreement to the terms of use.

The fetal brain 3D T2w MRI segmentation protocol that was used for the manual segmentations can be found at <http://neuroimaging.ch/feta>.

Code availability

The code for our trustworthy AI segmentation algorithm of fetal brain 3D MRI is publicly available at `trustworthy-ai-fetal-brain-segmentation`. It included our implementation of the fallback algorithm. For medical image registration we have used the NiftyReg software that is publicly available at <https://github.com/KCL-BMEIS/niftyreg>. For the backbone AI algorithm we have used the nnU-Net software that is publicly available at <https://github.com/MIC-DKFZ/nnUNet>. nnU-Net is based on the PyTorch deep learning library and we have

used PyTorch version 1.10.

We have used python 3.8.10 for all experiments. The plots in Fig.3,3c,3d,9,10, including the confidence intervals, were computed using seaborn version 0.11.2. The statistical test were computed using scipy version 1.6.3.

Acknowledgements

This project has received funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement TRABIT No 765148. Tom Vercauteren is supported by a Medtronic / RAEng Research Chair [RCSRF1819\7\34]. Florian Kofler is supported through the SFB 824, subproject B12. Supported by Deutsche Forschungsgemeinschaft (DFG) through TUM International Graduate School of Science and Engineering (IGSSE), GSC 81.

Author contributions

LF and TV developed the theory in this study. LF developed the code and algorithms used in this study. LF developed the mathematical proofs. TV, SO, and JD supervised the overall study. LF, MA, FK, and TV designed the experiments. AD, AJ, JD, GK, AM, SO, DP allowed access to the fetal brain MRI data used in this study. MA, LF, TD, DE, FG, NM, and EV performed the manual segmentation of the fetal brain MRIs used in this study. MA, AB, AJ, PK, and MS performed the trustworthiness scoring of the segmentations. GK, DP, IP supervised the scoring. LF wrote the manuscript. TV, FK, MA, AB, AJ, GK, PK, BM, IP, DP, MS, and NM revised the manuscript.

Competing interests

The authors declare no competing interests.

References

- [1] Aertsen, M., Verduyck, J., De Keyzer, F., Vercauteren, T., Van Calenbergh, F., De Catte, L., Dymarkowski, S., Demaerel, P., Deprest, J.: Reliability of MR imaging-based posterior fossa and brain stem measurements in open spinal dysraphism in the era of fetal surgery. *American Journal of Neuroradiology* **40**(1), 191–198 (2019)
- [2] Alis, D., Yergin, M., Alis, C., Topel, C., Asmakutlu, O., Bagcilar, O., Senli, Y.D., Ustundag, A., Salt, V., Dogan, S.N., et al.: Inter-vendor performance of deep learning in segmenting acute ischemic lesions on diffusion-weighted imaging: a multicenter study. *Scientific reports* **11**(1), 1–10 (2021)
- [3] Allen, B., Agarwal, S., Coombs, L., Wald, C., Dreyer, K.: 2020 ACR data science institute artificial intelligence survey. *Journal of the American College of Radiology* **18**(8), 1153–1159 (2021)

- [4] Bakas, S., Reyes, M., Jakab, A., Bauer, S., Rempfler, M., Crimi, A., Shinohara, R.T., Berger, C., Ha, S.M., Rozycki, M., et al.: Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. arXiv preprint arXiv:1811.02629 (2018)
- [5] Cabitza, F.: Biases affecting human decision making in AI-supported second opinion settings. In: International Conference on Modeling Decisions for Artificial Intelligence. pp. 283–294. Springer (2019)
- [6] Cardoso, M.J., Modat, M., Wolz, R., Melbourne, A., Cash, D., Rueckert, D., Ourselin, S.: Geodesic information flows: spatially-variant graphs and their application to segmentation and fusion. *IEEE transactions on medical imaging* **34**(9), 1976–1988 (2015)
- [7] Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3D U-Net: learning dense volumetric segmentation from sparse annotation. In: International conference on medical image computing and computer-assisted intervention. pp. 424–432. Springer (2016)
- [8] Danzer, E., Johnson, M.P., Bebbington, M., Simon, E.M., Wilson, R.D., Bilaniuk, L.T., Sutton, L.N., Adzick, N.S.: Fetal head biometry assessed by fetal magnetic resonance imaging following in utero myelomeningocele repair. *Fetal diagnosis and therapy* **22**(1), 1–6 (2007)
- [9] Dietvorst, B.J., Simmons, J.P., Massey, C.: Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science* **64**(3), 1155–1170 (2018)
- [10] Doyle, A.C.: *The sign of four*. Broadview Press (2010)
- [11] Ebner, M., Wang, G., Li, W., Aertsen, M., Patel, P.A., Aughwane, R., Melbourne, A., Doel, T., Dymarkowski, S., De Coppi, P., et al.: An automated framework for localization, segmentation and super-resolution reconstruction of fetal brain MRI. *NeuroImage* **206**, 116324 (2020)
- [12] European Commission: *Ethics guidelines for trustworthy AI*. Report, European Commission (2019)
- [13] European Commission: *Proposal for a regulation of the european parliament and of the council laying down harmonised rules on artificial intelligence (artificial intelligence act) and amending certain union legislative acts com/2021/206 final*. Report, European Commission (2021)
- [14] European Union: *The European Union Medical Device Regulation – Regulation (EU) 2017/745 (EU MDR)*. Regulation, European Union (2017)
- [15] Fetit, A.E., Alansary, A., Cordero-Grande, L., Cupitt, J., Davidson, A.B., Edwards, A.D., Hajnal, J.V., Hughes, E., Kamnitsas, K., Kyriakopoulou, V., et al.: A deep learning approach to segmentation of the developing cortex in fetal brain MRI with minimal manual labeling. In: *Medical Imaging with Deep Learning*. pp. 241–261. PMLR (2020)

- [16] Fidon, L., Aertsen, M., Emam, D., Mufti, N., Guffens, F., Deprest, T., Demaerel, P., David, A.L., Melbourne, A., Ourselin, S., et al.: Label-set loss functions for partial supervision: application to fetal brain 3D MRI parcellation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 647–657. Springer (2021)
- [17] Fidon, L., Aertsen, M., Mufti, N., Deprest, T., Emam, D., Guffens, F., Schwartz, E., Ebner, M., Prayer, D., Kasprian, G., et al.: Distributionally robust segmentation of abnormal fetal brain 3D MRI. In: Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, and Perinatal Imaging, Placental and Preterm Image Analysis, pp. 263–273. Springer (2021)
- [18] Fidon, L., Aertsen, M., Shit, S., Demaerel, P., Ourselin, S., Deprest, J., Vercauteren, T.: Partial supervision for the feta challenge 2021. arXiv preprint arXiv:2111.02408 (2021)
- [19] Fidon, L., Ebner, M., Garcia-Peraza-Herrera, L.C., Modat, M., Ourselin, S., Vercauteren, T.: Incompressible image registration using divergence-conforming B-splines. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 438–446. Springer (2019)
- [20] Fidon, L., Viola, E., Mufti, N., David, A., Melbourne, A., Demaerel, P., Ourselin, S., Vercauteren, T., Deprest, J., Aertsen, M.: A spatio-temporal atlas of the developing fetal brain with spina bifida aperta. Open Research Europe **1**(123) (2021). <https://doi.org/10.12688/openreseurope.13914.1>
- [21] Gholipour, A., Rollins, C.K., Velasco-Annis, C., Ouaalam, A., Akhondi-Asl, A., Afacan, O., Ortinau, C.M., Clancy, S., Limperopoulos, C., Yang, E., et al.: A normative spatiotemporal MRI atlas of the fetal brain for automatic segmentation and analysis of early brain growth. Scientific reports **7**(1), 1–13 (2017)
- [22] Gonzalez, C., Gotkowski, K., Bucher, A., Fischbach, R., Kaltenborn, I., Mukhopadhyay, A.: Detecting when pre-trained nnU-Net models fail silently for Covid-19 lung lesion segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 304–314. Springer (2021)
- [23] Haenni, R.: Are alternatives to Dempster’s rule of combination real alternatives?: Comments on “about the belief function combination and the conflict management problem”—Lefevre et al. Information Fusion **3**(3), 237–239 (2002)
- [24] He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In: Proceedings of the IEEE international conference on computer vision. pp. 1026–1034 (2015)
- [25] Hoffman, R.R.: A taxonomy of emergent trusting in the human–machine relationship. Cognitive Systems Engineering pp. 137–164 (2017)
- [26] Hong, J., Yun, H.J., Park, G., Kim, S., Laurentys, C.T., Siqueira, L.C., Tarui, T., Rollins, C.K., Ortinau, C.M., Grant, P.E., et al.: Fetal cortical

- plate segmentation using fully convolutional networks with multiple plane aggregation. *Frontiers in neuroscience* p. 1226 (2020)
- [27] Iglesias, J.E., Sabuncu, M.R.: Multi-atlas segmentation of biomedical images: a survey. *Medical image analysis* **24**(1), 205–219 (2015)
- [28] Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods* **18**(2), 203–211 (2021)
- [29] Jacovi, A., Marasović, A., Miller, T., Goldberg, Y.: Formalizing trust in artificial intelligence: Prerequisites, causes and goals of human trust in AI. In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. pp. 624–635 (2021)
- [30] Kamraoui, R.A., Ta, V.T., Tourdias, T., Mansencal, B., Manjon, J.V., Coupé, P.: Deeplesionbrain: Towards a broader deep-learning generalization for multiple sclerosis lesion segmentation. *Medical Image Analysis* **76**, 102312 (2022)
- [31] Khalili, N., Lessmann, N., Turk, E., Claessens, N., de Heus, R., Kolk, T., Viergever, M., Benders, M., Išgum, I.: Automatic brain tissue segmentation in fetal MRI using convolutional neural networks. *Magnetic resonance imaging* **64**, 77–89 (2019)
- [32] Kofler, F., Ezhov, I., Fidon, L., Pirkl, C.M., Paetzold, J.C., Burian, E., Pati, S., El Hussein, M., Navarro, F., Shit, S., et al.: Robust, primitive, and unsupervised quality estimation for segmentation ensembles. *Frontiers in Neuroscience* **15** (2021)
- [33] Kofler, F., Ezhov, I., Isensee, F., Balsiger, F., Berger, C., Koerner, M., Paetzold, J., Li, H., Shit, S., McKinley, R., et al.: Are we using appropriate segmentation metrics? identifying correlates of human expert perception for CNN training beyond rolling the DICE coefficient. *arXiv preprint arXiv:2103.06205* (2021)
- [34] van Leeuwen, K.G., Schalekamp, S., Rutten, M.J., van Ginneken, B., de Rooij, M.: Artificial intelligence in radiology: 100 commercially available products and their scientific evidence. *European radiology* **31**(6), 3797–3804 (2021)
- [35] Li, L., Sinclair, M., Makropoulos, A., Hajnal, J.V., David Edwards, A., Kainz, B., Rueckert, D., Alansary, A.: CAS-Net: Conditional atlas generation and brain segmentation for fetal MRI. In: *Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, and Perinatal Imaging, Placental and Preterm Image Analysis*, pp. 221–230. Springer (2021)
- [36] Makropoulos, A., Counsell, S.J., Rueckert, D.: A review on automatic fetal and neonatal brain MRI segmentation. *NeuroImage* **170**, 231–248 (2018)
- [37] Mårtensson, G., Ferreira, D., Granberg, T., Cavallin, L., Oppedal, K., Padovani, A., Rektorova, I., Bonanni, L., Pardini, M., Kramberger, M.G., et al.: The reliability of a deep learning model in clinical out-of-distribution MRI data: a multicohort study. *Medical Image Analysis* **66**, 101714 (2020)

- [38] Modat, M., Cash, D.M., Daga, P., Winston, G.P., Duncan, J.S., Ourselin, S.: Global image registration using a symmetric block-matching approach. *Journal of Medical Imaging* **1**(2), 024003 (2014)
- [39] Modat, M., Ridgway, G.R., Taylor, Z.A., Lehmann, M., Barnes, J., Hawkes, D.J., Fox, N.C., Ourselin, S.: Fast free-form deformation using graphics processing units. *Computer methods and programs in biomedicine* **98**(3), 278–284 (2010)
- [40] Naidich, T.P., Pudlowski, R.M., Naidich, J., Gornish, M., Rodriguez, F.: Computed tomographic signs of the chiari II malformation. part I: Skull and dural partitions. *Radiology* **134**(1), 65–71 (1980)
- [41] Niyazi, M., Brada, M., Chalmers, A.J., Combs, S.E., Erridge, S.C., Fiorentino, A., Grosu, A.L., Lagerwaard, F.J., Minniti, G., Mirimanoff, R.O., et al.: Estro-acrop guideline “target delineation of glioblastomas”. *Radiotherapy and oncology* **118**(1), 35–42 (2016)
- [42] Oakden-Rayner, L., Dunnmon, J., Carneiro, G., Ré, C.: Hidden stratification causes clinically meaningful failures in machine learning for medical imaging. In: *Proceedings of the ACM conference on health, inference, and learning*. pp. 151–159 (2020)
- [43] Payette, K., de Dumast, P., Kebiri, H., Ezhov, I., Paetzold, J.C., Shit, S., Iqbal, A., Khan, R., Kottke, R., Grehten, P., et al.: An automatic multi-tissue human fetal brain segmentation benchmark using the fetal tissue annotation dataset. *Scientific Data* **8**(1), 1–14 (2021)
- [44] Perone, C.S., Ballester, P., Barros, R.C., Cohen-Adad, J.: Unsupervised domain adaptation for medical imaging segmentation with self-ensembling. *NeuroImage* **194**, 1–11 (2019)
- [45] Pollenus, J., Lagae, L., Aertsen, M., Jansen, K.: The impact of cerebral anomalies on cognitive outcome in patients with spina bifida: A systematic review. *European Journal of Paediatric Neurology* (2020)
- [46] Ranzini, M., Fidon, L., Ourselin, S., Modat, M., Vercauteren, T.: MON-Aifbs: MONAI-based fetal brain MRI deep learning segmentation. arXiv preprint arXiv:2103.13314 (2021)
- [47] Redko, I., Morvant, E., Habrard, A., Sebban, M., Bennani, Y.: *Advances in domain adaptation theory*. Elsevier (2019)
- [48] Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H.R., Albarqouni, S., Bakas, S., Galtier, M.N., Landman, B.A., Maier-Hein, K., et al.: The future of digital health with federated learning. *NPJ digital medicine* **3**(1), 1–7 (2020)
- [49] Robinson, R., Valindria, V.V., Bai, W., Oktay, O., Kainz, B., Suzuki, H., Sanghvi, M.M., Aung, N., Paiva, J.M., Zemrak, F., et al.: Automated quality control in image segmentation: application to the UK Biobank cardiovascular magnetic resonance imaging study. *Journal of Cardiovascular Magnetic Resonance* **21**(1), 1–14 (2019)

- [50] Shafer, G.: A mathematical theory of evidence. Princeton university press (1976)
- [51] Smets, P.: Belief functions: the disjunctive rule of combination and the generalized bayesian theorem. *International Journal of approximate reasoning* **9**(1), 1–35 (1993)
- [52] Ulyanov, D., Vedaldi, A., Lempitsky, V.: Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022* (2016)
- [53] Wu, J., Sun, T., Yu, B., Li, Z., Wu, Q., Wang, Y., Qian, Z., Zhang, Y., Jiang, L., Wei, H.: Age-specific structural fetal brain atlases construction and cortical development quantification for chinese population. *NeuroImage* p. 118412 (2021)
- [54] Zadeh, L.A.: On the validity of Dempster’s rule of combination of evidence. *Infinite Study* (1979)
- [55] Zhao, L., Asis-Cruz, J., Feng, X., Wu, Y., Kapse, K., Largent, A., Quistorff, J., Lopez, C., Wu, D., Qing, K., et al.: Automated 3D fetal brain segmentation using an optimized deep learning approach. *American Journal of Neuroradiology* (2022)

Appendix A: class-wise quantitative evaluation.

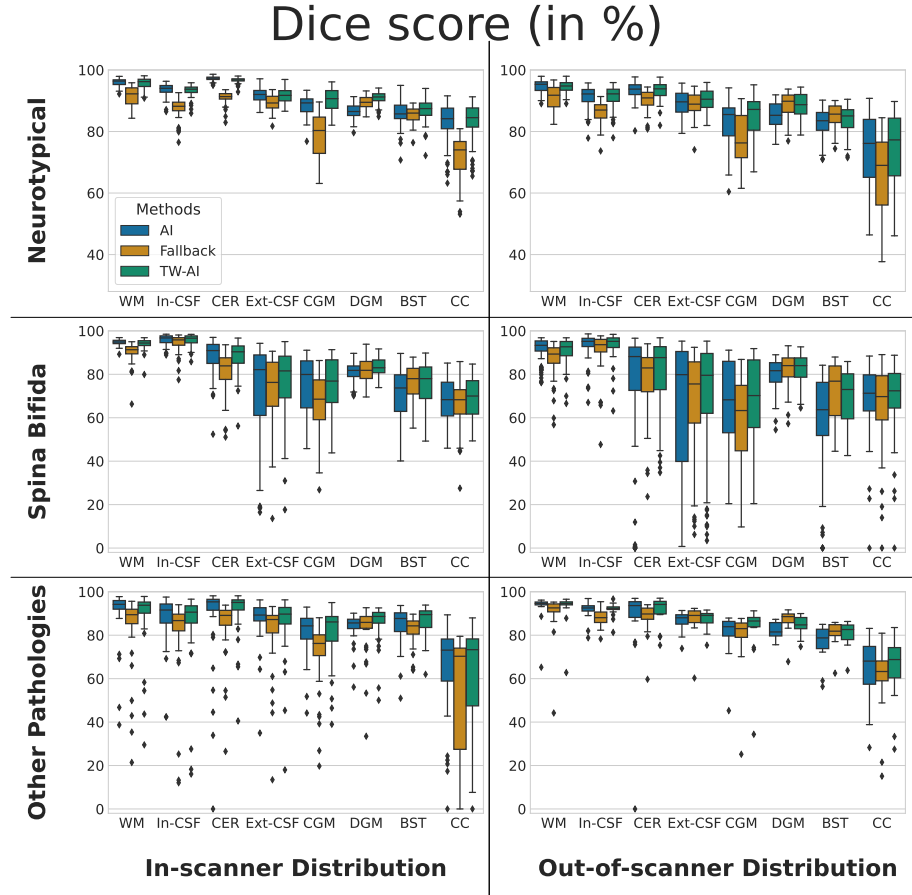


Figure 7: **Dice score (in %)** comparison of our **AI**, **fallback**, and **trust-worthy AI** segmentation algorithms for fetal brain 3D MRI segmentation. Dice scores are reported for all 3D MRI for 7 tissue types: white matter (**WM**), intra-axial cerebrospinal fluid (**in-CSF**), cerebellum (**Cer**), extra-axial cerebrospinal fluid (**Ext-CSF**), cortical gray matter (**CGM**), deep gray matter (**DGM**), brainstem (**BST**), and corpus callosum (**CC**). Box limits are the first quartiles and third quartiles. The central ticks are the median values. The whiskers extend the boxes to show the rest of the distribution, except for points that are determined to be outliers. Outliers are data points outside the range median $\pm 1.5 \times$ interquartile range. Fig. 3c was obtained from the same data after averaging the scores across regions of interest for each 3D MRI.

Hausdorff dist. 95% (in mm)

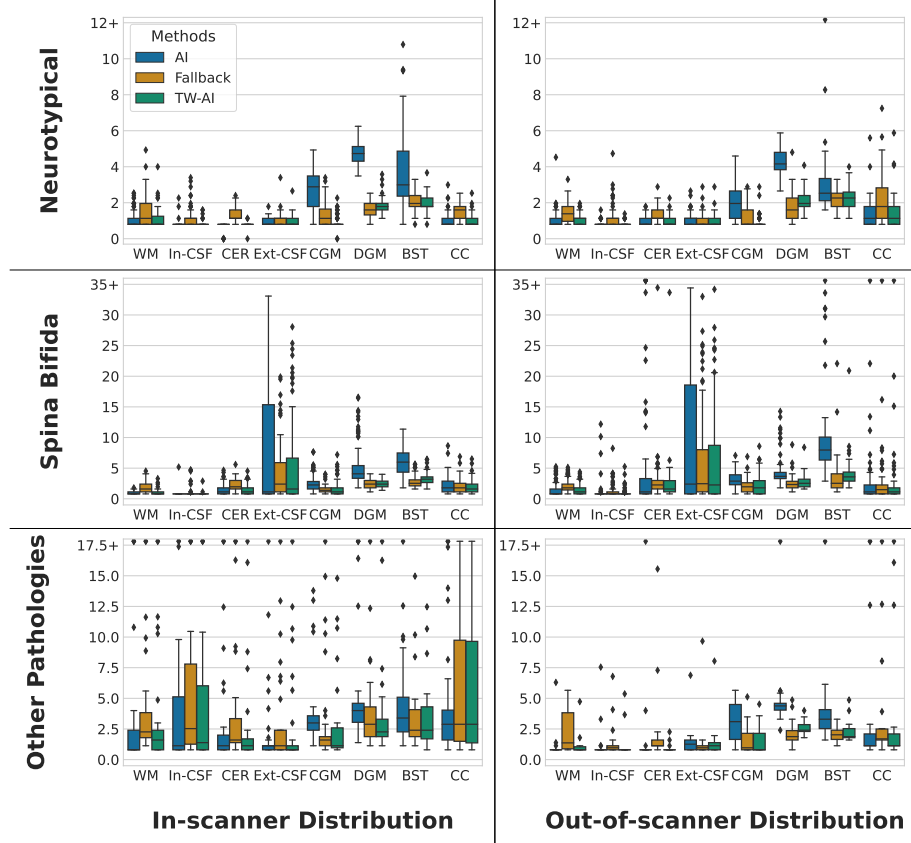


Figure 8: **Hausdorff distance (in mm) comparison of our AI, fallback, and trustworthy AI segmentation algorithms for fetal brain 3D MRI segmentation.** The organization and legend of this figure is the same as in Fig. 3a, except that here the Hausdorff distance at 95% percentile is reported in place of the Dice score. To improve the visualization we have clipped the distances to a maximum value. The clipped outliers are still visible on the top of each boxplot. Fig. 3b was obtained from the same data after averaging the scores across regions of interest for each 3D MRI. Box limits are the first quartiles and third quartiles. The central ticks are the median values. The whiskers extend the boxes to show the rest of the distribution, except for points that are determined to be outliers. Outliers are data points outside the range $\text{median} \pm 1.5 \times \text{interquartile range}$.

Mean and 95% CI for the Dice Score

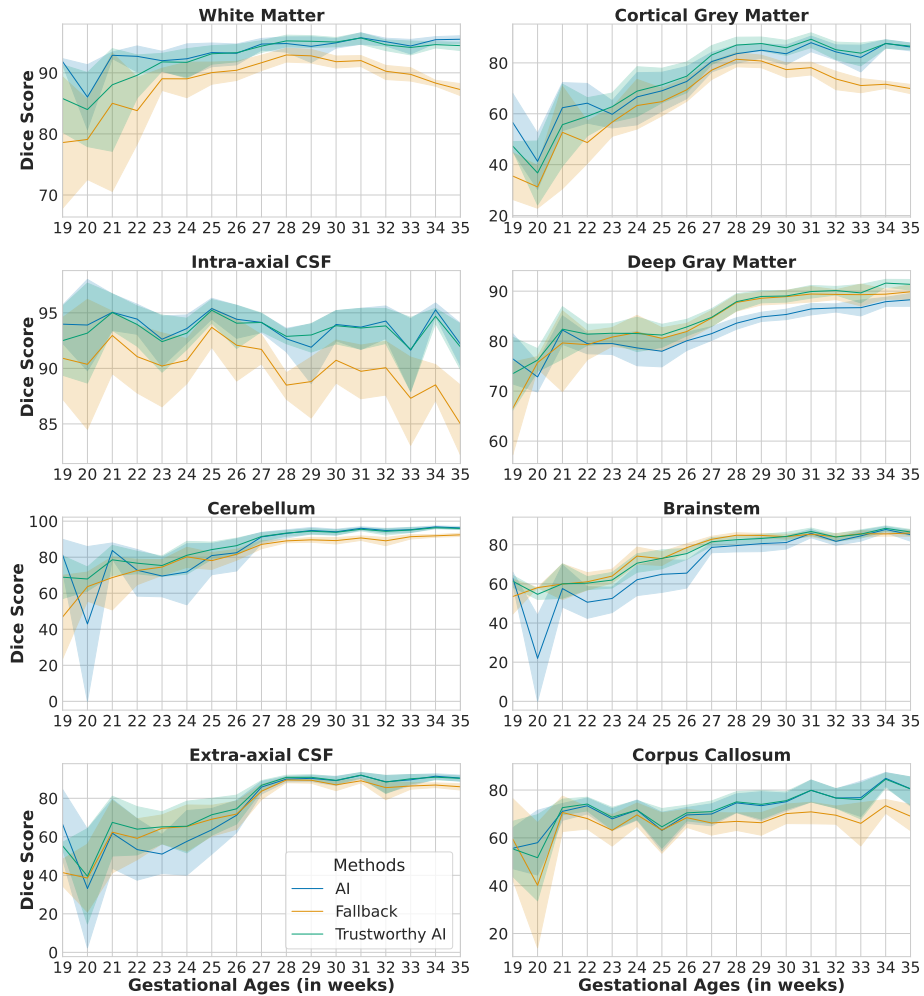


Figure 9: **Mean Dice score (in %) and 95% confidence interval as a function of the gestational age.** Here, we have used all the neurotypical and the spina bifida cases of the testing dataset. The trustworthy AI (in green) algorithm achieves similar or higher Dice scores than the best of the backbone AI (in blue) and the fallback (in orange) for all tissue type and all gestational age. Fig. 3c was obtained from the same data after averaging the scores across regions of interest for each 3D MRI. Box limits are the first quartiles and third quartiles. The central ticks are the median values. The whiskers extend the boxes to show the rest of the distribution, except for points that are determined to be outliers. Outliers are data points outside the range $\pm 1.5 \times$ interquartile range.

Mean and 95% CI for the Hausdorff Dist.

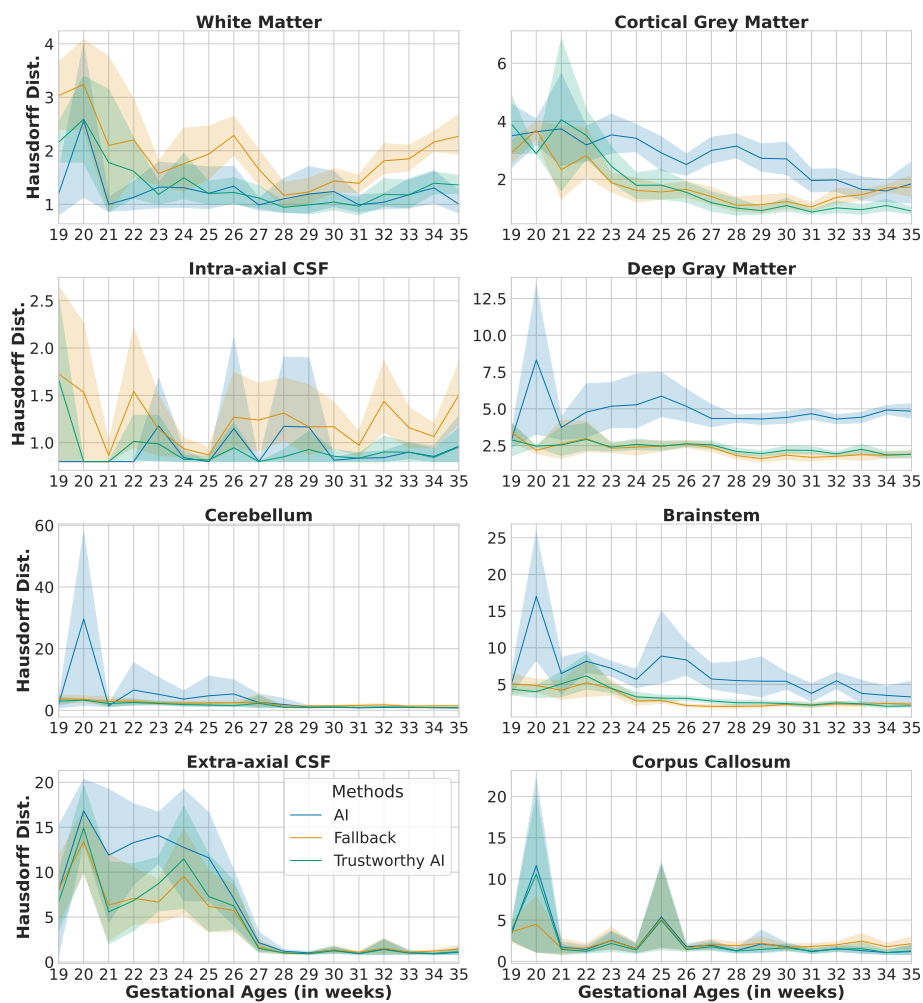


Figure 10: **Mean Hausdorff distance (in mm) and 95% confidence interval as a function of the gestational age.** Here, we have used all the neurotypical and the spina bifida cases of the testing dataset. The trustworthy IA (in green) algorithm achieves similar or lower Hausdorff distance than the best of the backbone AI (in blue) and the fallback (in orange) for all tissue type and all gestational age. Fig. 3d was obtained from the same data after averaging the scores across regions of interest for each 3D MRI.

Trustworthiness scores per ROI for out-of-distribution 3D MRIs

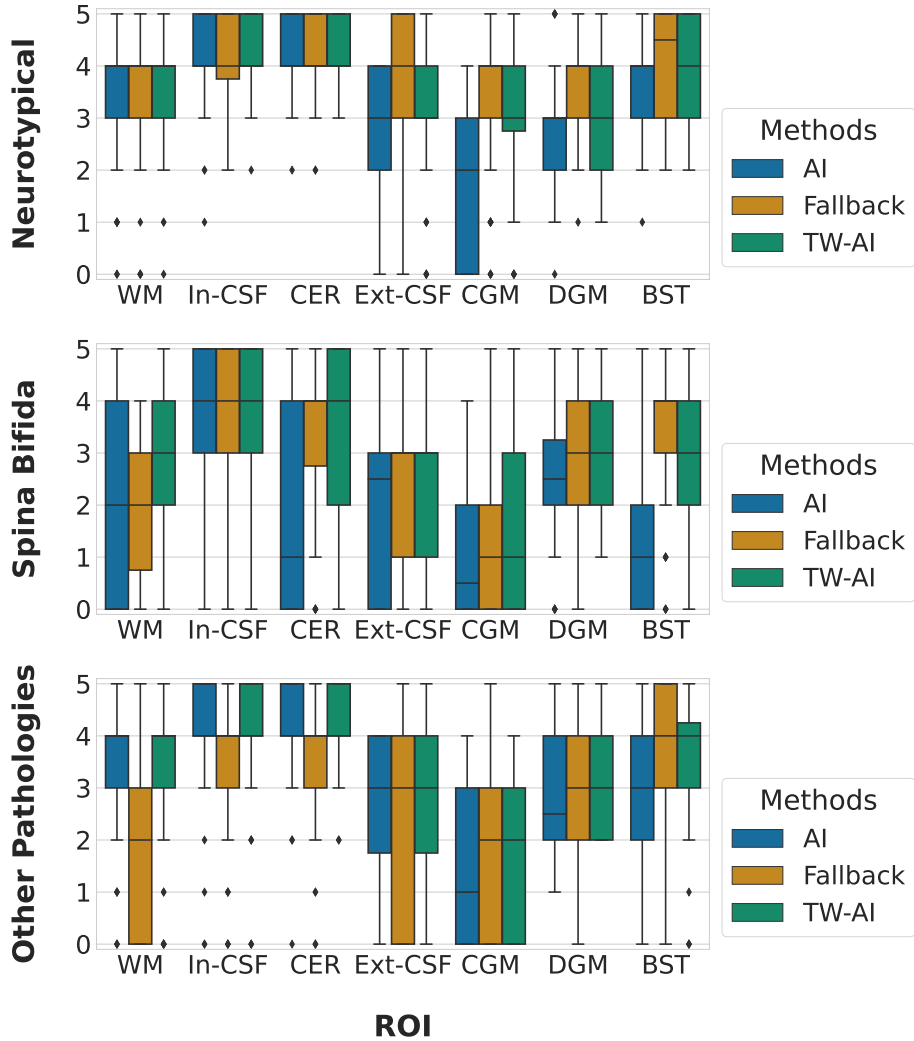


Figure 11: **Experts scores for out-of-distribution 3D MRIs.** The scores were evaluated by a panel of 8 experts for the three segmentation algorithms for each case and for all the region of interests. Experts performed 4 independent scoring and the algorithms were pseudonymized. The scores of different experts for a given region of interest and for a given 3D MRI were aggregated using averaging. Here, we have used 50 out-of-scanner distribution 3D MRIs from the FeTA dataset (20 neurotypical, 20 spina bifida, and 10 other pathologies). Box limits are the first quartiles and third quartiles. The central ticks are the median values. The whiskers extend the boxes to show the rest of the distribution, except for points that are determined to be outliers. Outliers are data points outside the range $\text{median} \pm 1.5 \times \text{interquartile range}$. Fig. 3e was obtained from the same data after averaging the scores across regions of interest for each 3D MRI.

Appendix B: tuning of the registration parameters

Table 1: **Tuning of the registration parameters.** We report the population average of the mean-class Dice score (DSC) in percentages. We also report the average number of volumes that need to be registered for each configuration. This number is approximately proportional to the computational time for the segmentation computation. The 3D MRIs used were the fold 0 of the training dataset. The row highlighted in **green** (resp. **orange**) correspond to the value of ΔGA selected for the neurotypical cases (resp. the spina bifida cases). A higher value of ΔGA leads to using more volumes in the fallback, registration-based segmentation method. Hence, this difference of ΔGA reflects the use of two neurotypical atlases [21, 53] while only one spina bifida atlas [20] is available.

Atlas fusion	Atlas selection	ΔGA	DSC Control	DSC Spina bifida	Average #volumes
Mean	Condition	0	83.9	70.0	1.6
Mean	Condition	1	84.7	72.4	4.8
Mean	Condition	2	84.8	73.1	7.7
Mean	Condition	3	84.9	73.3	10.4
Mean	Condition	4	85.1	73.3	13.0
Mean	All	0	82.0	66.4	3.1
Mean	All	1	84.6	67.8	9.4
Mean	All	2	85.0	66.9	15.3
Mean	All	3	85.1	66.8	20.8
Mean	All	4	85.2	66.8	26.0
GIF	Condition	0	84.0	72.0	1.6
GIF	Condition	1	84.8	76.1	4.8
GIF	Condition	2	85.0	76.9	7.7
GIF	Condition	3	85.2	77.6	10.4
GIF	Condition	4	85.4	77.7	13.0
GIF	All	0	84.1	72.0	3.1
GIF	All	1	84.9	74.6	9.4
GIF	All	2	85.2	75.1	15.3
GIF	All	3	85.3	75.4	20.8
GIF	All	4	85.5	75.3	26.0

Appendix C: proof of the formula (15) for the anatomical BPA.

In this section we give the proof for the formula (15): For all voxel \mathbf{x} and all $\mathbf{C}' \subset \mathbf{C}$,

$$\begin{aligned} m_{I,\mathbf{x}}^{\text{anatomy}}(\mathbf{C} \setminus \mathbf{C}') &= \left(\bigoplus_{c \in \mathbf{C}} m_{\mathbf{x}}^{(c)} \right) (\mathbf{C} \setminus \mathbf{C}') \\ &= \prod_{c \in \mathbf{C}} \left(\delta_c(\mathbf{C}') m_{\mathbf{x}}^{(c)}(\mathbf{C} \setminus \{c\}) + (1 - \delta_c(\mathbf{C}')) m_{\mathbf{x}}^{(c)}(\mathbf{C}) \right) \end{aligned} \quad (26)$$

To simplify the notations and without loss of generality, in this proof we assume that $\mathbf{C} = \{1, \dots, K\}$ with K the number of classes. This simply amounts to renaming the classes by the numbers from 1 to K .

Equation (26), that we want to prove, can then be rewritten as

$$m_{I,\mathbf{x}}^{\text{anatomy}}(\mathbf{C} \setminus \mathbf{C}') = \prod_{c=1}^K \left(\delta_c(\mathbf{C}') m_{\mathbf{x}}^{(c)}(\mathbf{C} \setminus \{c\}) + (1 - \delta_c(\mathbf{C}')) m_{\mathbf{x}}^{(c)}(\mathbf{C}) \right) \quad (27)$$

Let us first give the reader an intuition of the formula that we will prove by computing the Dempster's rule of combination for the first two BPAs $m_{\mathbf{x}}^{(1)}$ and $m_{\mathbf{x}}^{(2)}$. To simplify the calculations, we will write the Dempster's rule of combination for complements of sets like in (26).

Let $\mathbf{C}' \subsetneq \mathbf{C}$, using the definition of Dempster's rule of combination (3) and using the relation $\forall \mathbf{G}, \mathbf{H} \subset \mathbf{C}, (\mathbf{C} \setminus \mathbf{G}) \cap (\mathbf{C} \setminus \mathbf{H}) = \mathbf{C} \setminus (\mathbf{G} \cup \mathbf{H})$

$$m_{\mathbf{x}}^{(1)} \oplus m_{\mathbf{x}}^{(2)}(\mathbf{C} \setminus \mathbf{C}') = \frac{\sum_{\mathbf{G}, \mathbf{H} \subset \mathbf{C} | \mathbf{G} \cup \mathbf{H} = \mathbf{C}'} m_{\mathbf{x}}^{(1)}(\mathbf{C} \setminus \mathbf{G}) m_{\mathbf{x}}^{(2)}(\mathbf{C} \setminus \mathbf{H})}{1 - \sum_{\mathbf{G}, \mathbf{H} \subset \mathbf{C} | \mathbf{G} \cup \mathbf{H} = \mathbf{C}} m_{\mathbf{x}}^{(1)}(\mathbf{C} \setminus \mathbf{G}) m_{\mathbf{x}}^{(2)}(\mathbf{C} \setminus \mathbf{H})} \quad (28)$$

Using the definition of $m_{\mathbf{x}}^{(1)}$ and $m_{\mathbf{x}}^{(2)}$ in (12), $m_{\mathbf{x}}^{(1)}(\mathbf{C} \setminus \mathbf{G}) = 0$ if $\mathbf{G} \notin \{\emptyset, \{1\}\}$ and $m_{\mathbf{x}}^{(2)}(\mathbf{C} \setminus \mathbf{H}) = 0$ if $\mathbf{H} \notin \{\emptyset, \{2\}\}$.

This implies that the sum in the denominator is equal to zeros and that there are only four possible values of \mathbf{C}' such that the numerator is non zeros, i.e. $\mathbf{C}' \in \{\emptyset, \{1\}, \{2\}, \{1, 2\}\}$. This gives

$$\begin{aligned} \left(m_{\mathbf{x}}^{(1)} \oplus m_{\mathbf{x}}^{(2)} \right) (\mathbf{C}) &= m_{\mathbf{x}}^{(1)}(\mathbf{C}) m_{\mathbf{x}}^{(2)}(\mathbf{C}) \\ \left(m_{\mathbf{x}}^{(1)} \oplus m_{\mathbf{x}}^{(2)} \right) (\mathbf{C} \setminus \{1\}) &= m_{\mathbf{x}}^{(1)}(\mathbf{C} \setminus \{1\}) m_{\mathbf{x}}^{(2)}(\mathbf{C}) \\ \left(m_{\mathbf{x}}^{(1)} \oplus m_{\mathbf{x}}^{(2)} \right) (\mathbf{C} \setminus \{2\}) &= m_{\mathbf{x}}^{(1)}(\mathbf{C}) m_{\mathbf{x}}^{(2)}(\mathbf{C} \setminus \{2\}) \\ \left(m_{\mathbf{x}}^{(1)} \oplus m_{\mathbf{x}}^{(2)} \right) (\mathbf{C} \setminus \{1, 2\}) &= m_{\mathbf{x}}^{(1)}(\mathbf{C} \setminus \{1\}) m_{\mathbf{x}}^{(2)}(\mathbf{C} \setminus \{2\}) \end{aligned} \quad (29)$$

and $\left(m_{\mathbf{x}}^{(1)} \oplus m_{\mathbf{x}}^{(2)} \right) (\mathbf{C} \setminus \mathbf{C}') = 0$ for all other values of \mathbf{C}' .

A general formula is given by, $\mathbf{C}' \subset \mathbf{C}$,

$$\begin{aligned} \left(m_{\mathbf{x}}^{(1)} \oplus m_{\mathbf{x}}^{(2)} \right) (\mathbf{C} \setminus \mathbf{C}') &= \\ \prod_{c=3}^K (1 - \delta_c(\mathbf{C}')) \prod_{c=1}^2 \left(\delta_c(\mathbf{C}') m_{\mathbf{x}}^{(c)}(\mathbf{C} \setminus \{c\}) + (1 - \delta_c(\mathbf{C}')) m_{\mathbf{x}}^{(c)}(\mathbf{C}) \right) \end{aligned} \quad (\text{H}_2)$$

For clarity we remind that for all $c \in \mathbf{C}$, δ_c is the Dirac measure associated with c defined as

$$\forall \mathbf{C}' \subset \mathbf{C}, \quad \delta_c(\mathbf{C}') = \begin{cases} 1 & \text{if } c \in \mathbf{C}' \\ 0 & \text{if } c \notin \mathbf{C}' \end{cases} \quad (30)$$

The idea of the proof is to generalize formula (H₂) to all the combinations of the first k anatomical BPAs until reaching $k = K$.

For all $k \in \{2, \dots, K\}$, we defined (H_k) as

$$\begin{aligned} \left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)} \right) (\mathbf{C} \setminus \mathbf{C}') = \\ \prod_{c=k+1}^K (1 - \delta_c(\mathbf{C}')) \prod_{c=1}^k \left(\delta_c(\mathbf{C}') m_{\mathbf{x}}^{(c)} (\mathbf{C} \setminus \{c\}) + (1 - \delta_c(\mathbf{C}')) m_{\mathbf{x}}^{(c)} (\mathbf{C}) \right) \end{aligned} \quad (\text{H}_k)$$

When $k = K$, the set of indices for the first product is empty and the product is equal to 1 by convention. Therefore, H_K is exactly the same as relation (26) that we want to prove. We will prove this equality by induction on the variable k for k from 2 to K .

We have already proven (H₂). It remains to demonstrate that, for all k from 2 to $K - 1$, H_k holds true implies that H_{k+1} also holds true.

Let $k \in \{2, \dots, K - 1\}$, let us assume that H_k is true.

Let $\mathbf{C}' \subsetneq \mathbf{C}$, using the same formula as in (28)

$$\begin{aligned} \left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)} \right) \oplus m_{\mathbf{x}}^{(k+1)} (\mathbf{C} \setminus \mathbf{C}') = \\ \frac{\sum_{\mathbf{G}, \mathbf{H} \subset \mathbf{C} \mid \mathbf{G} \cup \mathbf{H} = \mathbf{C}'} \left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)} \right) (\mathbf{C} \setminus \mathbf{G}) m_{\mathbf{x}}^{(k+1)} (\mathbf{C} \setminus \mathbf{H})}{1 - \sum_{\mathbf{G}, \mathbf{H} \subset \mathbf{C} \mid \mathbf{G} \cup \mathbf{H} = \mathbf{C}} \left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)} \right) (\mathbf{C} \setminus \mathbf{G}) m_{\mathbf{x}}^{(k+1)} (\mathbf{C} \setminus \mathbf{H})} \end{aligned} \quad (31)$$

Let us denote

$$N = \sum_{\mathbf{G}, \mathbf{H} \subset \mathbf{C} \mid \mathbf{G} \cup \mathbf{H} = \mathbf{C}} \left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)} \right) (\mathbf{C} \setminus \mathbf{G}) m_{\mathbf{x}}^{(k+1)} (\mathbf{C} \setminus \mathbf{H}) \quad (32)$$

Let us first demonstrate that $N = 0$. Using the definition of $m_{\mathbf{x}}^{(k+1)}$ in (12), $m_{\mathbf{x}}^{(k+1)} (\mathbf{C} \setminus \mathbf{G}) = 0$ if $\mathbf{G} \notin \{\emptyset, \{k+1\}\}$. Therefore, we need only to study the cases $G \in \{\mathbf{C}, \mathbf{C} \setminus \{k+1\}\}$.

For $G = \mathbf{C}$, $\left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)} \right) (\mathbf{C} \setminus \mathbf{G}) = \left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)} \right) (\emptyset) = 0$ like for every basic probability assignment (BPA).

For $G = \mathbf{C} \setminus \{k+1\}$, according to (H_k), that we have assumed true,

$$\begin{aligned} \left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)} \right) (\mathbf{C} \setminus \mathbf{G}) &= \left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)} \right) (\{k+1\}) \\ &= \prod_{c=k+1}^K (1 - \delta_c(\mathbf{C} \setminus \{k+1\})) \prod_{c=1}^k m_{\mathbf{x}}^{(c)} (\mathbf{C} \setminus \{c\}) \end{aligned} \quad (33)$$

We have to consider two cases, $k+1 < K$ and $k+1 = K$.

If $k+1 < K$, the second term in the first product of (33) is equal to 0 and therefore $\left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)}\right) (\{k+1\}) m_{\mathbf{x}}^{(k+1)}(\mathbf{C} \setminus \{k+1\}) = 0$

If $k+1 = K$, we have

$$\left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)}\right) (\{k+1\}) m_{\mathbf{x}}^{(k+1)}(\mathbf{C} \setminus \{k+1\}) = \prod_{c=1}^K m_{\mathbf{x}}^{(c)}(\mathbf{C} \setminus \{c\}) \quad (34)$$

Voxel \mathbf{x} belongs to at least one of the class binary masks. Let us denote c_0 the binary mask to which is voxel \mathbf{x} belongs to. Using the definition of $m_{\mathbf{x}}^{(c_0)}$ in (12), we have $m_{\mathbf{x}}^{(c_0)}(\mathbf{C} \setminus \{c_0\}) = 0$. Therefore, the product above is equal to 0. This allows us to conclude, in every case, that $N = 0$.

Therefore, (31) becomes

$$\begin{aligned} \left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)}\right) \oplus m_{\mathbf{x}}^{(k+1)}(\mathbf{C} \setminus \mathbf{C}') = \\ \sum_{\mathbf{G}, \mathbf{H} \subset \mathbf{C} \mid \mathbf{G} \cup \mathbf{H} = \mathbf{C}'} \left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)}\right) (\mathbf{C} \setminus \mathbf{G}) m_{\mathbf{x}}^{(k+1)}(\mathbf{C} \setminus \mathbf{H}) \end{aligned} \quad (35)$$

Similarly as before, due to the definition of $m_{\mathbf{x}}^{(k+1)}$, we only need to study the cases of the sets \mathbf{G} that are solutions of $\mathbf{G} \cap \emptyset = \mathbf{C}'$ or $\mathbf{G} \cap \{k+1\} = \mathbf{C}'$. The first equality has the unique solution $\mathbf{G} = \mathbf{C}'$ and the second equality has either no solution, if $k+1 \notin \mathbf{C}'$, or two solutions $G \in \{\mathbf{C}' \setminus \{k+1\}, \mathbf{C}'\}$ if $k+1 \in \mathbf{C}'$. Using the Dirac measure, we can treat all the cases at once and (35) becomes

$$\begin{aligned} \left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)}\right) \oplus m_{\mathbf{x}}^{(k+1)}(\mathbf{C} \setminus \mathbf{C}') = \left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)}\right) (\mathbf{C} \setminus \mathbf{C}') m_{\mathbf{x}}^{(k+1)}(\mathbf{C}) \\ + \delta_{k+1}(\mathbf{C}') \left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)}\right) (\mathbf{C} \setminus \mathbf{C}') m_{\mathbf{x}}^{(k+1)}(\mathbf{C} \setminus \{k+1\}) \\ + \delta_{k+1}(\mathbf{C}') \left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)}\right) (\mathbf{C} \setminus (\mathbf{C}' \setminus \{k+1\})) m_{\mathbf{x}}^{(k+1)}(\mathbf{C} \setminus \{k+1\}) \end{aligned} \quad (36)$$

Using (H_k) , we can rewrite the second term of (36) as

$$\begin{aligned} \delta_{k+1}(\mathbf{C}') \left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)}\right) (\mathbf{C} \setminus \mathbf{C}') = \\ \delta_{k+1}(\mathbf{C}') \prod_{c=k+1}^K (1 - \delta_c(\mathbf{C}')) \prod_{c=1}^k \left(\delta_c(\mathbf{C}') m_{\mathbf{x}}^{(c)}(\mathbf{C} \setminus \{c\}) + (1 - \delta_c(\mathbf{C}')) m_{\mathbf{x}}^{(c)}(\mathbf{C})\right) \end{aligned} \quad (37)$$

The product of the first two terms of the product on the right-hand side is $\delta_{k+1}(\mathbf{C}')(1 - \delta_{k+1}(\mathbf{C}')) = 0$, independently to the value of \mathbf{C}' . Therefore, the second term of (36) is equal to zeros.

Using (H_k) , and by remarking that

$$\begin{aligned} \delta_{k+1}(\mathbf{C}' \setminus \{k+1\}) = 0 \\ \forall c \in \mathbf{C} \setminus \{K+1\}, \quad \delta_{k+1}(\mathbf{C}' \setminus \{k+1\}) = \delta_{k+1}(\mathbf{C}') \end{aligned} \quad (38)$$

we obtain

$$\begin{aligned} & \left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)} \right) (\mathbf{C} \setminus (\mathbf{C}' \setminus \{k+1\})) = \\ & \prod_{c=k+2}^K (1 - \delta_c(\mathbf{C}')) \prod_{c=1}^k \left(\delta_c(\mathbf{C}') m_{\mathbf{x}}^{(c)}(\mathbf{C} \setminus \{c\}) + (1 - \delta_c(\mathbf{C}')) m_{\mathbf{x}}^{(c)}(\mathbf{C}) \right) \end{aligned} \quad (39)$$

Using this equality and (\mathbf{H}_k) , we can rewrite (36) as

$$\begin{aligned} & \left(\bigoplus_{c=1}^k m_{\mathbf{x}}^{(c)} \right) \oplus m_{\mathbf{x}}^{(k+1)}(\mathbf{C} \setminus \mathbf{C}') = \\ & m_{\mathbf{x}}^{(k+1)}(\mathbf{C})(1 - \delta_{k+1}(\mathbf{C}')) \prod_{c=k+2}^K (1 - \delta_c(\mathbf{C}')) \prod_{c=1}^k \left(\delta_c(\mathbf{C}') m_{\mathbf{x}}^{(c)}(\mathbf{C} \setminus \{c\}) + (1 - \delta_c(\mathbf{C}')) m_{\mathbf{x}}^{(c)}(\mathbf{C}) \right) \\ & + \delta_{k+1}(\mathbf{C}') m_{\mathbf{x}}^{(k+1)}(\mathbf{C} \setminus \{k+1\}) \prod_{c=k+2}^K (1 - \delta_c(\mathbf{C}')) \prod_{c=1}^k \left(\delta_c(\mathbf{C}') m_{\mathbf{x}}^{(c)}(\mathbf{C} \setminus \{c\}) + (1 - \delta_c(\mathbf{C}')) m_{\mathbf{x}}^{(c)}(\mathbf{C}) \right) \end{aligned} \quad (40)$$

By grouping the two terms we eventually obtain that \mathbf{H}_{k+1} holds true, i.e.

$$\begin{aligned} & \left(\bigoplus_{c=1}^{k+1} m_{\mathbf{x}}^{(c)} \right) (\mathbf{C} \setminus \mathbf{C}') = \\ & \prod_{c=k+2}^K (1 - \delta_c(\mathbf{C}')) \prod_{c=1}^{k+1} \left(\delta_c(\mathbf{C}') m_{\mathbf{x}}^{(c)}(\mathbf{C} \setminus \{c\}) + (1 - \delta_c(\mathbf{C}')) m_{\mathbf{x}}^{(c)}(\mathbf{C}) \right) \end{aligned} \quad (41)$$

We have proved that \mathbf{H}_2 is true and we have proved that for all k from 2 to $K-1$, \mathbf{H}_k holds true implies that \mathbf{H}_{k+1} also holds true. Therefore, using the induction principle, we conclude that \mathbf{H}_K is true. ■

Appendix D: proof of equality (17).

In this section, we give a proof of (17). It states that for all $c \in \mathbf{C}$,

$$\left(p_{I,\mathbf{x}} \oplus m_{I,\mathbf{x}}^{\text{anatomy}}\right)(c) = \frac{p_{I,\mathbf{x}}(c)m_{\mathbf{x}}^{(c)}(\mathbf{C})}{\sum_{c' \in \mathbf{C}} p_{I,\mathbf{x}}(c')m_{\mathbf{x}}^{(c')}(\mathbf{C})} \quad (42)$$

We start the proof from the Dempster's rule of combination for a probability and a BPA (6).

$$\left(p_{I,\mathbf{x}} \oplus m_{I,\mathbf{x}}^{\text{anatomy}}\right)(c) = \frac{p_{I,\mathbf{x}}(c) \sum_{\mathbf{F} \subset \mathbf{C} | c \in \mathbf{F}} m_{I,\mathbf{x}}^{\text{anatomy}}(\mathbf{F})}{1 - \sum_{c' \in \mathbf{C}} \sum_{\mathbf{F} \subset (\mathbf{C} \setminus \{c'\})} p_{I,\mathbf{x}}(c')m_{I,\mathbf{x}}^{\text{anatomy}}(\mathbf{F})} \quad (43)$$

We now rewrite this equation using complement sets in the numerator to be able to use the formula (15) for the anatomical BPAs.

$$\left(p_{I,\mathbf{x}} \oplus m_{I,\mathbf{x}}^{\text{anatomy}}\right)(c) = \frac{p_{I,\mathbf{x}}(c) \sum_{\mathbf{G} \subset (\mathbf{C} \setminus \{c\})} m_{I,\mathbf{x}}^{\text{anatomy}}(\mathbf{C} \setminus \mathbf{G})}{1 - \sum_{c' \in \mathbf{C}} \sum_{\mathbf{F} \subset (\mathbf{C} \setminus \{c'\})} p_{I,\mathbf{x}}(c')m_{I,\mathbf{x}}^{\text{anatomy}}(\mathbf{F})} \quad (44)$$

Let us first simplify the numerator. Using (15) we obtain, for all $\mathbf{G} \subset (\mathbf{C} \setminus \{c\})$,

$$\begin{aligned} m_{I,\mathbf{x}}^{\text{anatomy}}(\mathbf{C} \setminus \mathbf{G}) &= \prod_{c' \in \mathbf{C}} \left(\delta_{c'}(\mathbf{C}')m_{\mathbf{x}}^{(c)}(\mathbf{C} \setminus \{c'\}) + (1 - \delta_{c'}(\mathbf{C}'))m_{\mathbf{x}}^{(c')}(\mathbf{C}) \right) \\ &= m_{\mathbf{x}}^{(c)}(\mathbf{C}) \left(\prod_{c' \in (\mathbf{C} \setminus \{c\})} \left(\delta_{c'}(\mathbf{C}')m_{\mathbf{x}}^{(c')}(\mathbf{C} \setminus \{c'\}) + (1 - \delta_{c'}(\mathbf{C}'))m_{\mathbf{x}}^{(c')}(\mathbf{C}) \right) \right) \end{aligned} \quad (45)$$

Therefore the term $m_{\mathbf{x}}^{(c)}(\mathbf{C})$ can be factorized outside of the sum in the numerator of the right-hand side of (44). Let us denote the sum of the numerator, after factorization, as

$$A_c = \sum_{\mathbf{G} \subset (\mathbf{C} \setminus \{c\})} \prod_{c' \in (\mathbf{C} \setminus \{c\})} \left(\delta_{c'}(\mathbf{C}')m_{\mathbf{x}}^{(c')}(\mathbf{C} \setminus \{c'\}) + (1 - \delta_{c'}(\mathbf{C}'))m_{\mathbf{x}}^{(c')}(\mathbf{C}) \right) \quad (46)$$

One can remark that the terms of the product are all independent of c . In addition, the c' 'th terms of the product is either $m_{\mathbf{x}}^{(c')}(\mathbf{C} \setminus \{c'\})$ or $m_{\mathbf{x}}^{(c')}(\mathbf{C})$ depending on \mathbf{G} and when summing over all $\mathbf{G} \subset (\mathbf{C} \setminus \{c\})$ we obtain all the possible products. Therefore, the sum can be factorized as

$$A_c = \prod_{c' \in (\mathbf{C} \setminus \{c\})} \left(m_{\mathbf{x}}^{(c')}(\mathbf{C} \setminus \{c'\}) + m_{\mathbf{x}}^{(c')}(\mathbf{C}) \right) \quad (47)$$

Using the definition of the anatomical BPAs, we have, for all $c \in \mathbf{C}$, $m_{\mathbf{x}}^{(c)}(\mathbf{C} \setminus \{c'\}) + m_{\mathbf{x}}^{(c')}(\mathbf{C}) = 1$. As a result, we obtain $A_c = 1$.

This proves, that

$$\left(p_{I,\mathbf{x}} \oplus m_{I,\mathbf{x}}^{\text{anatomy}}\right)(c) \propto p_{I,\mathbf{x}}(c)m_{\mathbf{x}}^{(c)}(\mathbf{C}) \quad (48)$$

And since $\sum_{c \in \mathbf{C}} \left(p_{I, \mathbf{x}} \oplus m_{I, \mathbf{x}}^{\text{anatomy}} \right) (c) = 1$, we can conclude without additional calculations that

$$\left(p_{I, \mathbf{x}} \oplus m_{I, \mathbf{x}}^{\text{anatomy}} \right) (c) = \frac{p_{I, \mathbf{x}}(c) m_{\mathbf{x}}^{(c)}(\mathbf{C})}{\sum_{c' \in \mathbf{C}} p_{I, \mathbf{x}}(c') m_{\mathbf{x}}^{(c')}(\mathbf{C})} \quad (49)$$

■