This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version except for minor formatting changes. The final published version is available on SpringerLink.

Supplementary Material: Pairing-Independent Cross-Field MRI Translation: Synthetic Supervision and Structure-Preserving Domain Adaptation for Ultra-Low-Field to High-Field Synthesis

Juan Andrés Cisneros Jácome[1], Lisset Andrea Zea Monsalve[1], Laura Valentina Garcia Mosquera[2], and Luis Fernando Torres Torres[1]

[1] LTSI, Université de Rennes, Rennes, France
[2] NeuroCo Student Research Group, University of Antioquia, Medellín, Colombia

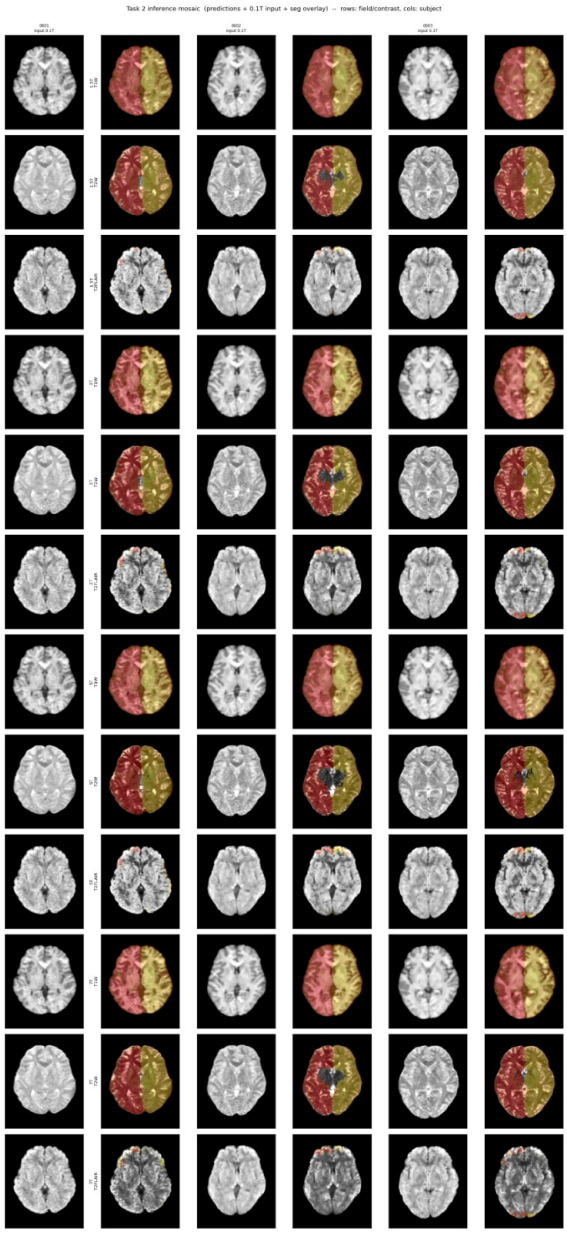

Fig. 1. Task 2 inference mosaic. Synthesized predictions, 0.1 T inputs, and SynthSeg segmentation overlays across all twelve (field, contrast) combinations (rows) and three validation subjects (columns, IDs 0001–0003). **Rows 1–8** (T1W, T2W): DGM nuclei are consistently delineated across all four target fields and subjects, confirming correct pipeline operation. **Rows 9–12** (T2FLAIR): central regions are largely black zero DGM voxels across all fields and two of three subjects, despite visually plausible macromorphology in the raw prediction panels. This is the defining visual signature of the tissue-contrast synthesis deficit reported in Sect. 3.4 of the main paper; it persists regardless of target field or training recipe (16 vs. 80 epochs).

T1W and T2W: synthesis succeeds (rows 1-8). DGM nuclei are sharply delineated and anatomically consistent across all four target fields and all three subjects, confirming that the 2.5D NAFNet backbone and SynthSeg pipeline operate correctly. These rows constitute the positive control: any failure in rows 9-12 is therefore specific to T2FLAIR, not a systemic pipeline artifact.

T2FLAIR: absent labels, not incorrect predictions (rows 9-12). Black central regions reflect zero DGM voxel assignments caused by compressed local tissue contrast (brain-masked intensity std. 0.129 vs. 0.158 for T1W; 18% reduction), not by global image failure. The raw prediction panels retain plausible sulcal and ventricular structure. The deficit is field-independent: T2FLAIR overlays are equally absent at 1.5 T (row 9, 221 training subjects) and at 5 T (row 11, 102 subjects), ruling out training-data shortage as the primary cause. It is also recipe-independent: the pattern is unchanged between the original (16-20 epochs, capped subjects) and revised (80/25 epochs, all subjects) submissions.

Subject 0002: partial trace and the CLAHE experiment (center column). Subject 0002 retains a small amygdala trace at 1.5 and 3 T (the "13/14 at zero" case in Sect. 3.4). Applying CLAHE contrast amplification before segmentation eliminated this trace entirely, taking the count to 14/14 at zero. This confirms that post-processing cannot recover tissue contrast the network did not synthesize; amplifying an absent signal difference only introduces artifacts.

Field-to-field visual similarity in T1W (rows 1-4). The four T1W rows appear more alike than four independently trained field-specific models would suggest. Without histogram matching, raw network outputs differ by only $\sim$0.08% in DGM voxel count across target fields (Sect. 2.8). The intensity differences visible in the submitted predictions therefore reflect post-hoc histogram matching rather than field-conditioned contrast learned during training.