

We used the Dhollander intermediate files from the failed runs to follow the problem through the tissue masks, safe_sdm, safe_fa, shell-wise signal/SDM, and finally back to the original DWI measurements at the affected voxels. We also investigated five subjects that completed normally as controls, after rerunning them to retain the intermediate files. Everything described below was run with MRtrix3 3.0.8.

The 3 failed cases

Subject	safe_mask	Non-finite FA in crude CSF	Crude CSF	Refined CSF	Non-finite / refined CSF	Refined SDM min / median / max
CASE01	403,618	1	13,577	1	1/1	9.98 / 9.98 / 9.98
CASE02	331,644	3	14,037	2	2/2	6.51 / 8.26 / 10.00
CASE03	402,936	2	2	1	1/1	7.21 / 7.21 / 7.21

The “abnormal” voxels dominated the final refinement. CASE03 is the cleanest example: crude CSF contained only two voxels in total, both with non-finite FA and SDM around 7.0–7.2.

Subject	Non-finite-FA voxels	Non-positive b≈1000 measurements	% of 80 high-b volumes	Refined-CSF SDM range
CASE01	1	79/80	98.8–98.8%	9.98–9.98
CASE02	3	76/80, 79/80, 77/80	95.0–98.8%	6.51–10.00
CASE03	2	79/80, 78/80	97.5–98.8%	7.21–7.21

The original DWI measurements at all six non-finite-FA locations showed the same pattern: 76–79 of the 80 b≈1000 measurements were ≤0. The low-b signal was much better preserved in these voxels, resulting in a very large apparent signal decay and therefore a very high SDM.

Controls

We ran the same diagnostics on five subjects where Dhollander completed normally. None had non-finite FA inside safe_mask. Their refined-CSF pools contained 4,279–10,679 voxels, with median SDM around 2.48–2.55, and the final 10% CSF selections contained 428–1,068 voxels.

Control	safe_mask	Non-finite FA	Refined CSF	Final CSF	Refined SDM min / median / max
CONTROL01	350,943	0	4,674	467	2.11 / 2.48 / 3.94
CONTROL02	363,259	0	4,435	444	2.14 / 2.49 / 4.81
CONTROL03	375,292	0	4,279	428	2.15 / 2.52 / 3.92
CONTROL04	394,282	0	5,338	534	2.16 / 2.53 / 4.91
CONTROL05	387,024	0	10,679	1,068	2.18 / 2.55 / 3.78

Rerunning failed subjects

We reran the Dhollander command with the default settings as used earlier, after removing the non-finite voxels from the original response-estimation input.

Subject	Original refined CSF	Rerun refined CSF	Rerun final CSF	Median SDM	Max SDM	safe_mask size
CASE01	1	4,828	483	9.98 → 2.38	9.98 → 4.12	403,618 → 403,555
CASE02	2	4,562	456	8.26 → 2.41	10.00 → 3.98	331,644 → 331,514
CASE03	1	7,090	709	7.21 → 2.48	7.21 → 4.77	402,936 → 402,885

The refined-CSF pools increased to 4,828, 4,562 and 7,090 voxels, with final CSF selections of 483, 456 and 709 voxels. Their SDM distributions approximated those of the successful controls.





Figure 1. Refined-CSF SDM median and minimum-to-maximum range.

CASE01 shows the change clearly. In the failed run, refined CSF was literally one non-finite-FA voxel at $SDM \approx 9.98$. After excluding that location from the input response mask and rerunning, refined CSF contained 4,828 voxels with median SDM 2.38 and a range of 2.03–4.12.

A. Original CASE01 failure



B. CASE01 after the rerun




Figure 2. CASE01 refined-CSF SDM histograms before and after the rerun.

Blue = all voxels; red = finite FA > 0.2; dark grey = non-finite FA (in A only).

Conclusion

As we removed the bad locations from the input response mask, those “new holes” went through Dhollander’s normal three erosion passes, meaning neighbouring voxels were also lost: safe_mask changed by 63 voxels in CASE01, 130 in CASE02 and 51 in CASE03, rather than only 1/3/2 voxels.


While this masking approach seems to do the trick for calculating response functions, we think it would be preferable to exclude voxels with non-finite safe_fa values from the subsequent crude tissue classification in DHollander.

Additional observation: FA >0.2 voxels in CSF pool

In every run with a normal crude-CSF population, all finite crude-CSF voxels were FA ≤0.2, as expected. But since during refinement, high-SDM voxels from crude WM can be added back as CSF candidates, some finite FA >0.2 voxels make it into refined CSF and the final CSF response selection.

Subject	Final CSF voxels	Finite FA >0.2	Finite FA >1.0
CASE01	483	10 (2.07%)	4 (0.83%)
CASE02	456	9 (1.97%)	1 (0.22%)
CASE03	709	13 (1.83%)	2 (0.28%)
CONTROL01	467	4 (0.86%)	0 (0.00%)
CONTROL02	444	21 (4.73%)	0 (0.00%)
CONTROL03	428	17 (3.97%)	0 (0.00%)
CONTROL04	534	9 (1.69%)	0 (0.00%)
CONTROL05	1068	27 (2.53%)	1 (0.09%)

The numbers are small - 0.86–4.73% of the final CSF voxels- but we did see occasional FA values >1.0.



Values are mean subject-level percentages, with the observed subject range in parentheses. Percentages were calculated within each subject before averaging across subjects.

CSF pool	Group	FA >0.2	FA >0.3	FA >0.5	FA >1.0
CSF refinement candidates	Corrected cases (n=3)	11.54% (9.59–12.64%)	3.95% (3.06–4.68%)	0.11% (0.08–0.13%)	0.02% (<0.01–0.04%)
CSF refinement candidates	Controls (n=5)	11.71% (6.91–14.74%)	3.70% (2.20–5.02%)	0.08% (0.03–0.16%)	<0.01% (0.00–<0.01%)
Refined CSF	Corrected cases (n=3)	3.34% (1.92–4.82%)	0.41% (0.28–0.55%)	0.14% (0.12–0.15%)	0.05% (0.02–0.10%)
Refined CSF	Controls (n=5)	2.16% (1.35–3.54%)	0.24% (0.07–0.50%)	0.07% (0.04–0.09%)	<0.01% (0.00–0.02%)
Final CSF response voxels	Corrected cases (n=3)	1.96% (1.83–2.07%)	1.58% (1.54–1.66%)	1.09% (1.04–1.13%)	0.44% (0.22–0.83%)
Final CSF response voxels	Controls (n=5)	2.75% (0.86–4.73%)	0.93% (0.21–2.03%)	0.38% (0.00–0.68%)	0.02% (0.00–0.09%)

Of note, the number of voxels with FA > 0.5 is proportionally larger in the final CSF pool as compared to the other pools. We are curious whether this number of high-FA voxels is expected and desired behaviour.