NISTIR XXXX Draft

Ongoing Face Recognition Vendor Test (FRVT)

Part 1: Verification

Patrick Grother
Mei Ngan
Kayee Hanaoka
Information Access Division
Information Technology Laboratory

This publication is available free of charge from:
https://www.nist.gov/programs-projects/face-recognition-vendor-test-frvt-ongoing

2018/02/15

National Institute of Standards and Technology
U.S. Department of Commerce

DISCLAIMER

Specific hardware and software products identified in this report were used in order to perform the evaluations described in this document. In no case does identification of any commercial product, trade name, or vendor, imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the products and equipment identified are necessarily the best available for the purpose.

ABOUT THIS REPORT

This report is a draft NIST Interagency Report, and is open for comment. It documents the verification-track of the ongoing Face Recognition Vendor Test. The report will be updated continuously as new algorithms are evaluated, as new datasets are added, and as new analyses are included. Comments and suggestions should be directed to frvt@nist.gov.

Contents

List of Tables

List of Figures

1 News

2017-01-05

- ▷ We expect to produce another edition of this FRVT report on February 13.
- ▷ NIST's evaluation of 1:1 algorithms is now closed; it will reopen in May.
- ▷ NIST's evaluation of 1:N identification algorithms has Phase 1 algorithm submission deadline February 16, and Phase 2 deadline in June..
- ▷ The 1:N API and concept document is linked from here.

2 Changelog

2018-02-15

- ▷ Added results for algorithms from new developers Gorilla, Megvii.
- ▷ Added results for new algorithms from Dermalog, Neurotechnology, Visionlabs, Vocord.
- ▷ Retired results for early algorithms from those suppliers because FRVT lists only the two most recent algorithms per organization.
- ▷ Added a column to the FNMR summary Table 2 that gives FNMR for visa photos when FMR is set to 0.0001 for non-zero effort impostors i.e. those of the same sex, age group, and country of origin. This produces markedly higher FNMR over the zero-effort case because rejection of impostors becomes more difficult.

2018-01-24

- ▷ Added results for first algorithms from Gemalto Cogent, Intellivision and Ultinous.
- ▷ Added full or partial results for new algorithms from Idemia (Morpho), NTechLab, Shaman, Tevian, and Toshiba, and from Fudan and ITMO universities.
- ▷ Added entries for Fudan University which had incorrectly been omitted from Tables 1 and 3.
- ▷ Retired results for itmo-001, ntechlab-001, vocord-001, and vigilantsolutions-001 - FRVT lists only the two most recent algorithms per organization.
- ▷ Added GPU vs CPU shape designators to the tradeoff summaries in Figures 1 and 2.

2017-12-14

- ▷ Added results for algorithms from Aware, RankOne, Shaman, and Tevian.
- ▷ Retired results for RankOne-000 - FRVT lists only the two most recent algorithms per organization.
- ▷ New description of wild images in section 5.2
- ▷ New Figure 23 showing FMR for impostors of same age, sex and worst-case region.

2017-11-16

- ▷ Added results for algorithms from 3DiVi, Dermalog, Neurotechnology and Ping An.
- ▷ Retired results for 3DiVi-000, dermalog-002, and neurotechnology-000 - FRVT lists only two algorithms per organization.
- ▷ Retired results for vcognition-001 as the algorithm was inoperable on at least one dataset.
- ▷ Added cross-pose recognition heatmaps for the wild images, Figures 10 and 11.

- ▷ Added Figure 14 to compare effect of providing whole vs. face-cropped child exploitation images.

2017-10-03

- ▷ Added results for algorithms from 3DiVi, Camvi, Idemia, Noblis, N-TechLab, and Visionlabs.
- ▷ Added partial results for two algorithms from Zhuhai Yisheng.
- ▷ The ntechlab-000 algorithm has been retired - FRVT lists only two algorithms per organization.
- ▷ Corrected fixed FMR operating point in the legends of some DET plots.

2017-08-22

- ▷ Added results for three additional algorithms, rankone-002, neurotechnology-001, and itmo-002.
- ▷ The algorithms dermalog-001 and rank-001 have been retired - FRVT lists only two algorithms per organization
- ▷ The algorithm tupel-001 has been retired as it is not operable on all datasets
- ▷ Clarified the tradeoff Figures 1 and 2 plot only genuine comparison durations.
- ▷ Corrected image type label in section 4.5

2017-08-07

- ▷ Added results for 5 new algorithms
- ▷ Added Figure 3 giving simulated example images.
- ▷ Added Figure 1 showing an alternative view of the same tradeoff data in Figure 2
- ▷ Added Figure 5 showing accuracy on visa images just for low FMR.
- ▷ Added Figure 24 showing impostor distribution shifts from certain country pairs. Section 5.6.1 in this and prior reports documents high false match rates for individuals from certain countries. That effect, however, is often not confined to anomalously high impostor scores in the tails of the distribution, but arises from systematic shifts of the whole distribution. These shifts sometimes reach 2σ.

2017-07-29

- ▷ Added results for 8 new algorithms
- ▷ Added results for a child-exploitation dataset
- ▷ Added Table 2 a standalone tabulation of false non-match rates
- ▷ We have received additional CPU algorithms - Results should appear August 4, 2017
- ▷ We have received additional GPU algorithms - Results to appear as computational resources are released from the Face Recognition Prize Challenge

2017-06-19

- ▷ Added five new algorithms, three of which remain in-process
- ▷ Added results for a "wild" dataset of images similar to non-cooperative photojournalism images
- ▷ Added Table 3 a standalone tabulation of failure to enrol rates
- ▷ Added Fig. 2 showing tradeoff between FNMR, template size, template generation time, and match duration.
- ▷ Added Fig. 19 showing how FMR is concentrated in certain images.
- ▷ Restated cross-region false match rates at nominal FMR = 0.0001 instead of 0.001
- ▷ Improved DET legends.

	Developer	Short	Seq.	Validation	Config[1]	Template		GPU	Comparison Time (ns)[3]	
	Name	Name	Num.	Date	Data (KB)	Size (B)	Time (ms)[2]		Genuine	Impostor
1	3DiVi	3divi	001	2017-06-22	190867	$^{47}4096 \pm 0$	$^{31}274 \pm 47$	Yes	$^{13}636 \pm 19$	$^{14}634 \pm 16$
2	3DiVi	3divi	002	2017-10-20	190867	$^{49}4096 \pm 0$	$^{32}279 \pm 48$	Yes	$^{15}692 \pm 22$	$^{17}707 \pm 26$
3	AnyVision	anyvision	001	2018-01-31	662659	$^{20}1024 \pm 0$	$^{26}244 \pm 0$	No	$^{55}74285 \pm 212$	$^{55}74145 \pm 200$
4	AnyVision	anyvision	002	2018-01-31	662659	$^{19}1024 \pm 0$	$^{27}248 \pm 0$	No	$^{54}74069 \pm 188$	$^{54}74019 \pm 198$
5	Aware	aware	000	2017-10-27	240240	$^{28}1572 \pm 0$	$^{52}655 \pm 26$	No	$^{34}4030 \pm 83$	$^{33}2984 \pm 152$
6	Aware	aware	001	2017-10-27	240240	$^{29}1572 \pm 0$	$^{53}656 \pm 26$	No	$^{30}2902 \pm 51$	$^{31}2810 \pm 111$
7	Ayonix	ayonix	000	2017-06-22	58505	$^{23}1036 \pm 0$	$^{1}18 \pm 2$	No	$^{10}621 \pm 23$	$^{13}620 \pm 26$
8	Camvi Technologies	camvitech	001	2017-09-13	118759	$^{21}1024 \pm 0$	$^{21}181 \pm 6$	No	$^{6}481 \pm 16$	$^{9}487 \pm 23$
9	Gemalto Cogent	cogent	000	2018-01-12	633812	$^{11}439 \pm 0$	$^{43}511 \pm 7$	No	$^{39}9541 \pm 54$	$^{40}9500 \pm 51$
10	Cyberextruder	cyberex	001	2017-08-02	121211	$^{6}256 \pm 0$	$^{59}893 \pm 25$	No	$^{21}1083 \pm 16$	$^{23}1079 \pm 19$
11	Cyberextruder	cyberex	002	2018-01-30	168909	$^{33}2048 \pm 0$	$^{45}532 \pm 6$	No	$^{26}1803 \pm 14$	$^{27}1779 \pm 22$
12	Dermalog	dermalog	004	2017-10-26	0	$^{2}128 \pm 0$	$^{16}149 \pm 5$	No	$^{7}490 \pm 26$	$^{8}471 \pm 26$
13	Dermalog	dermalog	005	2018-02-02	0	$^{3}128 \pm 0$	$^{13}130 \pm 11$	No	$^{8}499 \pm 22$	$^{10}500 \pm 22$
14	Digital Barriers	barriers	000	2017-05-31	157794	$^{36}2056 \pm 0$	$^{8}104 \pm 0$	No	$^{41}13232 \pm 166$	$^{42}13226 \pm 146$
15	Digital Barriers	barriers	001	2017-07-20	236915	$^{38}2056 \pm 0$	$^{33}294 \pm 1$	No	$^{40}12311 \pm 164$	$^{41}12347 \pm 197$
16	Fudan University	fudan	000	2017-11-22	202296	$^{16}534 \pm 0$	$^{6}84 \pm 0$	Yes	$^{25}1713 \pm 28$	$^{26}1715 \pm 20$
17	Fudan University	fudan	001	2018-01-11	345583	$^{30}1808 \pm 0$	$^{14}131 \pm 0$	Yes	$^{23}1288 \pm 26$	$^{24}1281 \pm 22$
18	Gorilla Technology	gorilla	000	2018-02-13	98290	$^{53}4204 \pm 0$	$^{28}256 \pm 14$	No	$^{33}3875 \pm 29$	$^{35}3872 \pm 51$
19	ID3 Technology	id3	001	2017-08-04	225574	$^{14}520 \pm 0$	$^{24}238 \pm 19$	No	$^{20}1058 \pm 32$	$^{22}1049 \pm 28$
20	ID3 Technology	id3	002	2017-08-04	225574	$^{13}520 \pm 0$	$^{42}482 \pm 34$	No	$^{22}1100 \pm 59$	$^{21}1048 \pm 32$
21	Innovatrics	innova	001	2017-07-25	0	$^{9}284 \pm 0$	$^{50}645 \pm 5$	No	$^{37}5506 \pm 131$	$^{38}4975 \pm 308$
22	Innovatrics	innova	002	2018-02-05	0	$^{15}530 \pm 0$	$^{17}153 \pm 0$	Yes	$^{42}14384 \pm 260$	$^{43}14497 \pm 207$
23	Intellivision	intellivision	001	2017-10-10	43692	$^{39}2056 \pm 0$	$^{4}62 \pm 2$	No	$^{28}2573 \pm 91$	$^{30}2544 \pm 38$
24	Is It You	isityou	000	2017-06-26	48010	$^{58}19200 \pm 0$	$^{10}113 \pm 5$	No	$^{58}237517 \pm 1318$	$^{58}237374 \pm 1279$
25	Innovation Systems	isystems	000	2018-01-11	209270	$^{18}1024 \pm 0$	$^{18}154 \pm 4$	No	$^{3}398 \pm 14$	$^{4}396 \pm 16$
26	ITMO University	itmo	002	2017-08-07	1923215	$^{52}4162 \pm 0$	$^{48}611 \pm 17$	No	$^{38}7423 \pm 96$	$^{39}7451 \pm 94$
27	Megvii	megvii	000	2018-02-13	1358915	$^{35}2048 \pm 0$	$^{56}785 \pm 0$	No	$^{45}19816 \pm 158$	$^{47}19805 \pm 174$
28	Idemia	morpho	000	2017-07-11	100806	$^{1}116 \pm 0$	$^{9}109 \pm 1$	No	$^{19}993 \pm 31$	$^{20}1000 \pm 34$
29	Idemia	morpho	002	2017-12-28	383021	$^{10}352 \pm 0$	$^{41}411 \pm 7$	No	$^{32}3620 \pm 94$	$^{29}2481 \pm 65$
30	Neurotechnology	neurotech	002	2017-11-02	280771	$^{44}2718 \pm 0$	$^{47}581 \pm 6$	No	$^{51}68738 \pm 748$	$^{52}68905 \pm 993$
31	Neurotechnology	neurotech	003	2018-02-02	339962	$^{54}5214 \pm 0$	$^{46}580 \pm 5$	No	$^{53}73589 \pm 2548$	$^{53}72837 \pm 1546$
32	Noblis	noblis	000	2017-07-05	713573	$^{24}1061 \pm 0$	$^{19}174 \pm 20$	Yes	$^{24}1343 \pm 45$	$^{25}1373 \pm 56$
33	N-Tech Lab	ntech	002	2017-08-23	894169	$^{55}6744 \pm 1$	$^{37}330 \pm 11$	No	$^{56}82508 \pm 213$	$^{56}82524 \pm 220$
34	N-Tech Lab	ntech	003	2017-12-29	1422738	$^{45}2906 \pm 1$	$^{12}128 \pm 3$	Yes	$^{47}35327 \pm 117$	$^{48}35371 \pm 123$
35	Ping An Technology	Pingan	002	2017-09-25	442564	$^{57}18436 \pm 0$	$^{15}132 \pm 7$	Yes	$^{31}3088 \pm 32$	$^{34}3051 \pm 28$
36	Rank One Computing	rankone	002	2017-08-18	0	$^{4}224 \pm 0$	$^{5}75 \pm 1$	No	$^{52}69113 \pm 3802$	$^{2}369 \pm 26$
37	Rank One Computing	rankone	003	2017-11-28	0	$^{5}228 \pm 0$	$^{2}37 \pm 1$	No	$^{1}322 \pm 20$	$^{1}324 \pm 19$
38	Samtech InfoNet Limited	samtech	000	2017-05-02	109774	$^{37}2056 \pm 0$	$^{29}262 \pm 2$	No	$^{36}4550 \pm 26$	$^{37}4541 \pm 28$
39	Shaman Software	shaman	000	2017-12-05	0	$^{48}4096 \pm 0$	$^{51}653 \pm 16$	No	$^{2}380 \pm 25$	$^{3}379 \pm 31$
40	Shaman Software	shaman	001	2018-01-13	0	$^{46}4096 \pm 0$	$^{34}294 \pm 2$	No	$^{12}635 \pm 19$	$^{6}441 \pm 25$
41	Smilart	smilart	002	2018-02-06	111826	$^{22}1024 \pm 0$	$^{20}176 \pm 16$	No	$^{44}18784 \pm 136$	$^{46}18795 \pm 151$
42	Synesis	synesis	000	2018-02-09	340264	$^{12}512 \pm 0$	$^{11}117 \pm 0$	No	$^{16}775 \pm 206$	$^{11}568 \pm 48$
43	Tevian	tevian	001	2018-01-19	455746	$^{34}2048 \pm 0$	$^{25}242 \pm 14$	No	$^{4}430 \pm 38$	$^{5}427 \pm 23$
44	TongYi Transportation Technology	tongyi	001	2017-04-01	625339	$^{40}2058 \pm 0$	$^{36}310 \pm 20$	No	$^{43}17769 \pm 74$	$^{44}17750 \pm 63$
45	TongYi Transportation Technology	tongyi	002	2017-07-15	625336	$^{41}2058 \pm 0$	$^{39}356 \pm 35$	No	$^{46}29816 \pm 281$	$^{45}17799 \pm 127$
46	Toshiba	toshiba	000	2018-01-11	3893310	$^{31}1812 \pm 0$	$^{44}528 \pm 3$	Yes	$^{27}2255 \pm 60$	$^{28}2251 \pm 53$
47	Toshiba	toshiba	001	2018-01-11	3893310	$^{43}2580 \pm 0$	$^{49}615 \pm 3$	Yes	$^{29}2900 \pm 80$	$^{32}2881 \pm 60$
48	Ultinous	ultinous	000	2017-12-18	90803	$^{50}4100 \pm 0$	$^{3}53 \pm 0$	Yes	$^{35}4263 \pm 50$	$^{36}4262 \pm 42$
49	Ultinous	ultinous	001	2018-02-02	403572	$^{56}12304 \pm 0$	$^{22}200 \pm 0$	Yes	$^{49}36357 \pm 223$	$^{50}36332 \pm 210$
50	VCognition	vcog	002	2017-06-12	3229434	$^{59}61504 \pm 5$	$^{40}357 \pm 25$	No	$^{59}296154 \pm 3077$	$^{59}296436 \pm 4183$
51	Vigilant Solutions	vigilant	002	2017-09-28	344137	$^{42}2060 \pm 0$	$^{58}844 \pm 2$	No	$^{14}679 \pm 10$	$^{16}680 \pm 7$
52	Vigilant Solutions	vigilant	003	2018-01-23	343048	$^{27}1548 \pm 0$	$^{57}824 \pm 3$	No	$^{11}634 \pm 19$	$^{15}638 \pm 17$
53	VisionLabs	visionlabs	002	2017-09-08	591936	$^{8}264 \pm 0$	$^{30}265 \pm 8$	No	$^{17}796 \pm 35$	$^{18}799 \pm 31$
54	VisionLabs	visionlabs	003	2018-02-02	639849	$^{7}260 \pm 0$	$^{23}231 \pm 5$	No	$^{5}439 \pm 23$	$^{7}454 \pm 20$
55	Vocord	vocord	002	2017-06-07	918292	$^{25}1330 \pm 0$	$^{55}782 \pm 36$	Yes	$^{57}83063 \pm 517$	$^{57}83072 \pm 714$
56	Vocord	vocord	003	2018-01-23	900909	$^{17}817 \pm 0$	$^{35}296 \pm 3$	Yes	$^{50}46683 \pm 378$	$^{49}36007 \pm 292$
57	Zhuhai Yisheng Electronics Technology	yisheng	001	2017-08-17	120112	$^{32}2048 \pm 0$	$^{7}103 \pm 1$	Yes	$^{18}906 \pm 31$	$^{19}905 \pm 25$
58	Zhuhai Yisheng Electronics Technology	yisheng	002	2018-02-01	604097	$^{26}1348 \pm 0$	$^{38}332 \pm 1$	Yes	$^{9}613 \pm 40$	$^{12}588 \pm 26$
59	Shanghai Yitu Technology	yitu	000	2017-05-23	2211068	$^{51}4130 \pm 0$	$^{54}672 \pm 2$	No	$^{48}35352 \pm 114$	$^{51}37848 \pm 1773$

Notes	
1	The configuration size does not capture static data included in libraries. We do not count these because some algorithms include common ancilliary libraries for image processing (e.g. openCV) or numerical computation (e.g. blas).
2	The median template creation times are measured on Intel®Xeon®CPU E5-2630 v4 @ 2.20GHz processors or, for GPU-enabled implementations, NVidia Tesla K40.
3	The comparison durations, in nanoseconds, are estimated using std::chrono::high_resolution_clock which on the machine in (2) counts 1ns clock ticks. Precision is somewhat worse than that however. The ± value is the median absolute deviation times 1.48 for Normal consistency.

Table 1: Summary of algorithms and properties included in this report. The red superscripts give ranking for the quantity in that column.

	Algorithm Name	FALSE NON-MATCH RATE (FNMR)															
		CONSTRAINED, COOPERATIVE										LESS CONSTRAINED, NON-COOPERATIVE					
		VISA		VISA		VISAMC		MUGSHOT		WEBCAM		SELFIE		WILD		CHILD EXP	
	FMR	1E-06		0.0001		0.0001		0.0001		0.0001		0.0001		0.0001		0.01	
1	3divi-001	0.154	24	0.020	12	0.078	18	0.030	20	0.001	15	0.046	25	0.492	13	0.549	3
2	3divi-002	0.154	23	0.021	14	0.079	19	0.033	23	0.001	16	0.049	30	0.495	14	0.561	5
3	anyvision-001	0.086	14	0.044	30	0.072	16	0.034	24	0.012	34	0.075	41	0.684	30	-	56
4	anyvision-002	0.090	16	0.039	27	0.066	15	0.034	25	0.012	33	0.075	39	0.689	31	-	43
5	aware-000	0.309	42	0.054	35	0.197	44	0.045	39	0.003	22	0.046	27	0.948	53	0.808	34
6	aware-001	0.309	43	0.054	34	0.197	43	0.045	40	0.004	24	0.046	26	0.993	56	0.952	39
7	ayonix-000	0.487	54	0.230	58	0.435	57	0.309	56	0.172	49	0.360	50	0.807	44	0.843	36
8	camvi-001	0.538	56	0.183	56	0.382	55	0.323	57	0.106	47	0.268	49	0.743	37	0.763	26
9	cogent-000	0.084	13	0.012	5	0.043	10	0.018	4	0.001	13	0.014	10	0.508	16	0.759	24
10	cyberextruder-001	0.255	37	0.076	43	0.197	45	0.165	54	0.029	42	0.144	45	0.853	46	0.780	32
11	cyberextruder-002	0.134	20	0.027	17	0.081	21	0.042	35	0.010	29	0.023	16	0.486	12	-	54
12	dermalog-004	0.240	35	0.093	49	0.181	39	0.131	53	0.016	36	0.121	44	0.996	58	0.960	40
13	dermalog-005	0.182	28	0.066	39	0.153	33	0.072	47	0.027	41	0.072	38	0.519	17	-	51
14	digitalbarriers-000	0.463	53	0.161	54	0.368	54	0.184	55	0.045	46	0.170	46	0.741	36	0.771	28
15	digitalbarriers-001	0.502	55	0.155	53	0.468	59	0.041	33	0.029	43	0.115	42	0.678	29	0.731	16
16	fdu-000	0.058	8	0.016	8	0.039	9	0.021	10	0.001	12	0.009	4	0.549	18	0.706	13
17	fdu-001	0.044	6	0.017	9	0.035	7	0.024	15	-	58	-	58	0.344	3	0.657	9
18	gorilla-000	-	62	-	62	0.080	20	0.079	49	0.249	50	0.187	47	0.913	52	-	61
19	id3-001	0.250	36	0.063	38	0.170	37	0.036	26	0.002	19	0.040	21	0.765	39	0.688	12
20	id3-002	0.239	34	0.057	37	0.161	36	0.032	22	0.003	23	0.037	17	0.810	45	0.686	11
21	innovatrics-001	0.183	29	0.034	22	0.108	28	0.043	36	0.001	14	0.043	22	0.643	25	0.815	35
22	innovatrics-002	0.614	58	0.086	47	0.405	56	0.024	14	-	56	-	56	0.768	40	-	49
23	intellivision-001	0.221	33	0.042	28	0.133	32	0.037	27	0.021	40	0.066	36	0.878	48	0.777	29
24	isityou-000	0.703	59	0.414	59	0.568	60	0.680	60	0.690	52	-	53	1.000	59	-	42
25	isystems-000	0.179	27	0.043	29	0.127	31	0.064	46	0.008	27	0.049	29	0.678	28	0.736	17
26	itmo-002	0.287	40	0.050	33	0.161	35	0.047	41	0.036	45	0.069	37	0.892	50	0.783	33
27	itmo-003	0.140	22	0.020	11	0.075	17	0.023	13	-	54	-	54	0.450	10	-	45
28	megvii-000	0.025	1	0.005	1	0.016	1	0.017	1	0.000	6	0.006	3	0.746	38	-	57
29	morpho-000	0.134	21	0.026	16	0.102	26	0.028	19	0.007	26	0.012	7	0.893	51	0.846	37
30	morpho-002	0.068	11	0.009	3	0.038	8	0.019	7	0.000	5	0.006	2	0.607	21	0.631	8
31	neurotechnology-002	0.166	26	0.036	23	0.112	29	0.041	32	0.000	8	0.020	15	0.427	7	0.683	10
32	neurotechnology-003	0.088	15	0.018	10	0.062	13	0.027	17	0.000	2	0.017	12	0.421	6	-	46
33	noblis-000	0.542	57	0.212	57	0.442	58	0.349	58	0.153	48	0.239	48	0.875	47	0.757	23
34	ntechlab-002	0.065	9	0.021	13	0.056	12	0.023	12	0.003	20	0.014	9	0.324	2	0.459	2
35	ntechlab-003	0.039	5	0.011	4	0.029	5	0.019	6	0.003	21	0.014	11	0.271	1	0.433	1
36	pa-002	0.286	39	0.086	46	0.185	40	0.041	34	0.010	30	0.043	24	0.633	24	0.722	14
37	rankone-002	0.217	32	0.049	32	0.155	34	0.032	21	0.000	9	0.052	31	0.705	33	0.778	31
38	rankone-003	0.184	30	0.038	26	0.117	30	0.028	18	0.000	7	0.040	19	0.674	26	0.777	30
39	samtech-000	0.443	51	0.161	55	0.342	53	0.044	37	0.021	39	0.063	35	0.878	49	0.765	27
40	shaman-000	0.977	61	0.913	61	0.930	62	0.968	62	0.992	53	0.997	52	0.995	57	0.962	41
41	shaman-001	0.462	52	0.136	51	0.335	52	0.083	50	0.030	44	0.115	43	0.775	41	0.899	38
42	smilart-002	0.353	47	0.082	45	0.244	49	0.098	52	-	61	-	61	-	62	-	59
43	synesis-000	0.324	45	0.045	31	0.175	38	0.040	30	0.021	38	0.063	34	1.000	61	-	48
44	tevian-000	0.129	19	0.036	24	0.087	23	0.022	11	0.002	18	0.009	5	0.500	15	0.601	7
45	tevian-001	0.127	18	0.033	21	0.081	22	0.021	9	0.002	17	0.006	1	0.448	9	-	44
46	tiger-001	0.397	49	0.137	52	0.312	51	0.524	59	0.015	35	0.046	28	0.593	20	-	62
47	tongyitrans-001	0.072	12	0.038	25	0.062	14	0.041	31	0.009	28	0.063	32	0.704	32	0.743	19
48	tongyitrans-002	0.066	10	0.030	18	0.044	11	0.039	28	0.010	31	0.063	33	0.725	35	0.746	20
49	toshiba-000	0.290	41	0.069	41	0.191	42	0.039	29	-	60	-	60	0.719	34	0.743	18
50	toshiba-001	0.193	31	0.033	20	0.106	27	0.044	38	-	57	-	57	0.795	43	-	50
51	ultinous-000	0.348	46	0.076	44	0.234	48	0.073	48	0.005	25	0.040	20	0.979	55	-	55
52	ultinous-001	0.400	50	0.075	42	0.249	50	0.049	43	-	62	-	62	0.486	11	-	60
53	vcog-002	0.903	60	0.504	60	0.752	61	0.692	61	0.559	51	0.666	51	0.778	42	0.752	21
54	vigilantsolutions-002	0.321	44	0.087	48	0.229	47	0.092	51	0.018	37	0.075	40	0.675	27	0.722	15
55	vigilantsolutions-003	0.267	38	0.068	40	0.188	41	0.057	45	0.001	10	0.037	18	0.571	19	-	58
56	visionlabs-002	0.051	7	0.014	7	0.033	6	0.020	8	0.001	11	0.017	14	0.397	5	0.556	4
57	visionlabs-003	0.029	2	0.009	2	0.018	2	0.018	3	0.000	3	0.017	13	0.369	4	-	52
58	vocord-002	0.034	4	0.013	6	0.023	3	0.019	5	0.011	32	0.012	8	0.948	54	0.762	25
59	vocord-003	0.109	17	0.094	50	0.100	25	0.056	44	-	59	-	59	0.620	22	-	53
60	yisheng-001	0.160	25	0.032	19	0.089	24	0.048	42	0.000	4	0.043	23	0.628	23	0.756	22
61	yisheng-002	0.380	48	0.055	36	0.224	46	0.024	16	-	55	-	55	-	60	-	47
62	yitu-000	0.033	3	0.021	15	0.028	4	0.017	2	0.000	1	0.012	6	0.431	8	0.586	6

Table 2: FNMR is the proportion of mated comparisons below a threshold set to achieve the FMR given in the header on the fourth row. FMR is the proportion of impostor comparisons at or above that threshold. The green column applies to impostors of the same sex, age group, and country of birth. All other columns use fully zero effort impostors Note that the webcam and selfie values apply to images collected on the same day, and that will often yield optimisticly low FNMR values.

FNMR(T) "False non-match rate"
FMR(T) "False match rate"

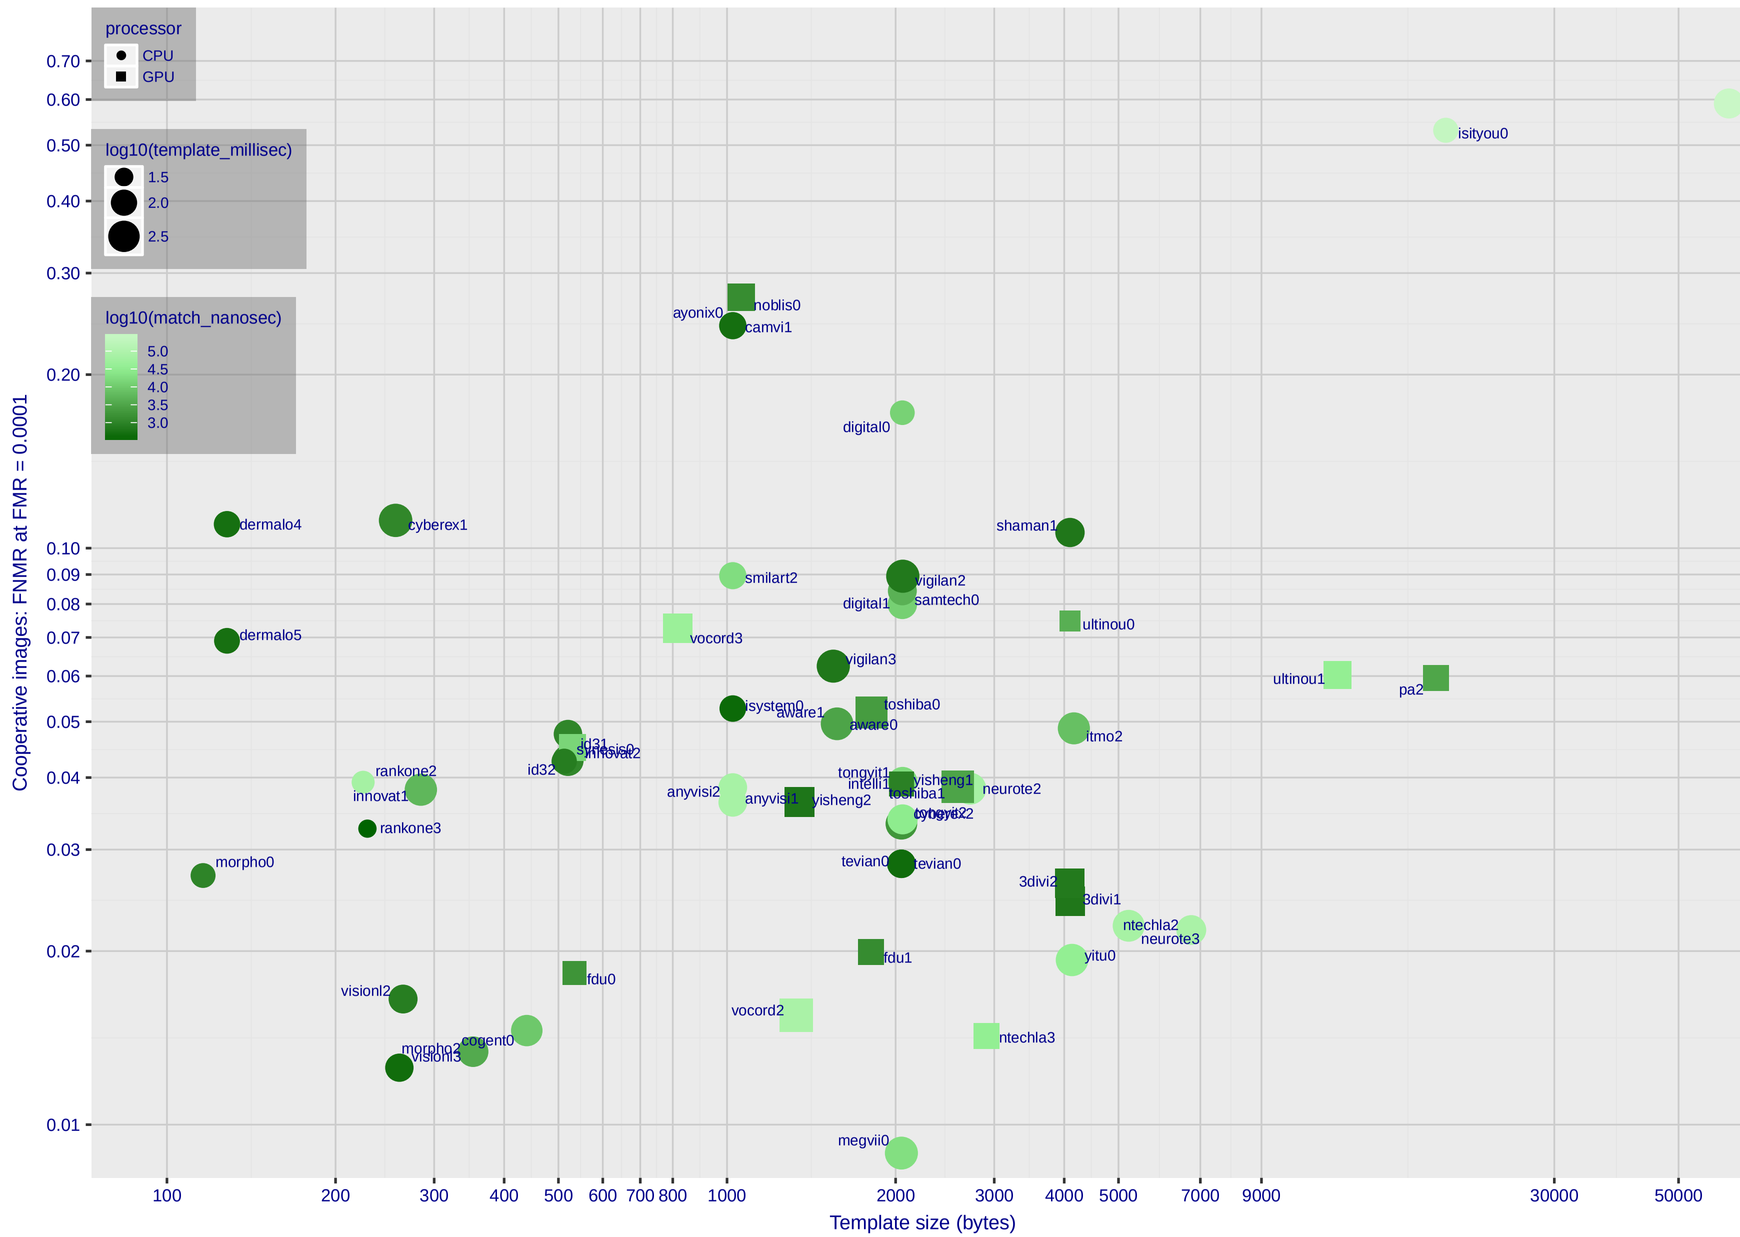

Figure 1: The points show false non-match rates (FNMR) versus the duration of the template generation operation. FNMR is the geometric mean of FNMR values for visa and mugshot images (from Figs. 4 and 6) at a false match rate (FMR) of 0.0001. Template generation time is a median estimated over 640 x 480 pixel portraits. The size of the points encodes template generation time - which spans one order of magnitude. The color of the points encodes one-to-one genuine template comparison duration - which span three orders of magnitude.

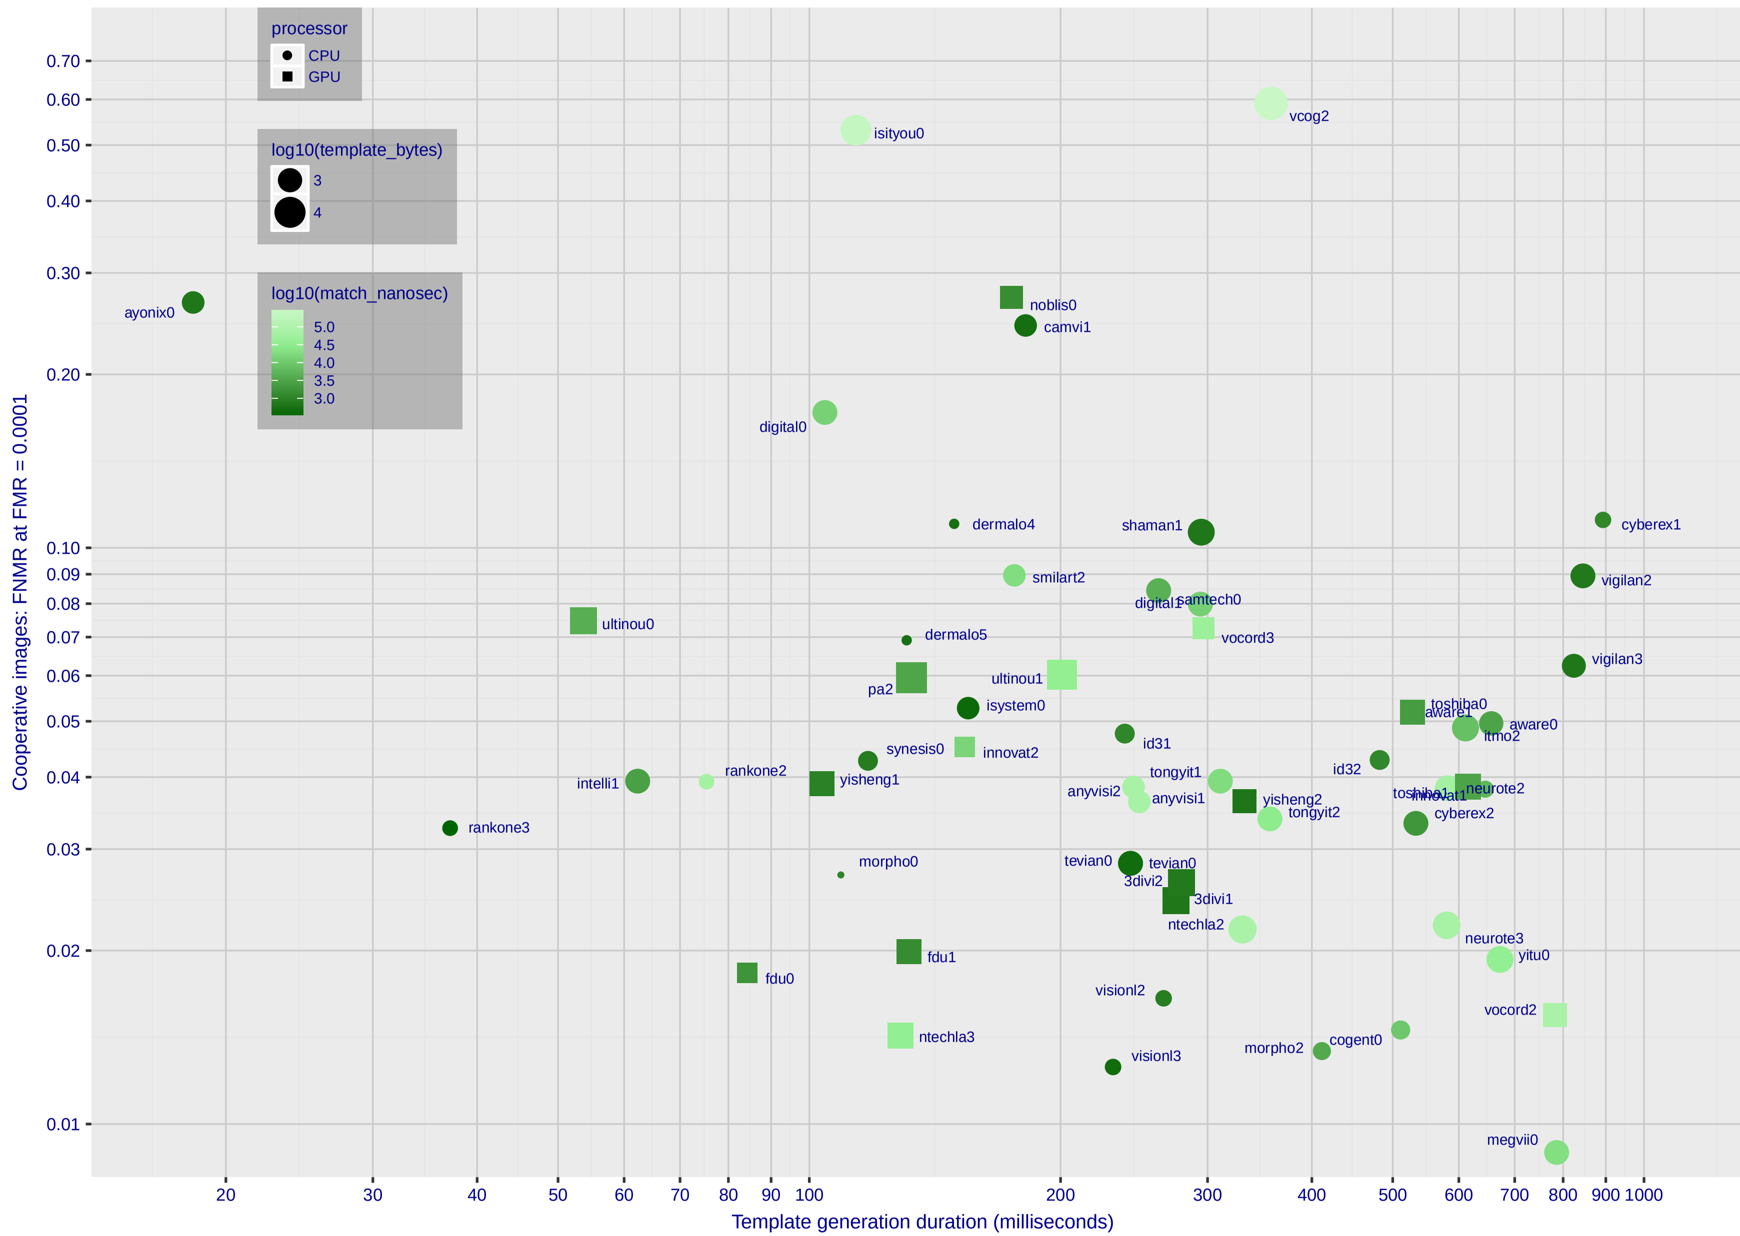

Figure 2: The points show false non-match rates (FNMR) versus the duration of the template generation operation. Some algorithms run on GPU, most on CPU - see Table 1. FNMR is the geometric mean of FNMR values for visa and mugshot images (from Figs. 4 and 6) at a false match rate (FMR) of 0.0001. Template generation time is a median estimated over 640 x 480 pixel portraits. The size of the points encodes template size - which span two orders of magnitude. The color of the points encodes one-to-one genuine template comparison duration - which span three orders of magnitude.

3 Metrics

3.1 Core accuracy

Given a vector of N genuine scores, u, the false non-match rate (FNMR) is computed as the proportion below some threshold, T:

$$\text{FNMR}(T) = 1 - \frac{1}{N}\sum_{i=1}^{N} H(u_i - T) \tag{1}$$

where $H(x)$ is the unit step function, and $H(0)$ taken to be 1.

Similarly, given a vector of N impostor scores, v, the false match rate (FMR) is computed as the proportion above T:

$$\text{FMR}(T) = \frac{1}{N}\sum_{i=1}^{N} H(v_i - T) \tag{2}$$

The threshold, T, can take on any value. We typically generate a set of thresholds from quantiles of the observed impostor scores, v, as follows. Given some interesting false match rate range, $[\text{FMR}_L, \text{FMR}_U]$, we form a vector of K thresholds corresponding to FMR measurements evenly spaced on a logarithmic scale

$$T_k = Q_v(1 - \text{FMR}_k) \tag{3}$$

where Q is the quantile function, and FMR_k comes from

$$\log_{10} \text{FMR}_k = \log_{10} \text{FMR}_L + \frac{k}{K}\ [\log_{10} \text{FMR}_U - \log_{10} \text{FMR}_L] \tag{4}$$

Error tradeoff characteristics are plots of FNMR(T) vs. FMR(T). These are plotted with $\text{FMR}_U \rightarrow 1$ and FMR_L as low as is sustained by the number of impostor comparisons, N. This is somewhat higher than the "rule of three" limit $3/N$ because samples are not independent, due to re-use of images.

4 Datasets

4.1 Child exploitation images

▷ The number of images is $O(10^4)$.

▷ The number of subjects is $O(10^3)$.

▷ The number of subjects with two images $O(10^3)$.

▷ The images are operational. They are taken from ongoing investigations of child exploitation crimes. The images are arbitrarily unconstrained. Pose varies considerably around all three axes, including subject lying down. Resolution varies very widely. Faces can be occluded by other objects, including hair and hands. Lighting varies, although the images are intended for human viewing. Mis-focus is rare. Images are given to the algorithm without any cropping; faces may occupy widely varying areas.

▷ The images are usually large from contemporary cameras. The mean interocular distance (IOD) is 70 pixels.

▷ The images are of subjects from several countries, due to the global production of this imagery.

▷ The images are of children, from infancy to late adolescence.

▷ All of the images are live capture, none are scanned. Many have been cropped.

▷ When these images are input to the algorithm, they are labelled as being of type "EXPLOITATION" - see Table 4 of the FRVT API.

4.2 Visa images

▷ The number of images is $O(10^5)$.

▷ The number of subjects is $O(10^5)$.

▷ The number of subjects with two images $O(10^4)$.

▷ The images have geometry in reasonable conformance with the ISO/IEC 19794-5 Full Frontal image type. Pose is generally excellent.

▷ The images are of size 252x300 pixels. The mean interocular distance (IOD) is 69 pixels.

▷ The images are of subjects from greater than 100 countries, with significant imbalance due to visa issuance patterns.

▷ The images are of subjects of all ages, including children, again with imbalance due to visa issuance demand.

▷ Many of the images are live capture. A substantial number of the images are photographs of paper photographs.

▷ When these images are input to the algorithm, they are labelled as being of type "ISO" - see Table 4 of the FRVT API.

4.3 Mugshot images

▷ The number of images is $O(10^6)$.

▷ The number of subjects is $O(10^5)$.

▷ The number of subjects with two images $O(10^5)$.

- ▷ The images have geometry in reasonable conformance with the ISO/IEC 19794-5 Full Frontal image type.
- ▷ The images are of variable sizes. The median IOD is 104 pixels. The mean IOD is 123 pixels.
- ▷ The images are of subjects from the United States.
- ▷ The images are of adults.
- ▷ The images are all live capture.
- ▷ When these images are input to the algorithm, they are labelled as being of type ”mugshot” - see Table 4 of the FRVT API.

4.4 Selfie images

- ▷ The number of images is below 500.
- ▷ The number of subjects is below 500.
- ▷ All subjects have a selfie image, and a portrait image.
- ▷ The portrait images are in reasonable conformance with the ISO/IEC 19794-5 Full Frontal image type.
- ▷ The selfie images vary: taken with camera above and below eye level, with one hand or two hands. Pitch angles vary more than yaw angles, which are frontal. Some perspective distortion is evident.
- ▷ The images have mean IOD of 140 pixels.
- ▷ The images are of subjects from the United States.
- ▷ The images are of adults.
- ▷ The images are all live capture.
- ▷ When these images are input to the algorithm, they are labelled as being of type ”WILD” - see Table 4 of the FRVT API.

4.5 Webcam images

- ▷ The number of images is below 1500.
- ▷ The number of subjects is below 1500.
- ▷ All subjects have a webcam image, and a portrait image.
- ▷ The portrait images are in reasonable conformance with the ISO/IEC 19794-5 Full Frontal image type.
- ▷ The webcam images are taken with camera at a typical head height, with mild pitch angles, low yaw angles, but some variation in range, such that low perspective distortion is sometimes evident.
- ▷ The images have mean IOD of 68 pixels (sd=12).
- ▷ The images are of subjects from the United States.
- ▷ The images are of adults.
- ▷ The images are all live capture.
- ▷ When these images are input to the algorithm, they are labelled as being of type ”MUGSHOT” - see Table 4 of the FRVT API.

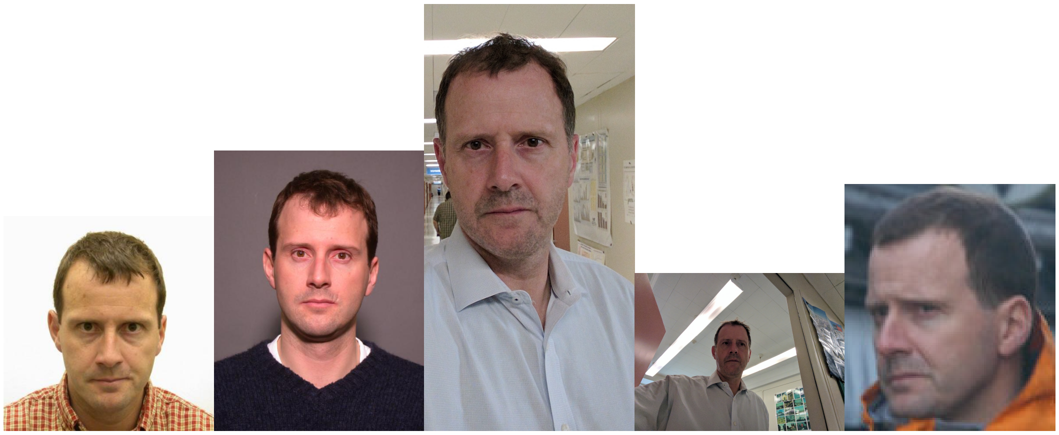

(a) Visa (b) Mugshot (c) Webcam (d) Selfie (e) Wild

Figure 3: The figure gives simulated samples of image types used in this report.

4.6 Wild images

- ▷ The number of images is O(10^5).
- ▷ The number of subjects is O(10^3).
- ▷ The number of subjects with two images O(10^3).
- ▷ The images include many photojournalism-style images. Images are given to the algorithm using a variable but generally tight crop of the head. Resolution varies very widely. The images are very unconstrained, with wide yaw and pitch pose variation. Faces can be occluded, including hair and hands.
- ▷ The images are of adults.
- ▷ All of the images are live capture, none are scanned.
- ▷ When these images are input to the algorithm, they are labelled as being of type ”WILD” - see Table 4 of the FRVT API.

5 Results

5.1 Test goals

- ▷ To state overall accuracy.
- ▷ To compare algorithms.

5.2 Test design

Method: For wild images:

- ▷ The comparisons are of wild photos against wild photos.
- ▷ The number of genuine comparisons is O(10^6).
- ▷ The number of impostor comparisons is O(10^7).

- ▷ The comparisons are fully zero-effort, meaning impostors are paired without attention to sex, age or other covariates.
- ▷ The number of persons is O(10^4).
- ▷ The number of images used to make 1 template is 1.
- ▷ The number of templates used to make each comparison score is two corresponding to simple one-to-one verification.

Method: For visa images:

- ▷ The comparisons are of visa photos against visa photos.
- ▷ The number of genuine comparisons is O(10^4).
- ▷ The number of impostor comparisons is O(10^{10}).
- ▷ The comparisons are fully zero-effort, meaning impostors are paired without attention to sex, age or other covariates. However, later analysis is conducted on subsets.
- ▷ The number of persons is O(10^5).
- ▷ The number of images used to make 1 template is 1.
- ▷ The number of templates used to make each comparison score is two corresponding to simple one-to-one verification.

For mugshot images:

- ▷ The comparisons are of mugshot photos against mugshot photos.
- ▷ The number of genuine comparisons is O(10^5).
- ▷ The number of impostor comparisons is O(10^7).
- ▷ The comparisons are fully zero-effort, meaning impostors are paired without attention to sex, age or other covariates.
- ▷ The number of persons is O(10^6).
- ▷ The number of images used to make 1 template is 1.
- ▷ The number of templates used to make each comparison score is two corresponding to simple one-to-one verification.

For selfie images:

- ▷ The comparisons are of selfie photos against portrait photos.
- ▷ The number of genuine comparisons is O(10^2).
- ▷ The number of impostor comparisons is O(10^8) selfies are compared with portraits of O(10^6) other subjects.
- ▷ The comparisons are fully zero-effort, meaning impostors are paired without attention to sex, age or other covariates.
- ▷ The number of persons is O(10^6).
- ▷ The number of images used to make 1 template is 1.

- ▷ The number of templates used to make each comparison score is two corresponding to simple one-to-one verification.

For webcam images:

- ▷ The comparisons are of webcam photos against portrait photos.
- ▷ The number of genuine comparisons is O(10^3).
- ▷ The number of impostor comparisons is O(10^9) webcams are compared with portraits of O(10^6) other subjects.
- ▷ The comparisons are fully zero-effort, meaning impostors are paired without attention to sex, age or other covariates.
- ▷ The number of persons is O(10^6).
- ▷ The number of images used to make 1 template is 1.
- ▷ The number of templates used to make each comparison score is two corresponding to simple one-to-one verification.

For child exploitation images:

- ▷ The comparisons are of unconstrained child exploitation photos against others of the same type.
- ▷ The number of genuine comparisons is O(10^4).
- ▷ The number of impostor comparisons is O(10^7).
- ▷ The comparisons are fully zero-effort, meaning impostors are paired without attention to sex, age or other covariates.
- ▷ The number of persons is O(10^3).
- ▷ The number of images used to make 1 template is 1.
- ▷ The number of templates used to make each comparison score is two corresponding to simple one-to-one verification.
- ▷ We produce two performance statements. First, is a DET as used for visa and mugshot images. The second is a cumulative match characteristic (CMC) summarizing a simulated one-to-many search process. This is done as follows.
  - We regard M enrollment templates as items in a gallery.
  - These M templates come from $M > N$ individuals, because multiple images of a subject are present in the gallery under separate identifiers.
  - We regard the verification templates as search templates.
  - For each search we compute the rank of the highest scoring mate.
  - This process should properly be conducted with a 1:N algorithm, such as those tested in NIST IR 8009. We use the 1:1 algorithms in a simulated 1:N mode here to a) better reflect what a child exploitation analyst does, and b) to do show algorithm efficacy is better than that revealed in the verification DETs.

5.3 Failure to enrol

Algorithm	Failure to Enrol Rate[1]											
Name	CHILD-EXPLOIT		MUGSHOT		SELFIES		VISA		WEBCAM		WILD	
3divi-001	0.2006	20	0.0019	37	0.0202	33	0.0007	31	0.0020	23	0.1898	39
3divi-002	0.3103	26	0.0039	51	0.0376	41	0.0010	34	0.0027	33	0.2413	45
anyvision-001	-	59	0.0070	60	0.0578	47	0.0091	62	0.0089	41	0.5569	61
anyvision-002	-	59	0.0070	59	0.0578	46	0.0090	61	0.0095	42	0.5575	62
aware-000	0.4663	38	0.0018	36	0.0231	35	0.0005	23	0.0055	37	0.2444	47
aware-001	0.1897	19	0.0010	22	0.0145	31	0.0002	11	0.0048	36	0.1431	26
ayonix-000	0.0000	3	0.0109	62	0.0751	50	0.0137	64	0.0109	44	0.0000	2
camvi-001	0.3931	36	0.0033	49	0.0231	36	0.0010	37	0.0027	34	0.3767	53
cogent-000	0.2914	24	0.0011	25	0.0029	14	0.0005	26	0.0020	26	0.2200	43
cyberextruder-001	0.5338	43	0.0036	50	0.0376	40	0.0029	50	0.0205	49	0.5730	63
cyberextruder-002	-	59	0.0026	46	0.0058	20	0.0028	49	0.0123	46	0.1164	24
dermalog-004	0.3110	27	0.0031	48	0.0087	27	0.0090	60	0.0020	31	0.2050	41
dermalog-005	-	59	0.0016	35	0.0000	8	0.0041	56	0.0014	20	0.1498	28
digitalbarriers-000	0.5469	44	0.0043	53	0.0925	51	0.0019	45	0.0184	48	0.5170	60
digitalbarriers-001	0.5102	42	0.0044	54	0.0925	52	0.0018	44	0.0232	50	0.4593	58
fdu-000	0.4992	41	0.0025	44	-	59	0.0011	40	0.0020	28	0.3437	52
fdu-001	0.1380	18	0.0015	30	-	59	0.0009	33	-	59	0.0674	20
gorilla-000	-	59	0.0009	19	0.0029	19	0.0004	21	0.0007	15	0.1644	32
id3-001	0.3411	31	0.0043	52	0.0260	38	0.0043	59	0.0014	18	0.2090	42
id3-002	0.3168	28	0.0030	47	0.0202	34	0.0032	52	0.0020	25	0.1496	27
innovatrics-001	0.3392	30	0.0013	29	0.0087	25	0.0004	20	0.0027	32	0.3337	51
innovatrics-002	-	59	0.0003	13	-	59	0.0004	17	-	59	0.2435	46
intellivision-001	0.5495	46	0.0052	55	0.0491	44	0.0042	57	0.0252	51	0.7025	66
isityou-000	0.4714	39	0.0022	41	0.0665	49	0.0010	36	0.0116	45	0.4586	57
isystems-000	0.4757	40	0.0025	45	0.0260	37	0.0010	35	0.0068	40	0.4544	56
itmo-002	0.5751	47	0.0068	58	0.0636	48	0.0029	51	0.0498	53	0.6985	64
megvii-000	-	59	0.0007	15	0.0000	13	0.0003	12	0.0007	16	0.0258	15
morpho-000	0.0000	8	0.0000	7	0.0000	12	0.0000	8	0.0000	9	0.0000	12
morpho-002	0.0572	14	0.0009	20	0.0000	4	0.0004	18	0.0007	10	0.0238	14
neurotechnology-002	0.2043	21	0.0000	8	0.0058	22	0.0000	6	0.0014	21	0.1525	29
neurotechnology-003	-	59	0.0011	24	0.0058	23	0.0005	22	0.0020	30	0.1698	34
noblis-000	0.0000	5	0.0000	6	0.0000	10	0.0000	7	0.0000	7	0.0000	10
ntechlab-002	0.0926	16	0.0009	17	0.0029	17	0.0005	24	0.0007	11	0.0584	18
ntechlab-003	0.0926	17	0.0009	18	0.0029	18	0.0005	25	0.0007	14	0.0584	19
pa-002	0.0000	1	0.0000	4	0.0000	9	0.0000	4	0.0000	6	0.0000	7
rankone-002	0.0009	10	0.0001	10	0.0000	11	0.0000	10	0.0000	8	0.0855	22
rankone-003	0.0009	9	0.0001	9	0.0000	5	0.0000	9	0.0000	4	0.0855	21
samtech-000	0.5474	45	0.0052	56	0.0491	45	0.0042	58	0.0252	52	0.7023	65
shaman-000	0.0000	4	0.0000	2	0.0000	3	0.0000	2	0.0000	3	0.0000	3
shaman-001	0.0000	2	0.0000	1	0.0000	1	0.0000	1	0.0000	1	0.0000	1
smilart-002	-	59	0.0012	28	-	59	0.0011	39	-	59	-	59
synesis-000	-	59	0.0022	42	0.0347	39	0.0007	30	0.0095	43	0.4111	54
tevian-000	0.3373	29	0.0011	27	0.0000	7	0.0012	42	0.0020	27	0.2661	48
tongyitrans-001	0.0000	7	0.0068	57	0.0462	43	0.0040	55	0.0055	39	0.0000	11
tongyitrans-002	0.3609	34	0.0078	61	0.0462	42	0.0040	54	0.0055	38	0.0000	5
toshiba-000	0.0000	6	0.0000	3	-	59	0.0000	3	-	59	0.0000	4
toshiba-001	-	59	0.0000	5	-	59	0.0000	5	-	59	0.0000	9
ultinous-000	-	59	0.0002	11	0.0000	2	0.0003	14	0.0000	2	0.0497	16
ultinous-001	-	59	0.0002	12	-	59	0.0003	15	-	59	0.0497	17
vcog-002	0.2209	22	0.0021	39	0.0087	28	0.0019	46	0.0007	13	0.1672	33
vigilantsolutions-002	0.3585	33	0.0010	23	0.0116	30	0.0004	19	0.0048	35	0.3262	50
vigilantsolutions-003	-	59	0.0008	16	0.0058	21	0.0004	16	0.0014	19	0.1336	25
visionlabs-002	0.3063	25	0.0021	40	0.0087	26	0.0026	48	0.0007	12	0.1865	38
visionlabs-003	-	59	0.0020	38	0.0087	29	0.0025	47	0.0014	22	0.1580	30
vocord-002	0.3782	35	0.0015	31	0.0029	16	0.0037	53	0.0171	47	0.1992	40
vocord-003	-	59	0.0256	63	-	59	0.0776	66	-	59	0.4808	59
yisheng-001	0.4277	37	0.0016	34	0.0173	32	0.0006	28	0.0020	29	0.1770	36
yisheng-002	-	59	0.0016	33	-	59	0.0006	27	-	59	0.1770	35
yitu-000	0.3475	32	0.0015	32	0.0029	15	0.0013	43	0.0014	17	0.1591	31

Table 3: FTE is the proportion of failed template generation attempts. Failures can occur because the software throws an exception, or because the software electively refuses to process the input image. This would typically occur if a face is not detected. FTE is measured as the number of function calls that give a non-zero error code, OR that give a "small" template. This is defined as one whose size is less than 0.3 times the median template size. This second rule is needed because some algorithms incorrectly fail to return a non-zero error code when template generation fails.

[1] The effects of FTE are included in the accuracy results of this report by regarding any template comparison involving a failed template enrollment to produce a low similarity score. Thus higher FTE results in higher FNMR.

5.4 Recognition accuracy

Core algorithm accuracy is stated via:

- ▷ **Cooperative subjects**
  - The summary table of Figure 2;
  - The visa image DETs of Figures 4 and 5;
  - The mugshot DETs of Figure 6 ;
  - The selfie-portrait DETs of Figure 7;
  - The webcam-portrait DETs of Figure 8;
- ▷ **Non-cooperative subjects**
  - The photojournalism DET of Figure 9
  - The sensitivity of photojournalism FNMR to relative yaw angles in Figure 10.
  - The sensitivity of photojournalism FMR to relative yaw angles in Figure 11.
  - The child-exploitation DET of Figure 12;
  - The child-exploitation CMC of Figure 13.

Figure 18 shows dependence of false match rate on algorithm score threshold. This allows a deployer to set a threshold to target a particular false match rate appropriate to the security objectives of the application.

Figure 20 likewise shows FMR(T) but for mugshots, and specially four subsets of the population.

Note that in both the mugshot and visa sets false match rates vary with the ethnicity, age, and sex, of the enrollee and impostor - see section 5.6.

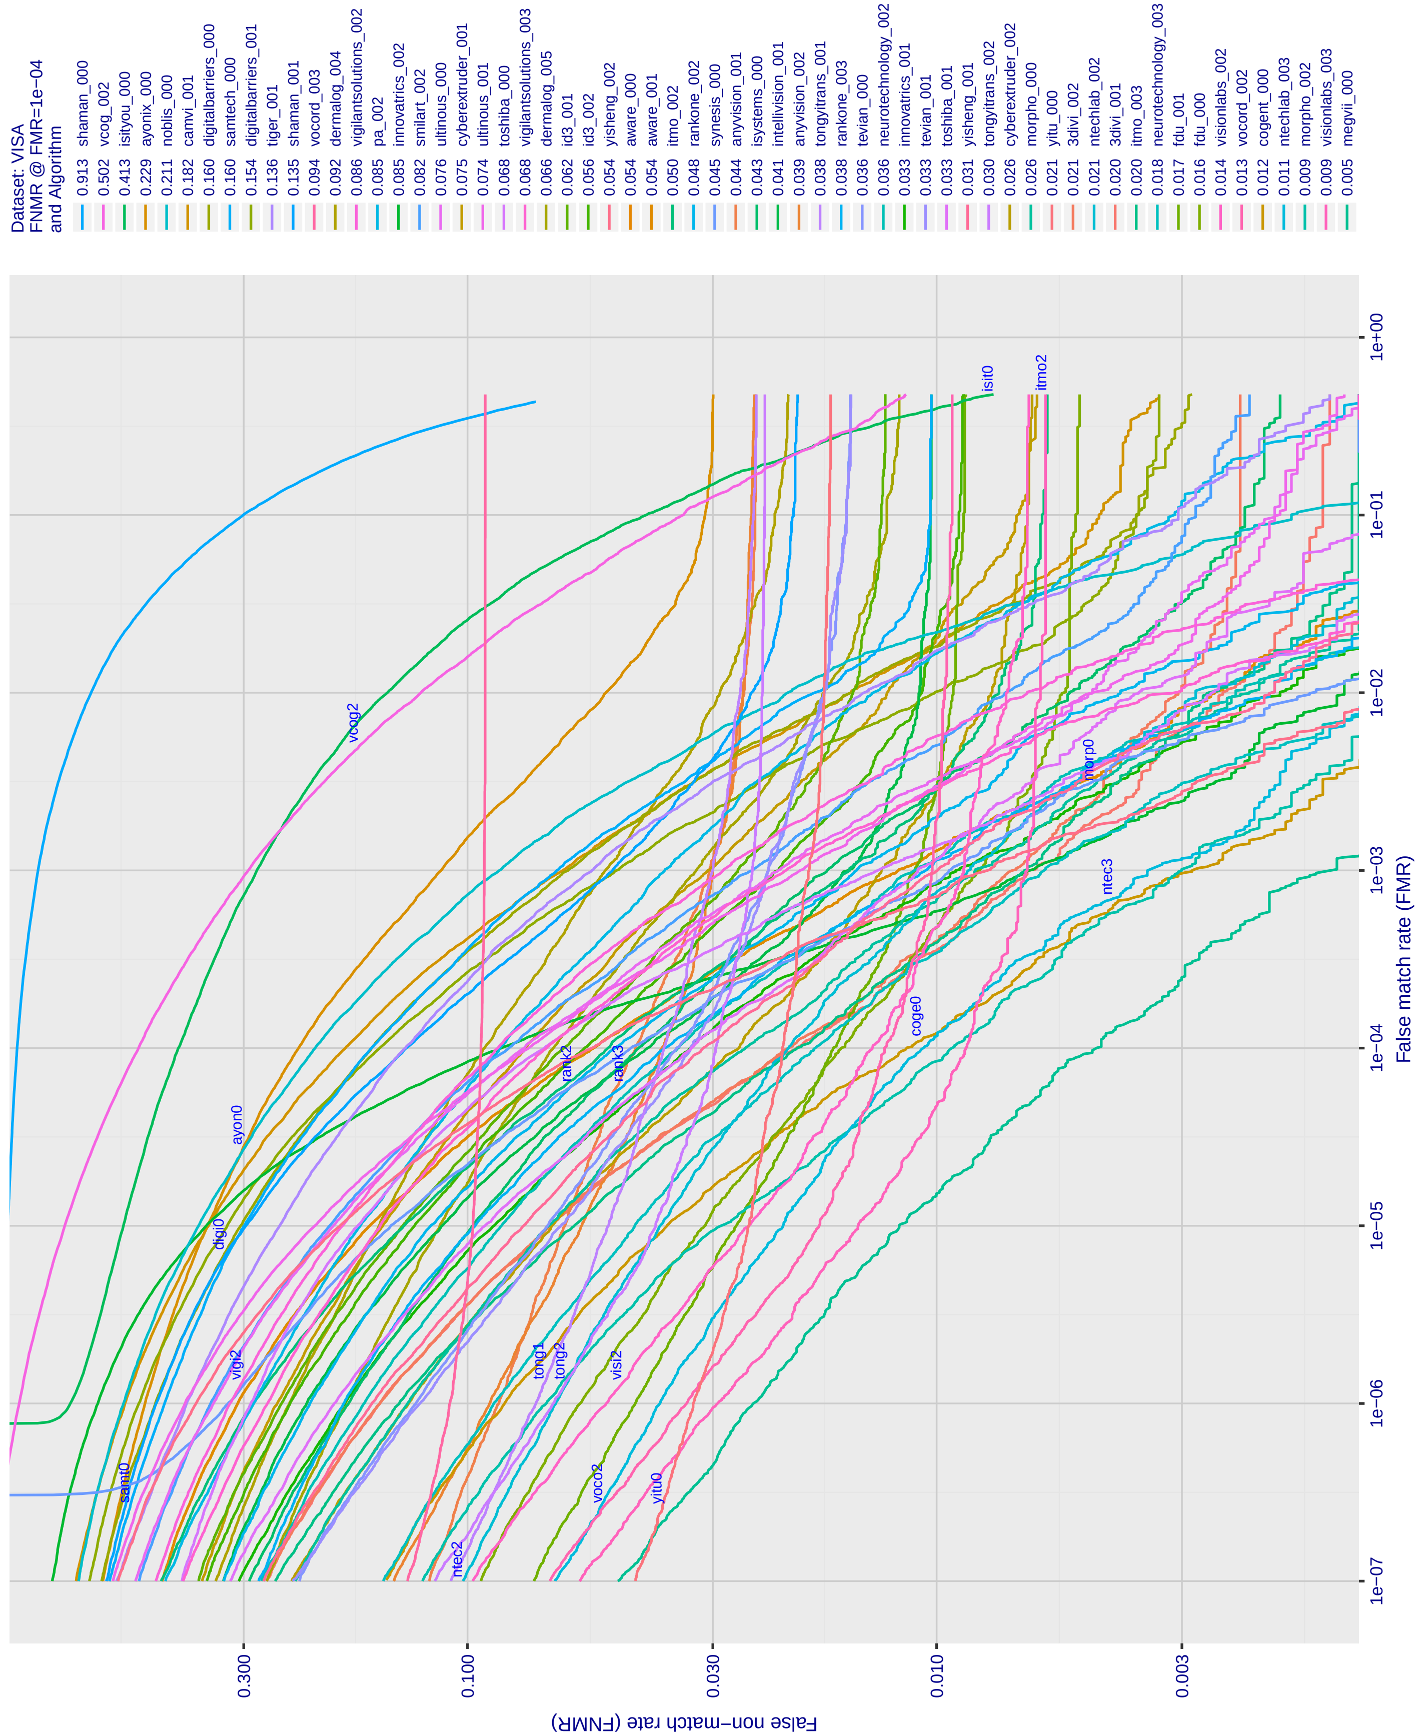

Figure 4: For the visa images, detection error tradeoff (DET) characteristics showing false non-match rate vs. false match rate plotted parametrically on threshold, T. The scales are logarithmic in order to show many decades of FMR.

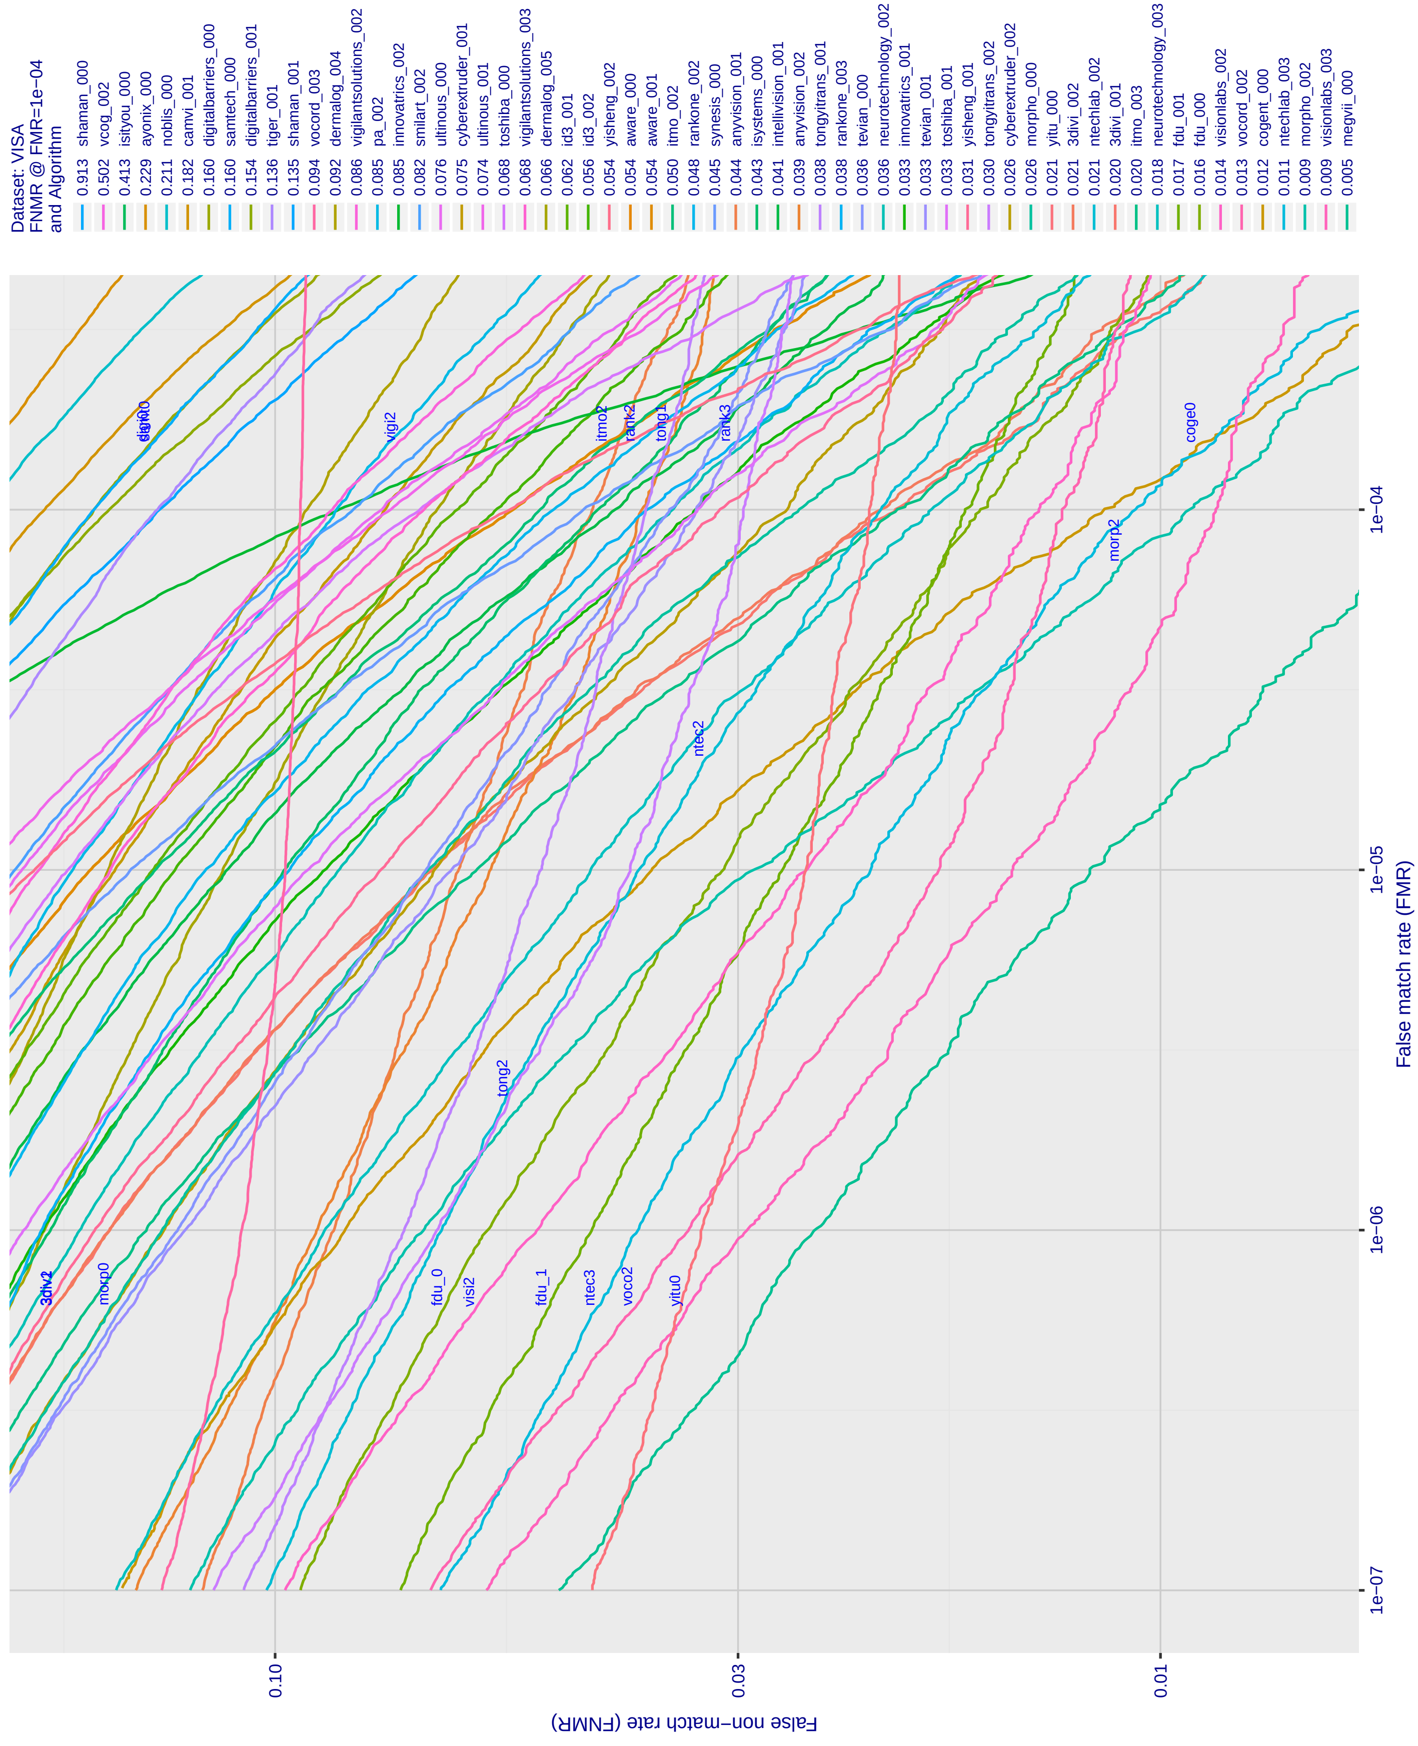

Figure 5: For the visa images, *but now just for low* FMR*, detection error tradeoff* (DET) *characteristics showing false non-match rate vs. false match rate plotted parametrically on threshold, T. The scales are logarithmic in order to show many decades of* FMR.

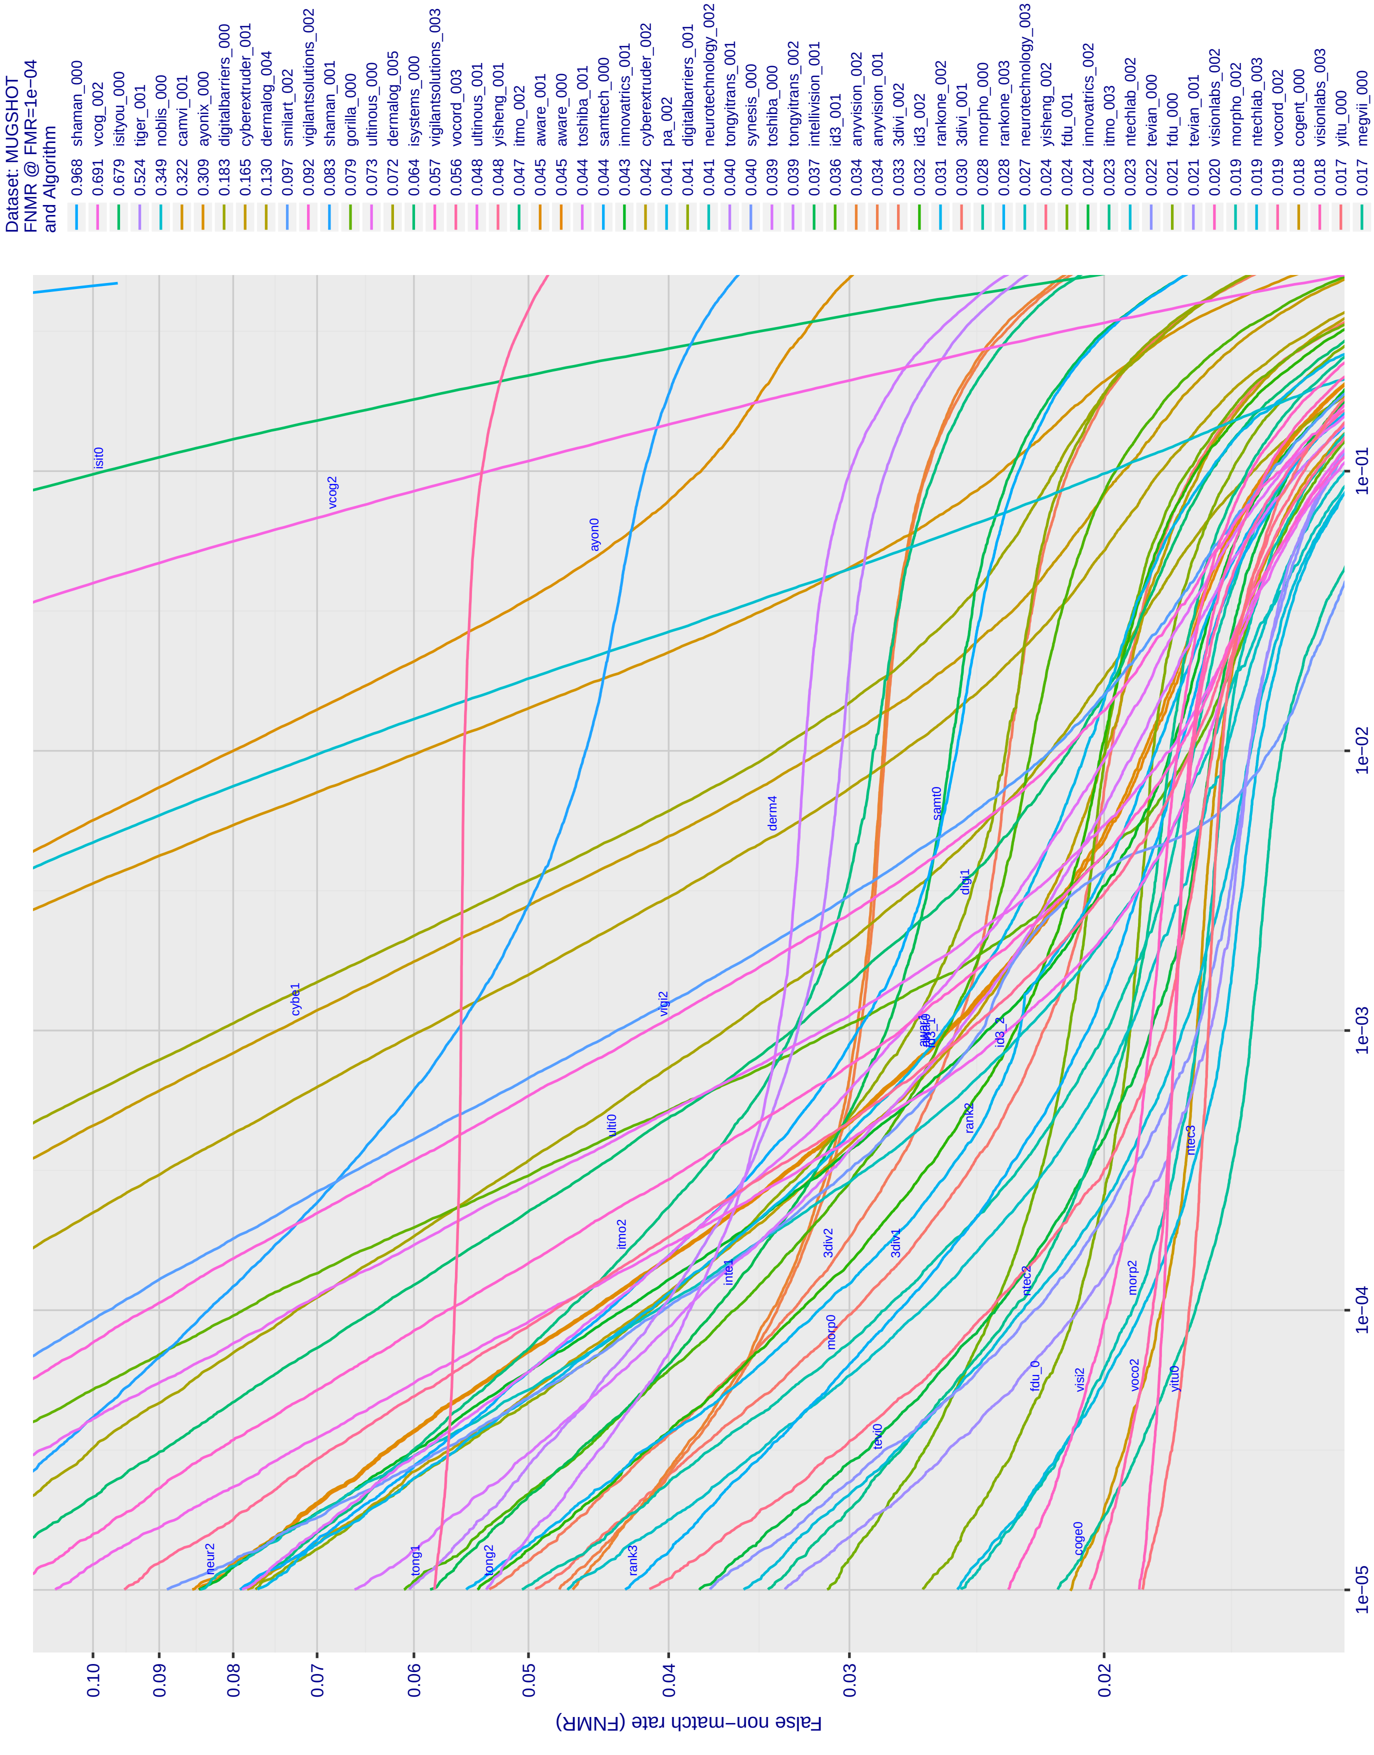

Figure 6: For the mugshot images, detection error tradeoff (DET) *characteristics showing false non-match rate vs. false match rate plotted parametrically on threshold, T. The scales are logarithmic in order to show decades of* FMR.

FNMR(T) "False non-match rate"
FMR(T) "False match rate"

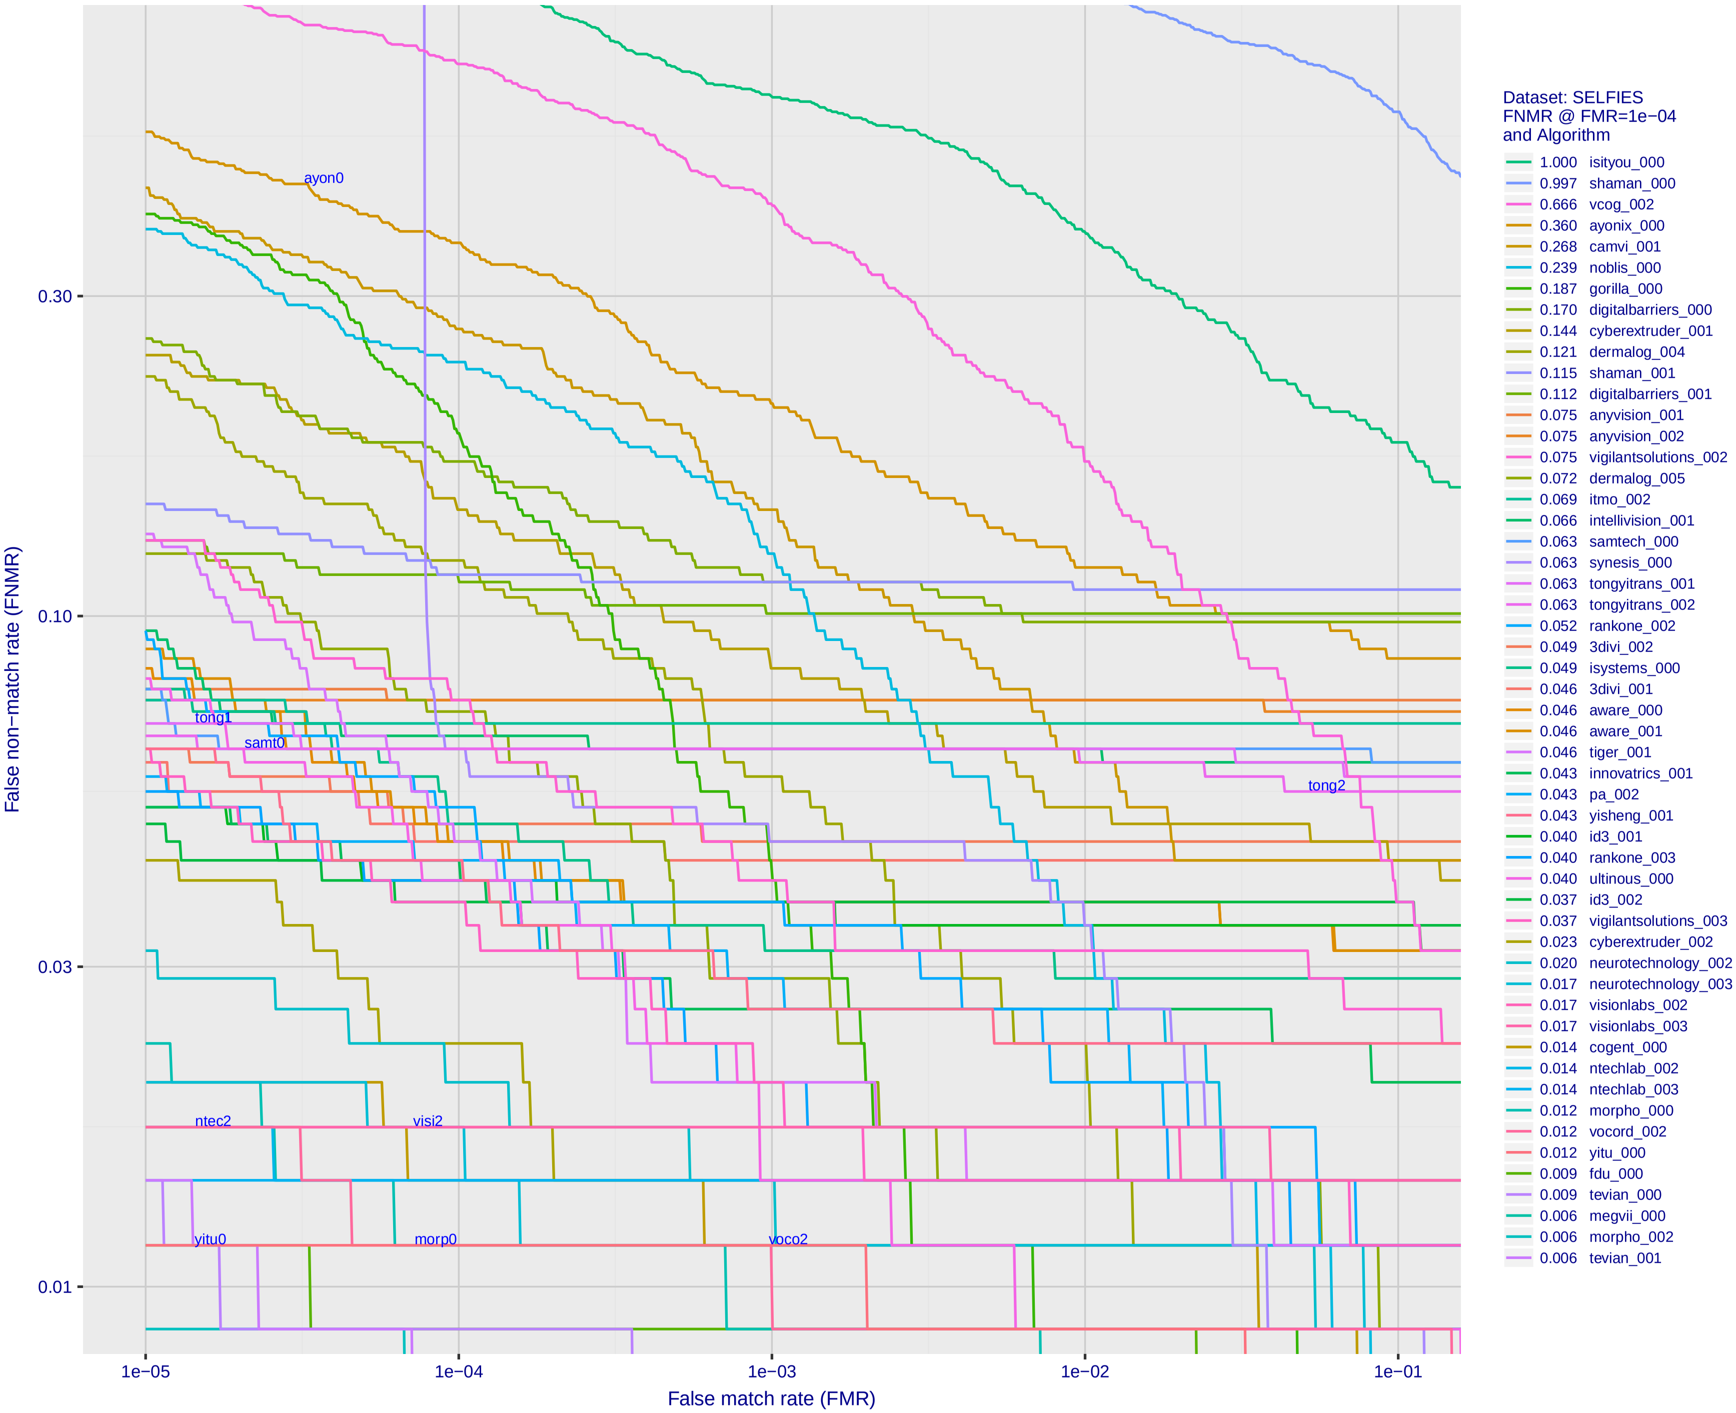

Figure 7: For the selfie-to-portrait comparisons, detection error tradeoff (DET) characteristics showing false non-match rate vs. false match rate plotted parametrically on threshold, T. The scales are logarithmic in order to show several decades of FMR. Caution: The FNMR values here are optimistic statements of accuracy because the image pairs were collected on the same day. This is known across biometrics to give better accuracy, and is operationally relevant only in special cases.

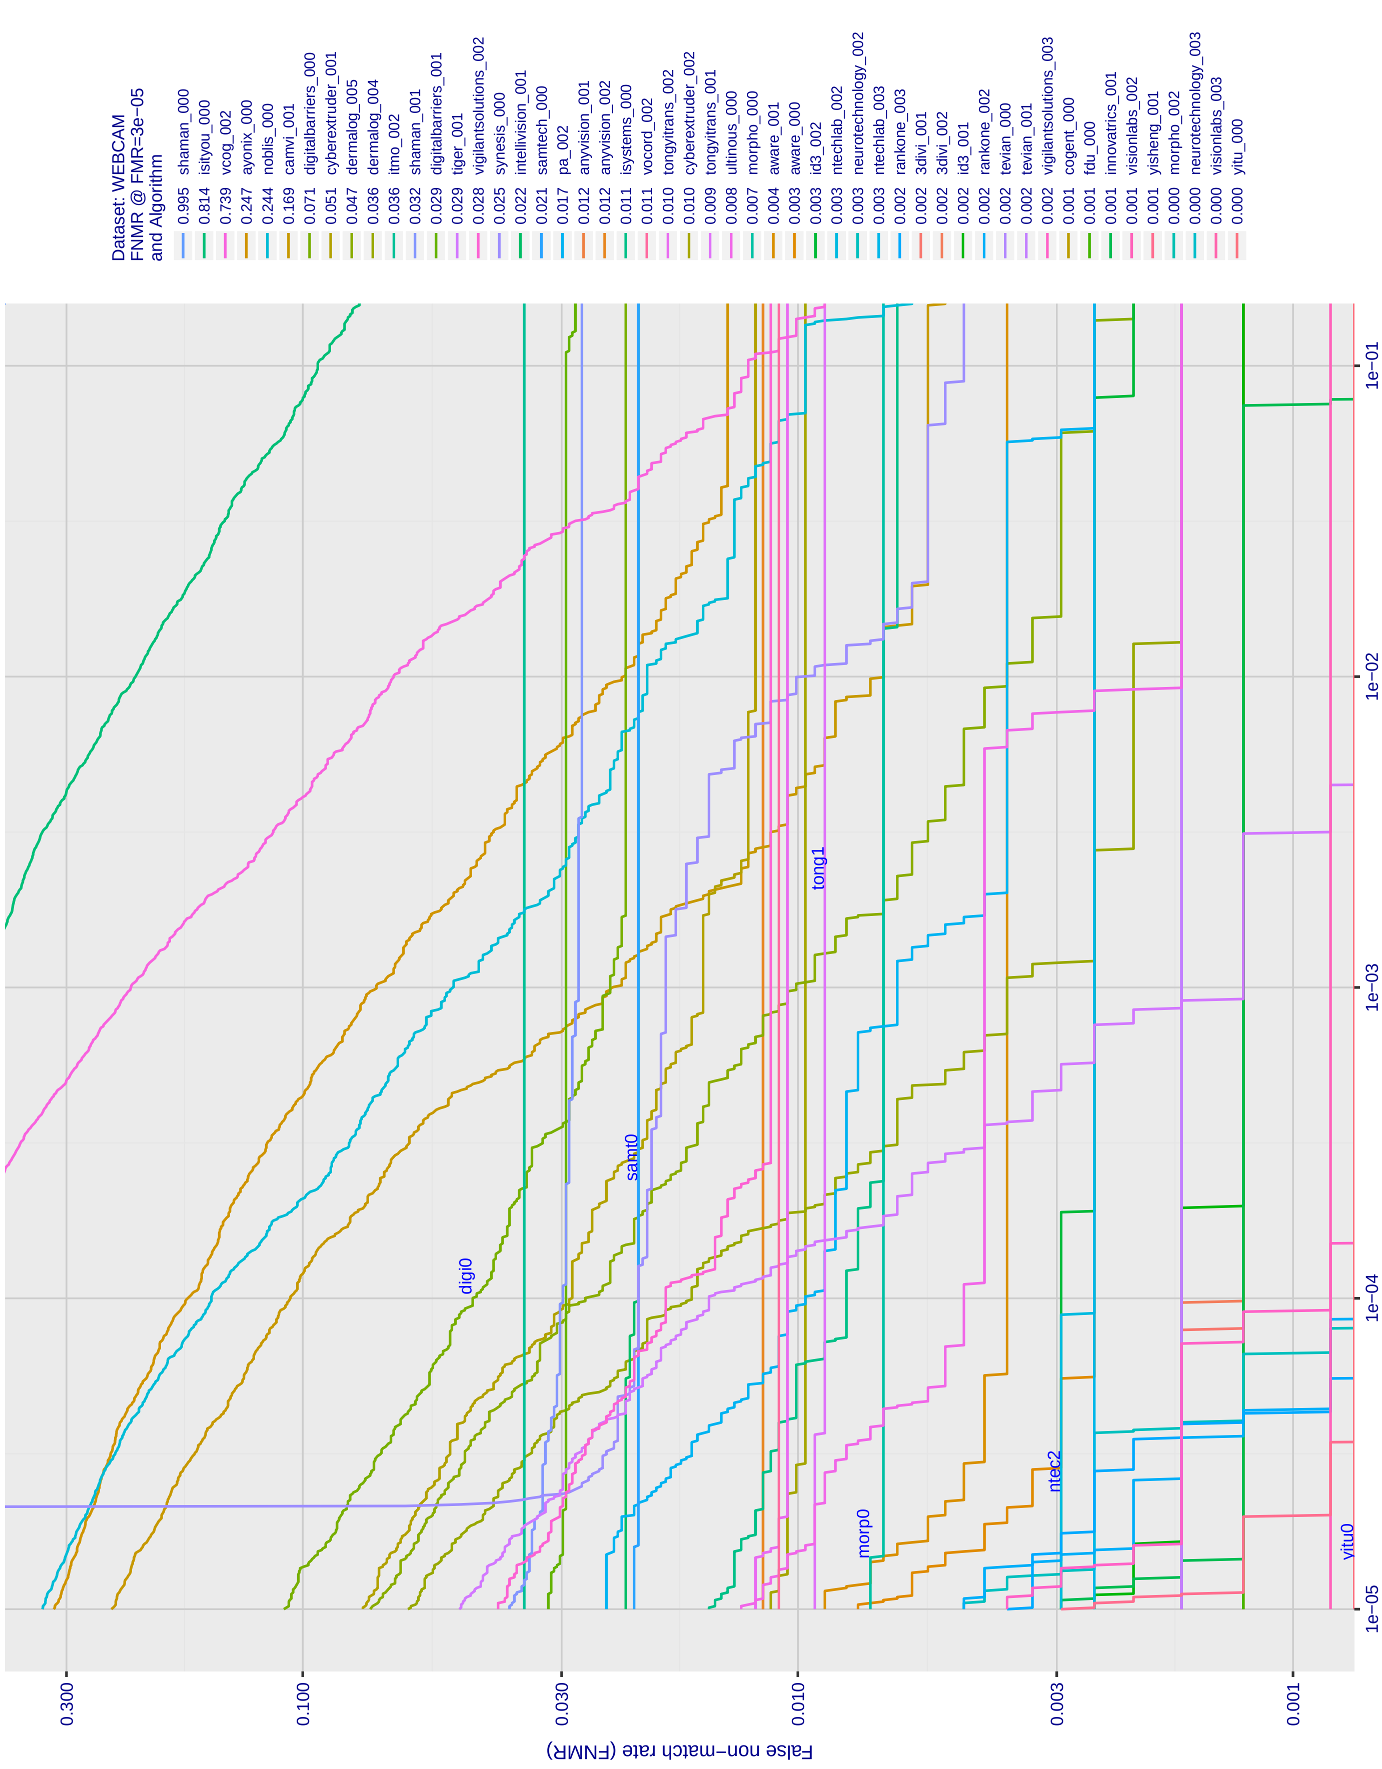

Figure 8: For the webcam-to-portrait comparisons, detection error tradeoff (DET) *characteristics showing false non-match rate vs. false match rate plotted parametrically on threshold, T. The scales are logarithmic in order to show several decades of* FMR. *Caution: The FNMR values here are optimistic statements of accuracy because the image pairs were collected on the same day. This is known across biometrics to give better accuracy, and is operationally relevant only in special cases.*

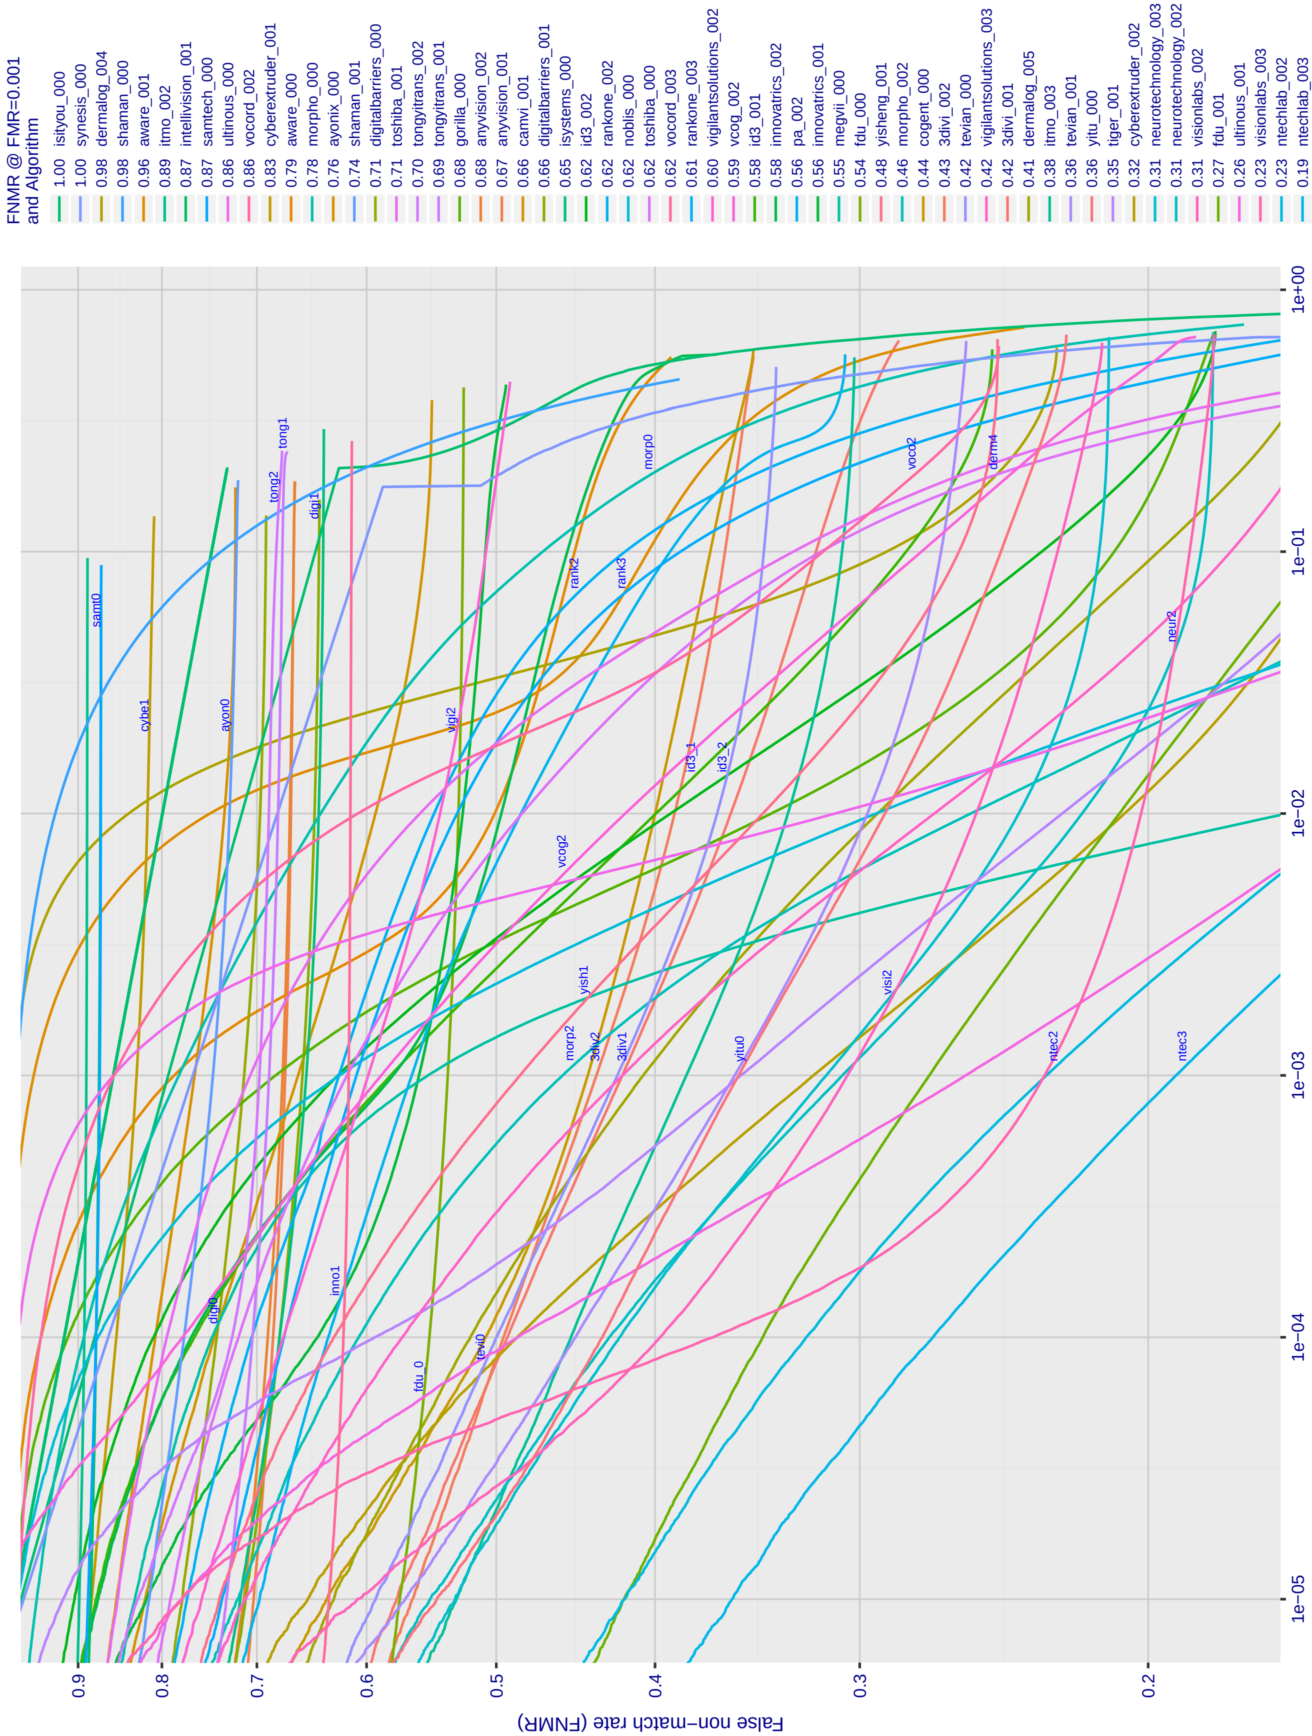

Figure 9: For the wild image comparisons, detection error tradeoff (DET) *characteristics showing false non-match rate vs. false match rate plotted parametrically on threshold, T. The scales are logarithmic in order to show several decades of* FMR.

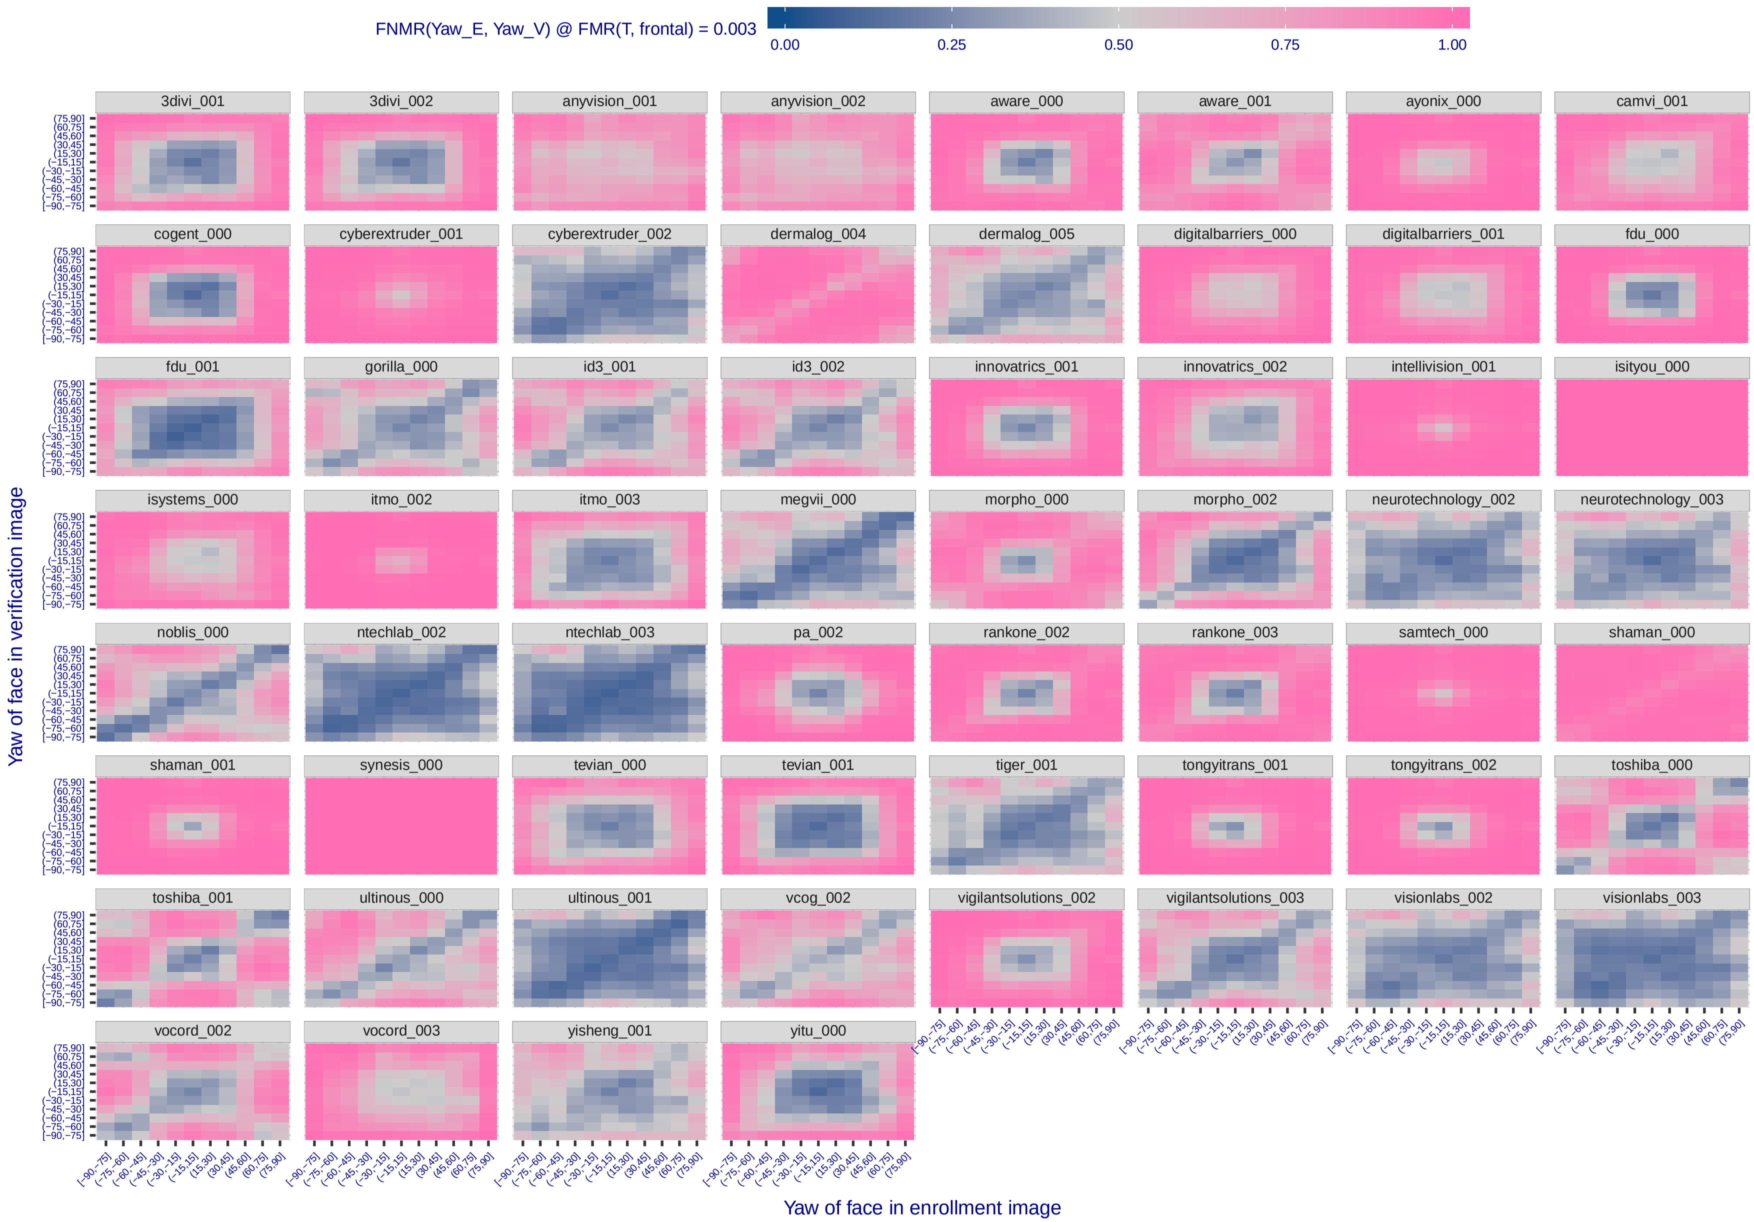

Figure 10: *For wild images, the heatmap show* FNMR *as a function of the yaw of the enrollment and verification images. The threshold is the same in all cells, and is set to the value that yeilds* FMR *= 0.003 on near frontal pairs i.e. where yaw is in the interval* $(-15, 15]$. *Poor algorithms give generally red figures. The better algorithms show a) diagonal dominance, indicating ability to authenticate when pairs have the same yaw angle, and b) off-diagonal cross-pose capability also. The yaw estimates are from an automated pose estimator, and are themselves noisy. The figure assumes that the pose estimates are not systematically incorrect.*

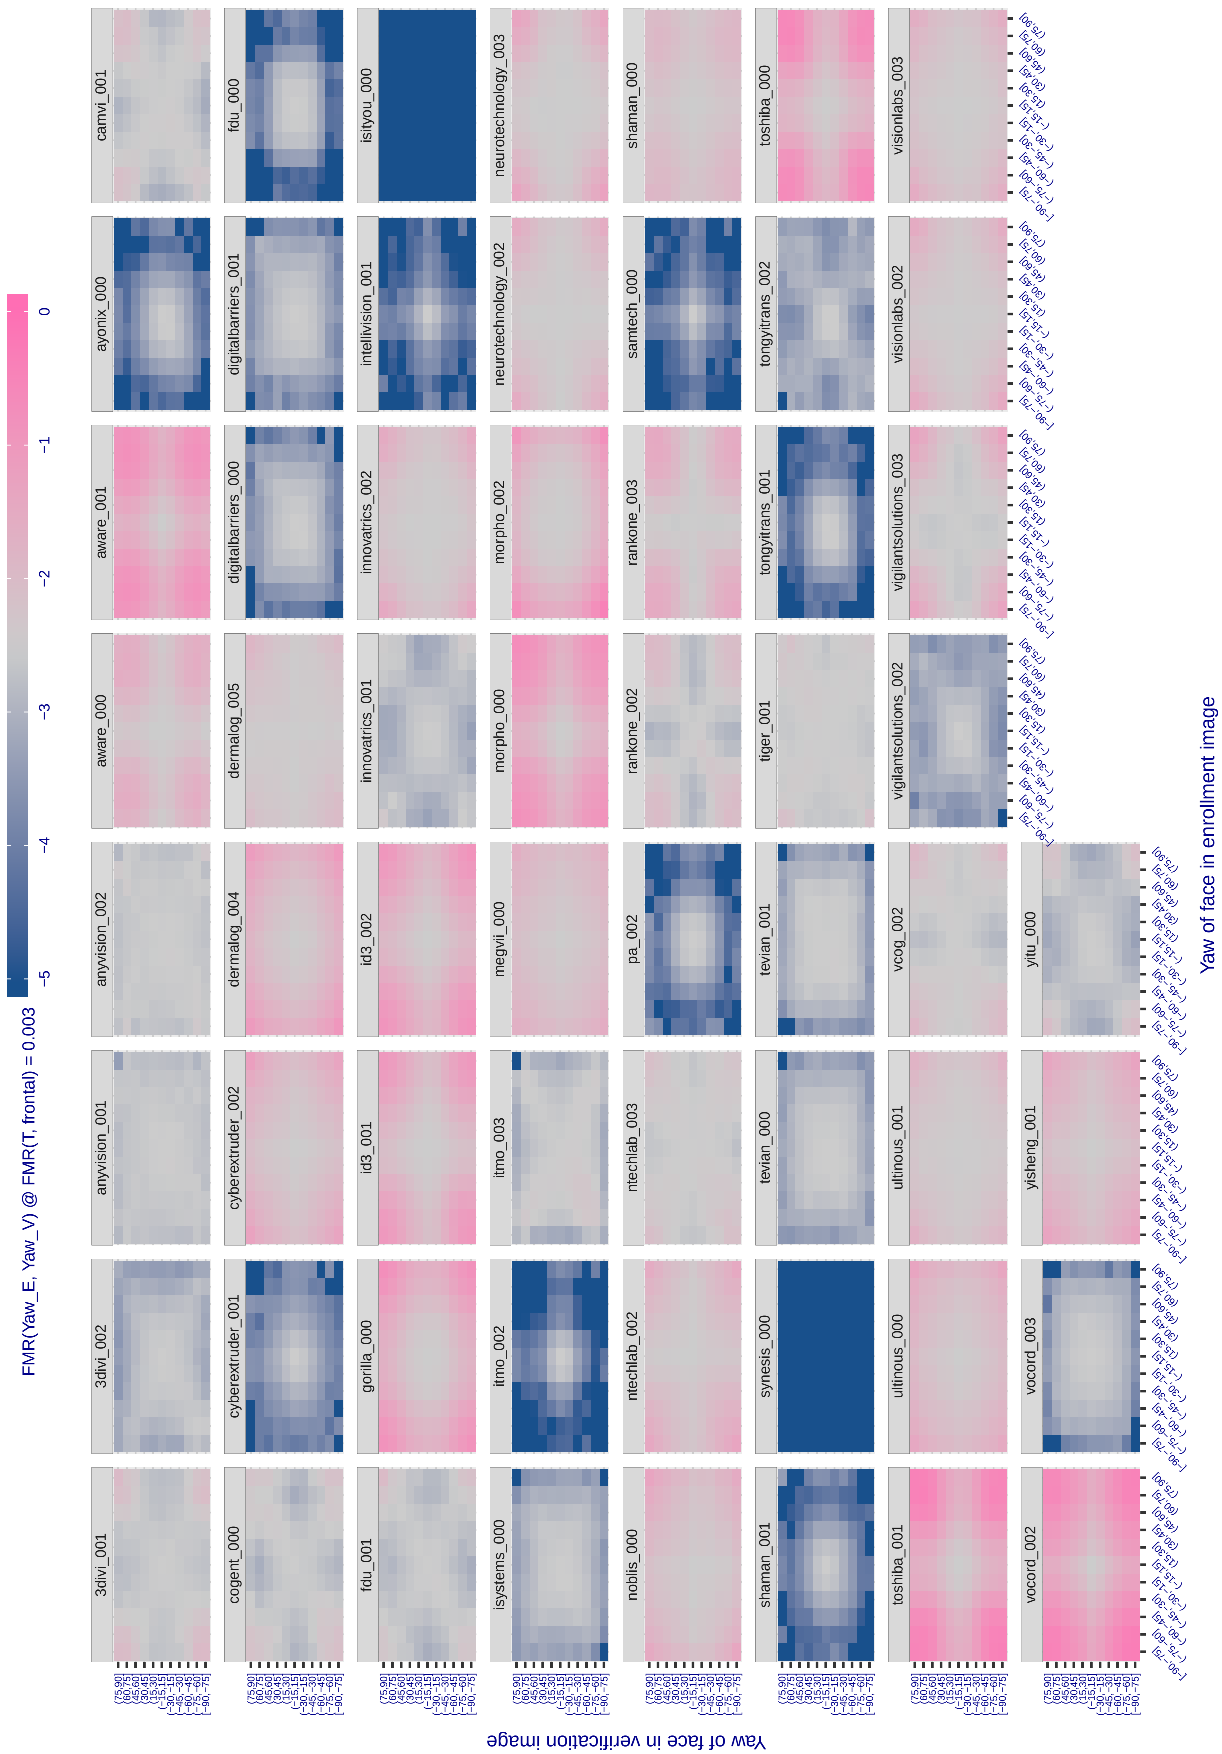

Figure 11: For wild images, the heatmaps show FMR *dependence on the yaw of the enrollment and verification images. The threshold is the same in all cells, and is set to the value that yields* FMR *= 0.003 on near frontal pairs, i.e. where yaw is in the interval* $(-15, 15]$. *Thus the center of each panel is grey. The desired behavior is that* FMR *does not vary with relative yaw. However, some algorithms gives elevated* FMR *when yaw differs. The yaw estimates are from an automated pose estimator, and are themselves noisy. The figure assumes that the pose estimates are not systematically incorrect.*

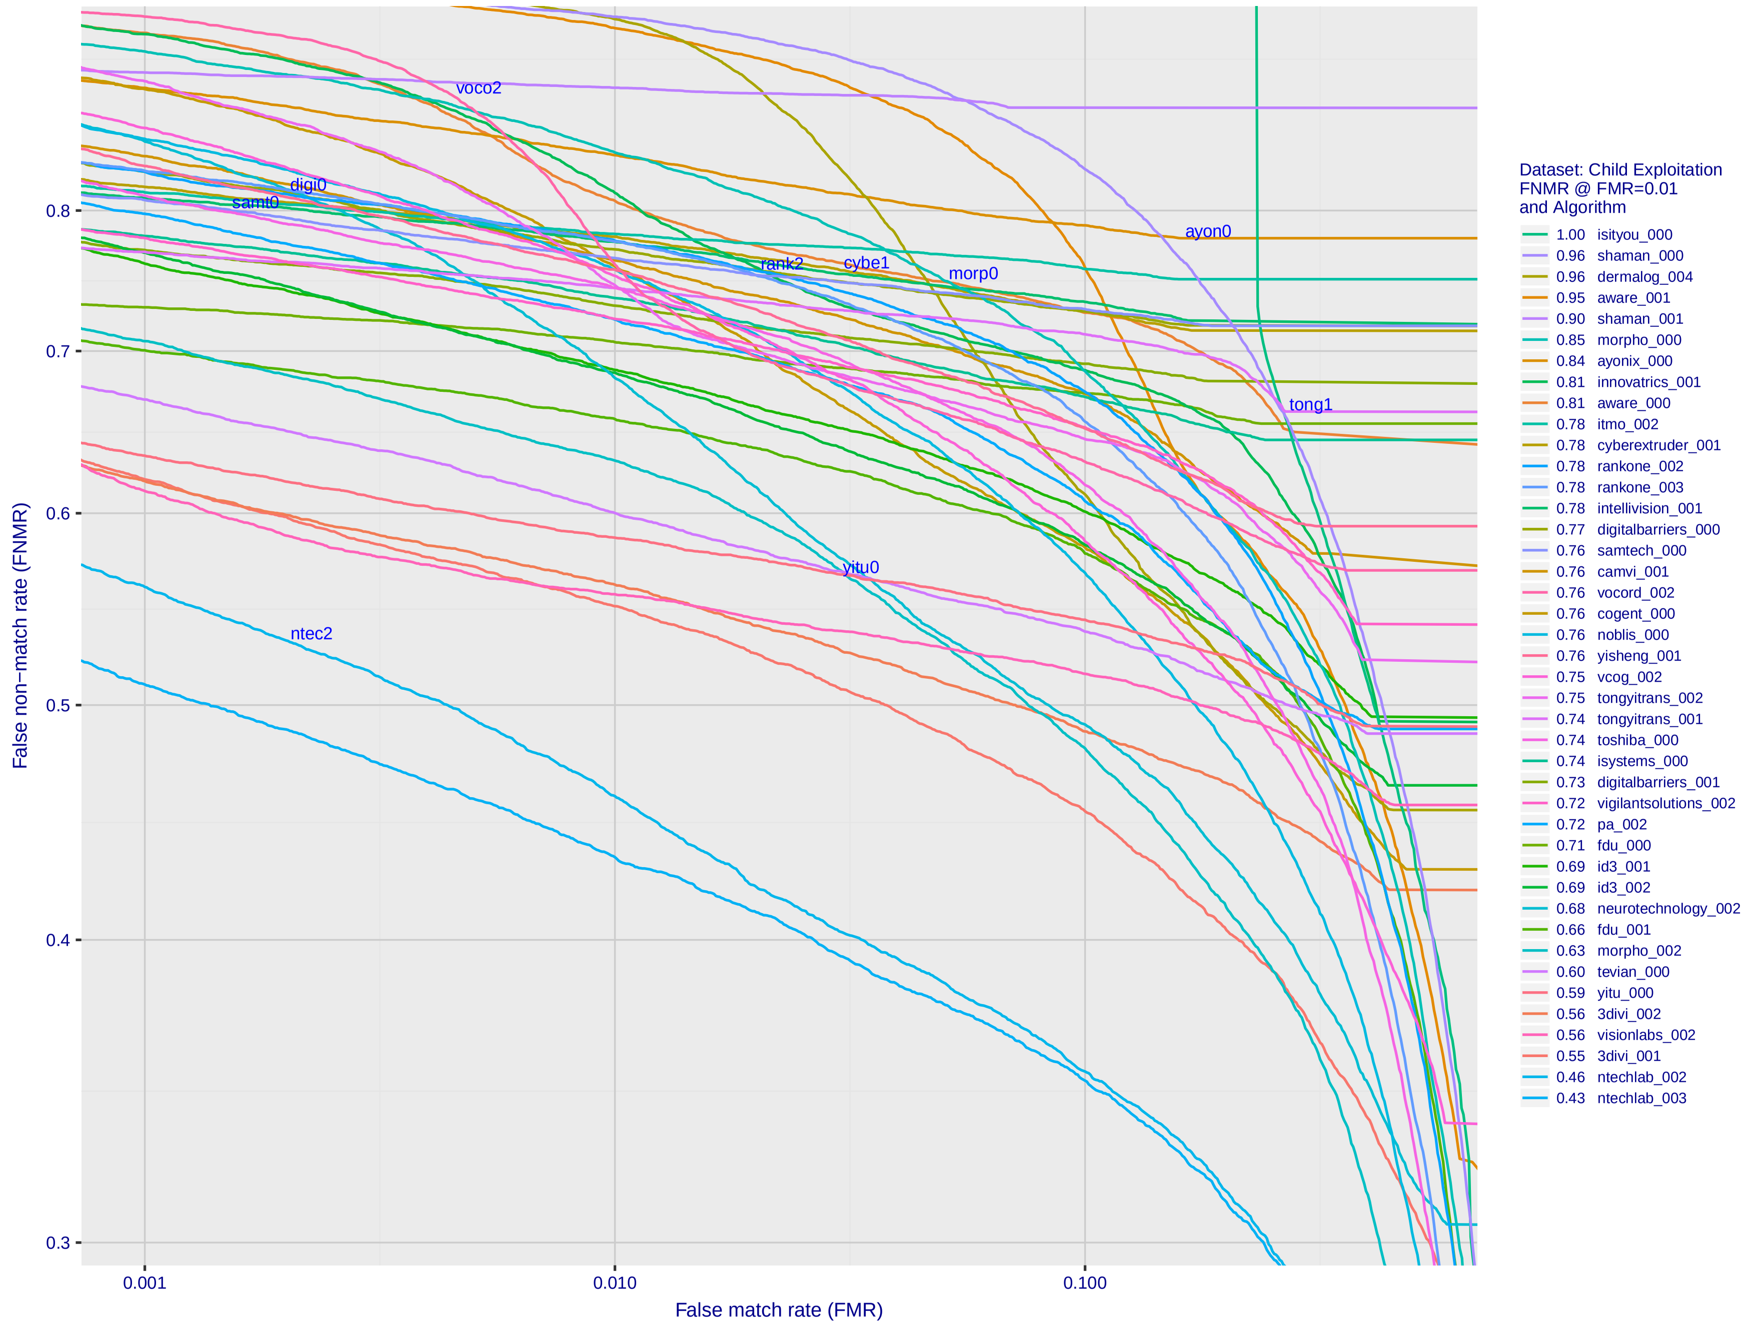

Figure 12: For child exploitation images, detection error tradeoff (DET) characteristics showing false non-match rate vs. false match rate plotted parametrically on threshold, T. The scales are logarithmic in order to show many decades of FMR. Accuracy is poor because many images have adverse quality characteristics, and because detection and enrollment fails.

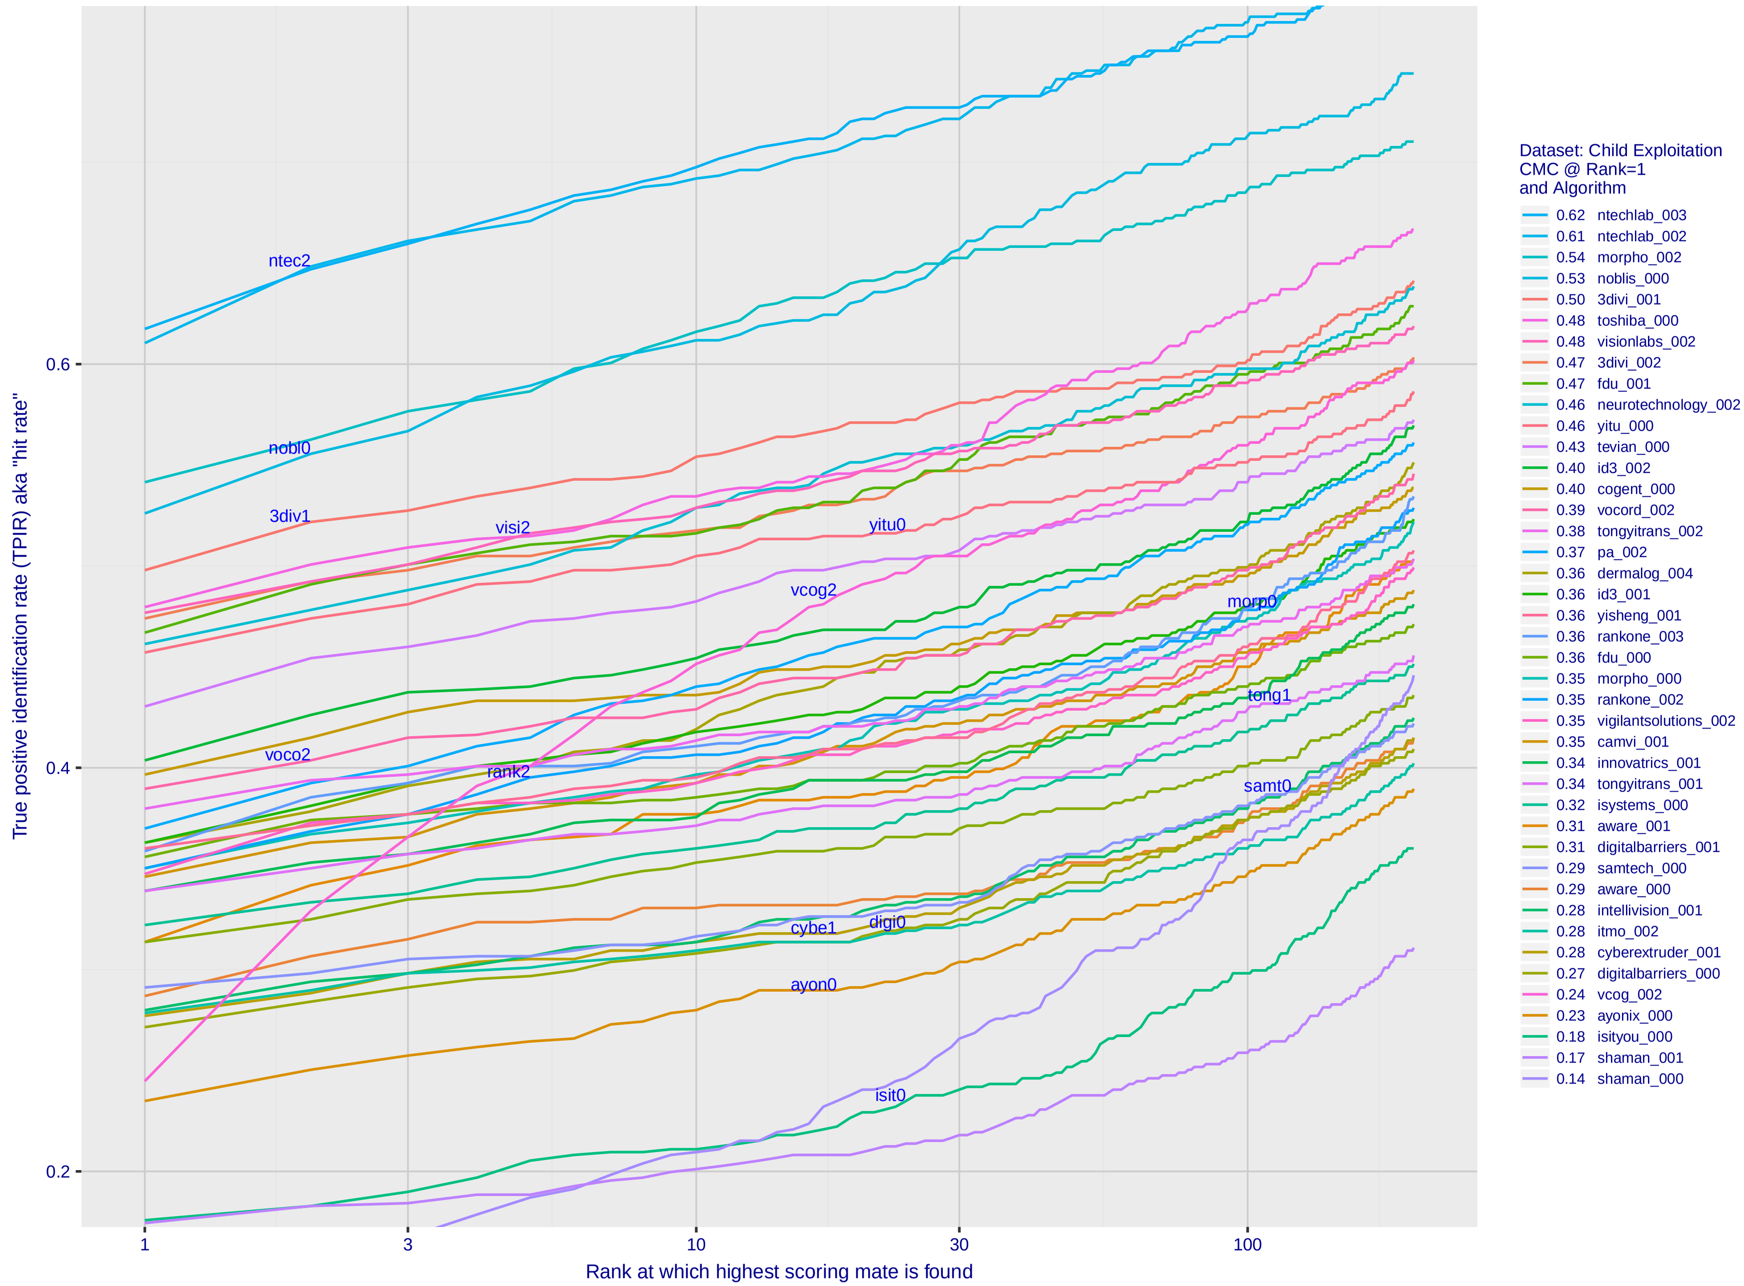

Figure 13: For child exploitation images, cumulative match characteristics (CMC) showing true positive identification rate vs. rank. This is simulation of a one-to-many search experiment - see discussion in section 5.2. The scales are logarithmic in order to show the effect of long candidate lists. Accuracy is poor but much improved relative to the 1:1 DETs of Fig. 12 because a search can succeed if any of a subject's several enrolled images matches the search image with a high score.

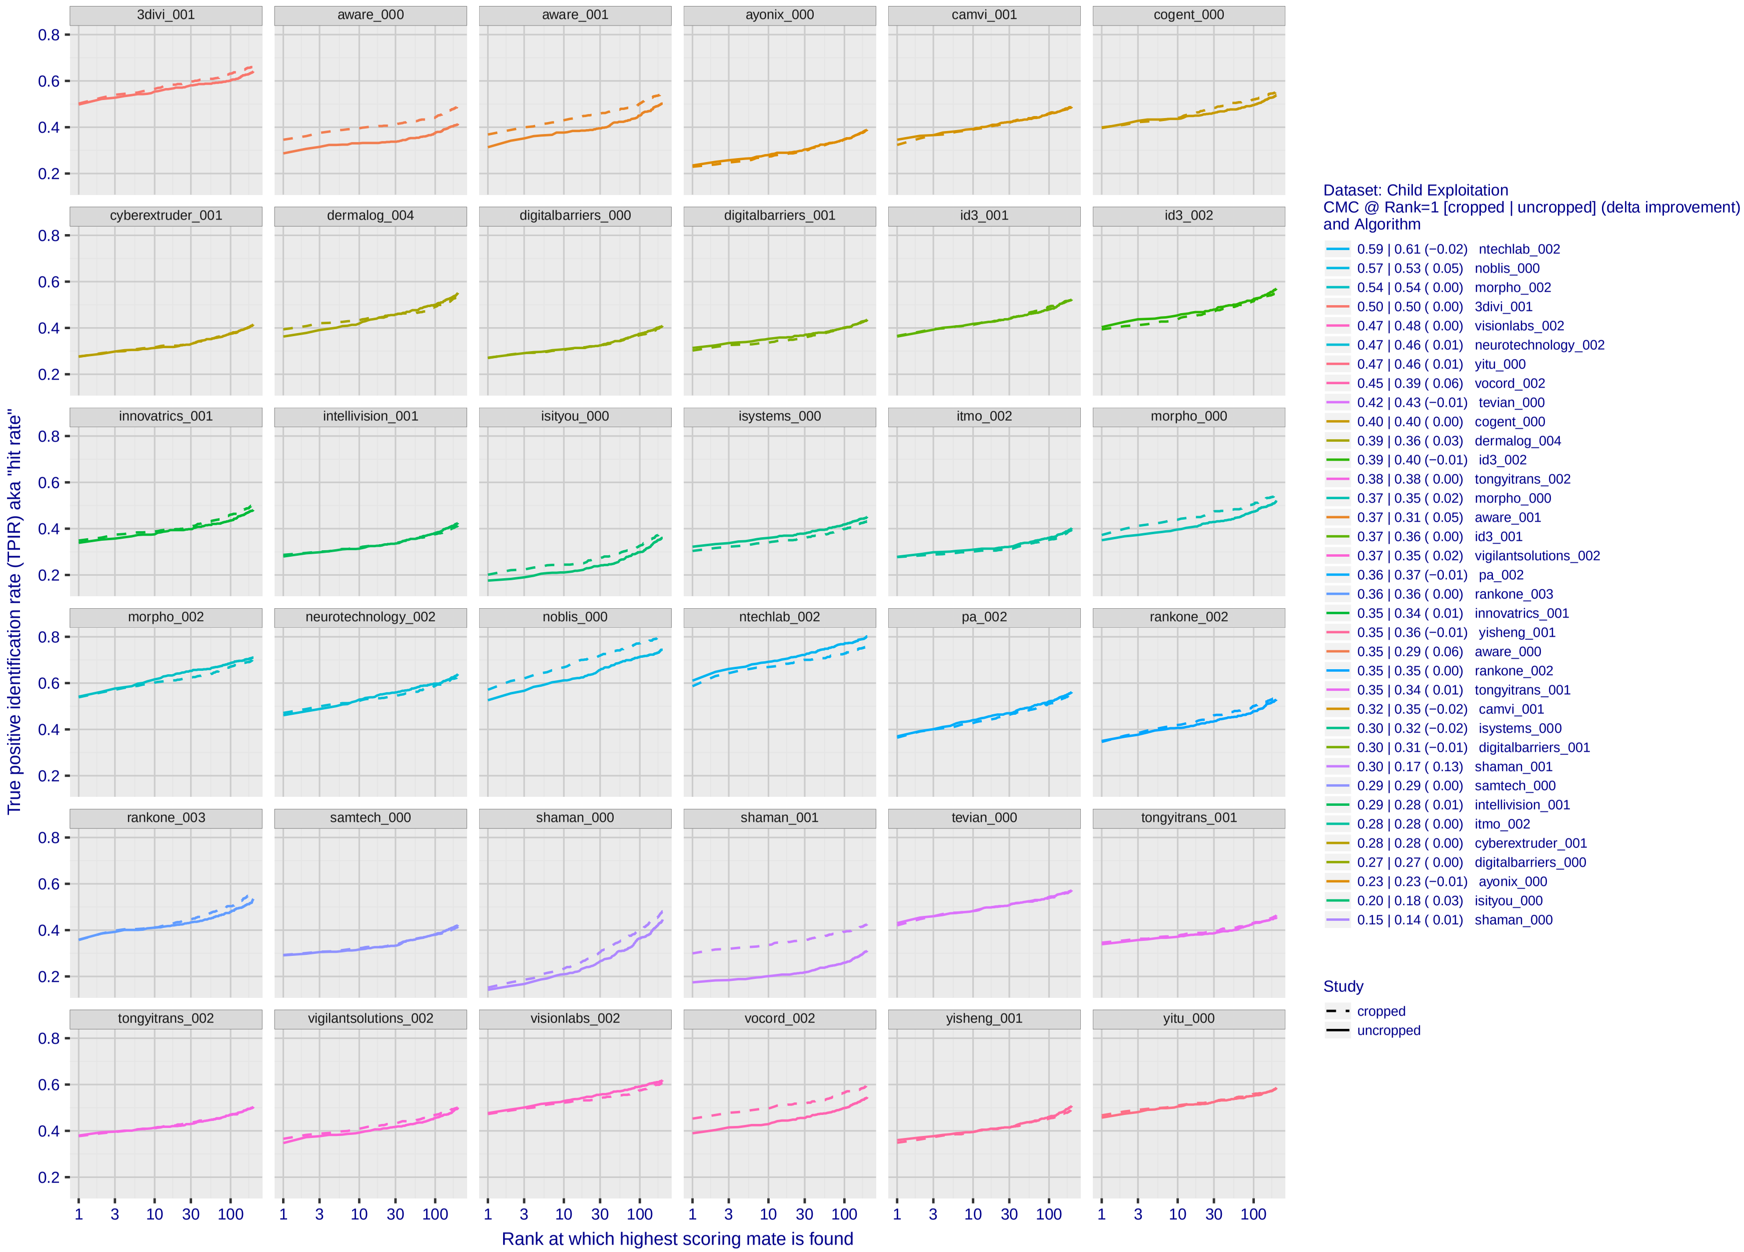

Figure 14: For child exploitation images, cumulative match characteristics (CMC) showing true positive identification rate vs. rank for two cases: 1. Whole image provided to the algorithm; 2. Human annotated rectangular region, cropped and provided to the algorithm. The difference between the traces is associated with detection of difficult faces, and fine localization. The CMC is simulation of a one-to-many search experiment - see discussion in section 5.2. The scales are logarithmic in order to show the effect of long candidate lists. Accuracy is poor but much improved relative to the 1:1 DETs of Fig. 12 because a search can succeed if any *of a subject's several enrolled images matches the search image with a high score.*

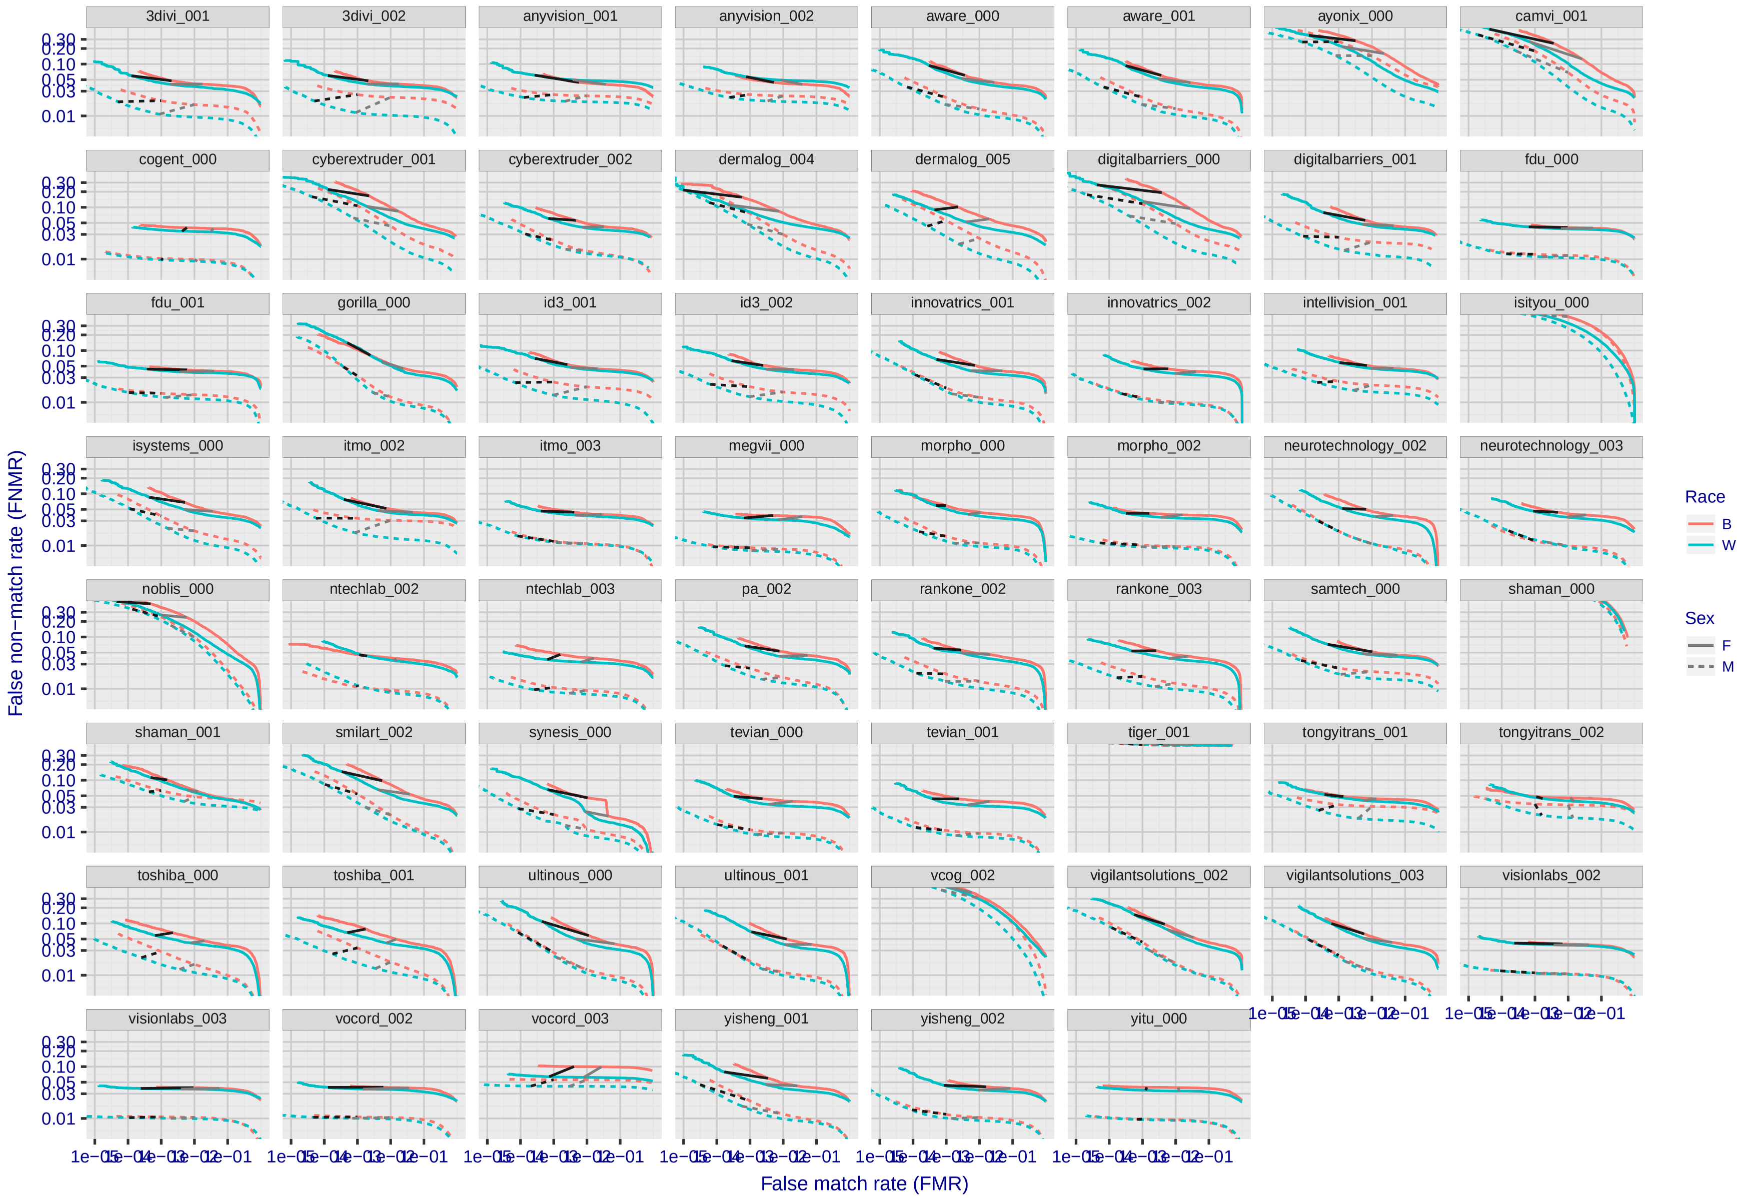

Figure 15: For the mugshot images, error tradeoff characteristics for white females, black females, black males and white males. The grey lines correspond to fixed thresholds, showing how both *FNMR and FMR vary at one operating threshold. Important: Many of the plots will naively be read as saying whites gives lower error rates than blacks because the blue traces lie beneath the red ones. However, this is misleading and incomplete: The grey lines show the traces are generally shifted horizontally. Thus for the dermalog-001 algorithm FNMR for whites is* higher than *for blacks* at a fixed threshold *but, at the same time, FMR is higher for blacks - see Figure 20. As access control systems almost always operate at a fixed threshold, the naive interpretation is incorrect.*

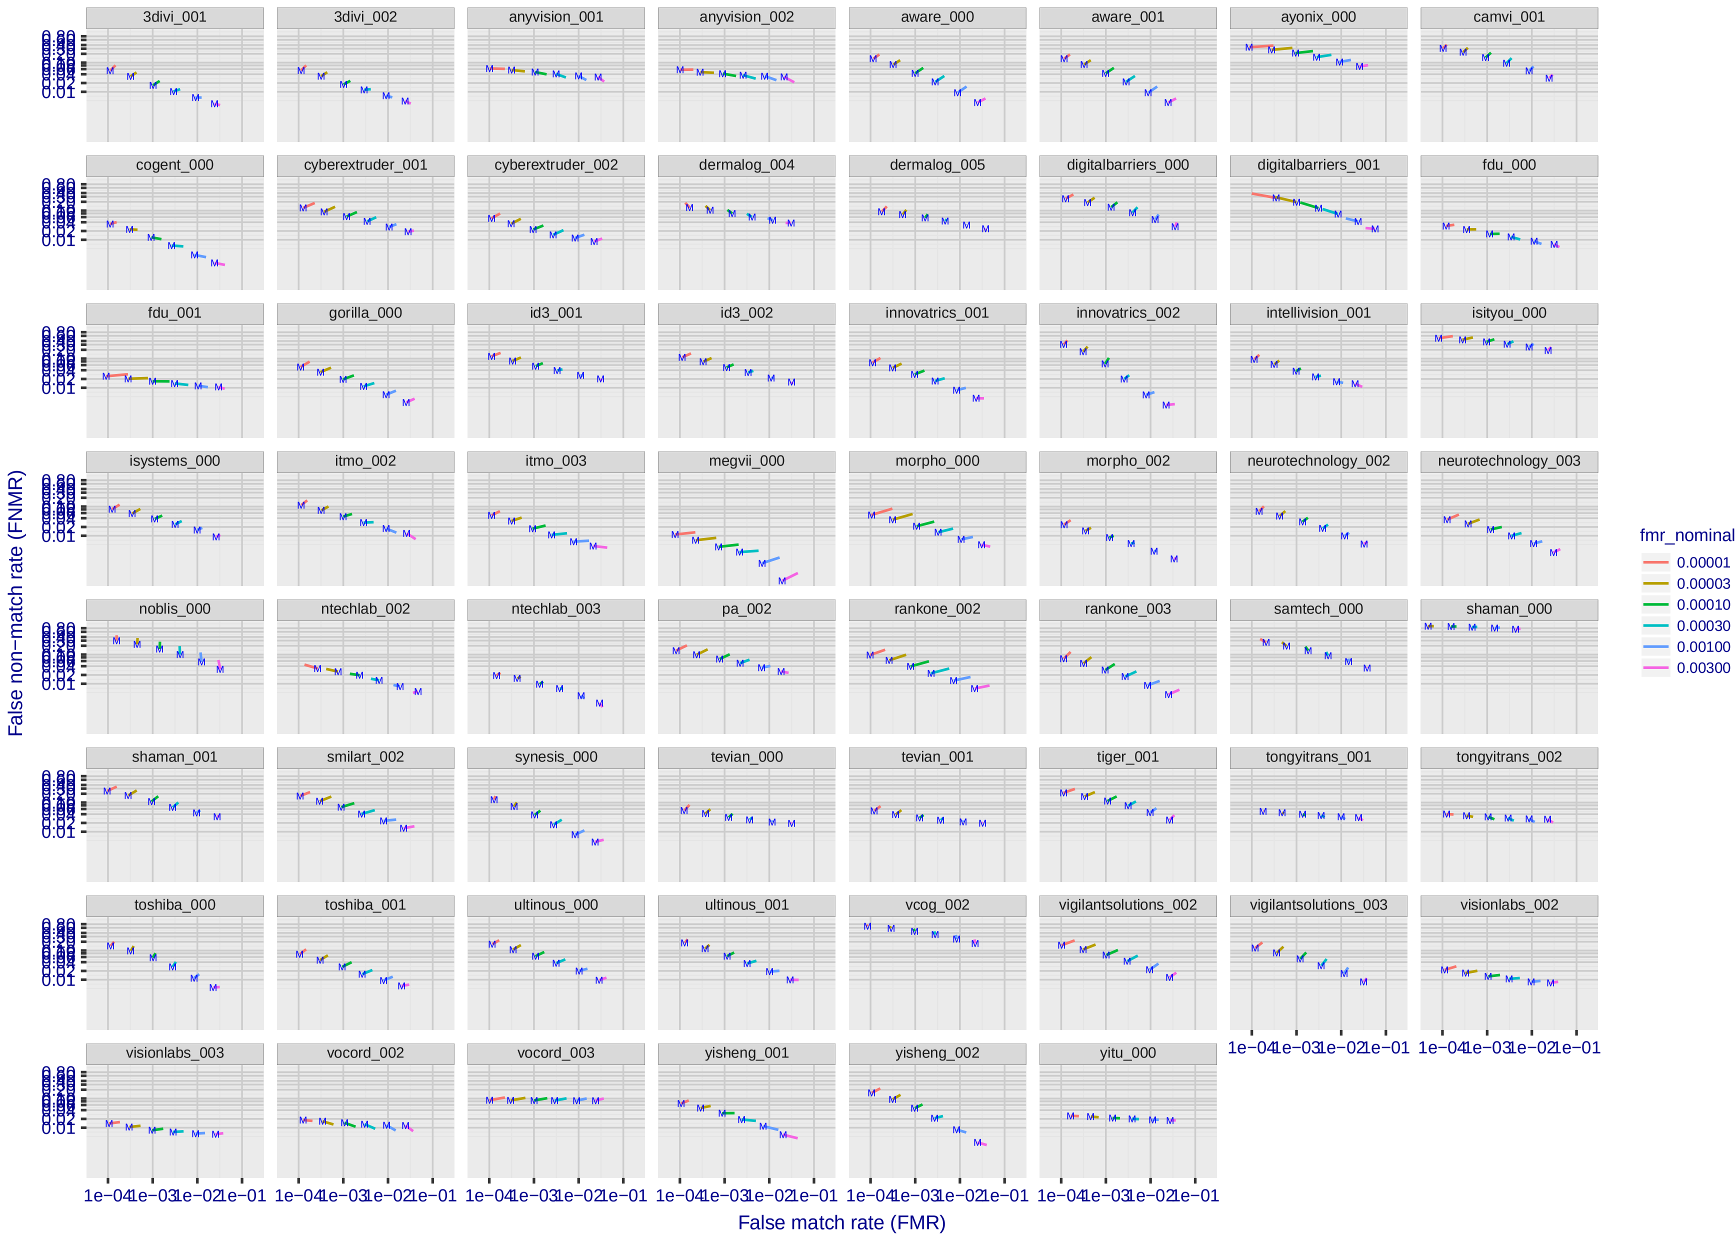

Figure 16: For the visa images, FNMR and FMR at six operating points along the DET characteristic. At each point a line is drawn between (FMR,FNMR)$_{\mathrm{MALE}}$ and (FMR,FNMR)$_{\mathrm{FEMALE}}$ showing how which sex has lower FMR and/or FNMR. The "M" label denotes male, the other end of the line corresponds to female. The six operating thresholds are selected to give the nominal false match rates given in the legend, and are computed over all impostor pairs regardless of age, sex, and place of birth. The plotted FMR values are broadly an order of magnitude larger than the nominal rates because FMR is computed over demographically-matched impostor pairs i.e individuals of the same sex, from the same geographic region (see section 5.6.1), and the same age group (see section 5.6.2).

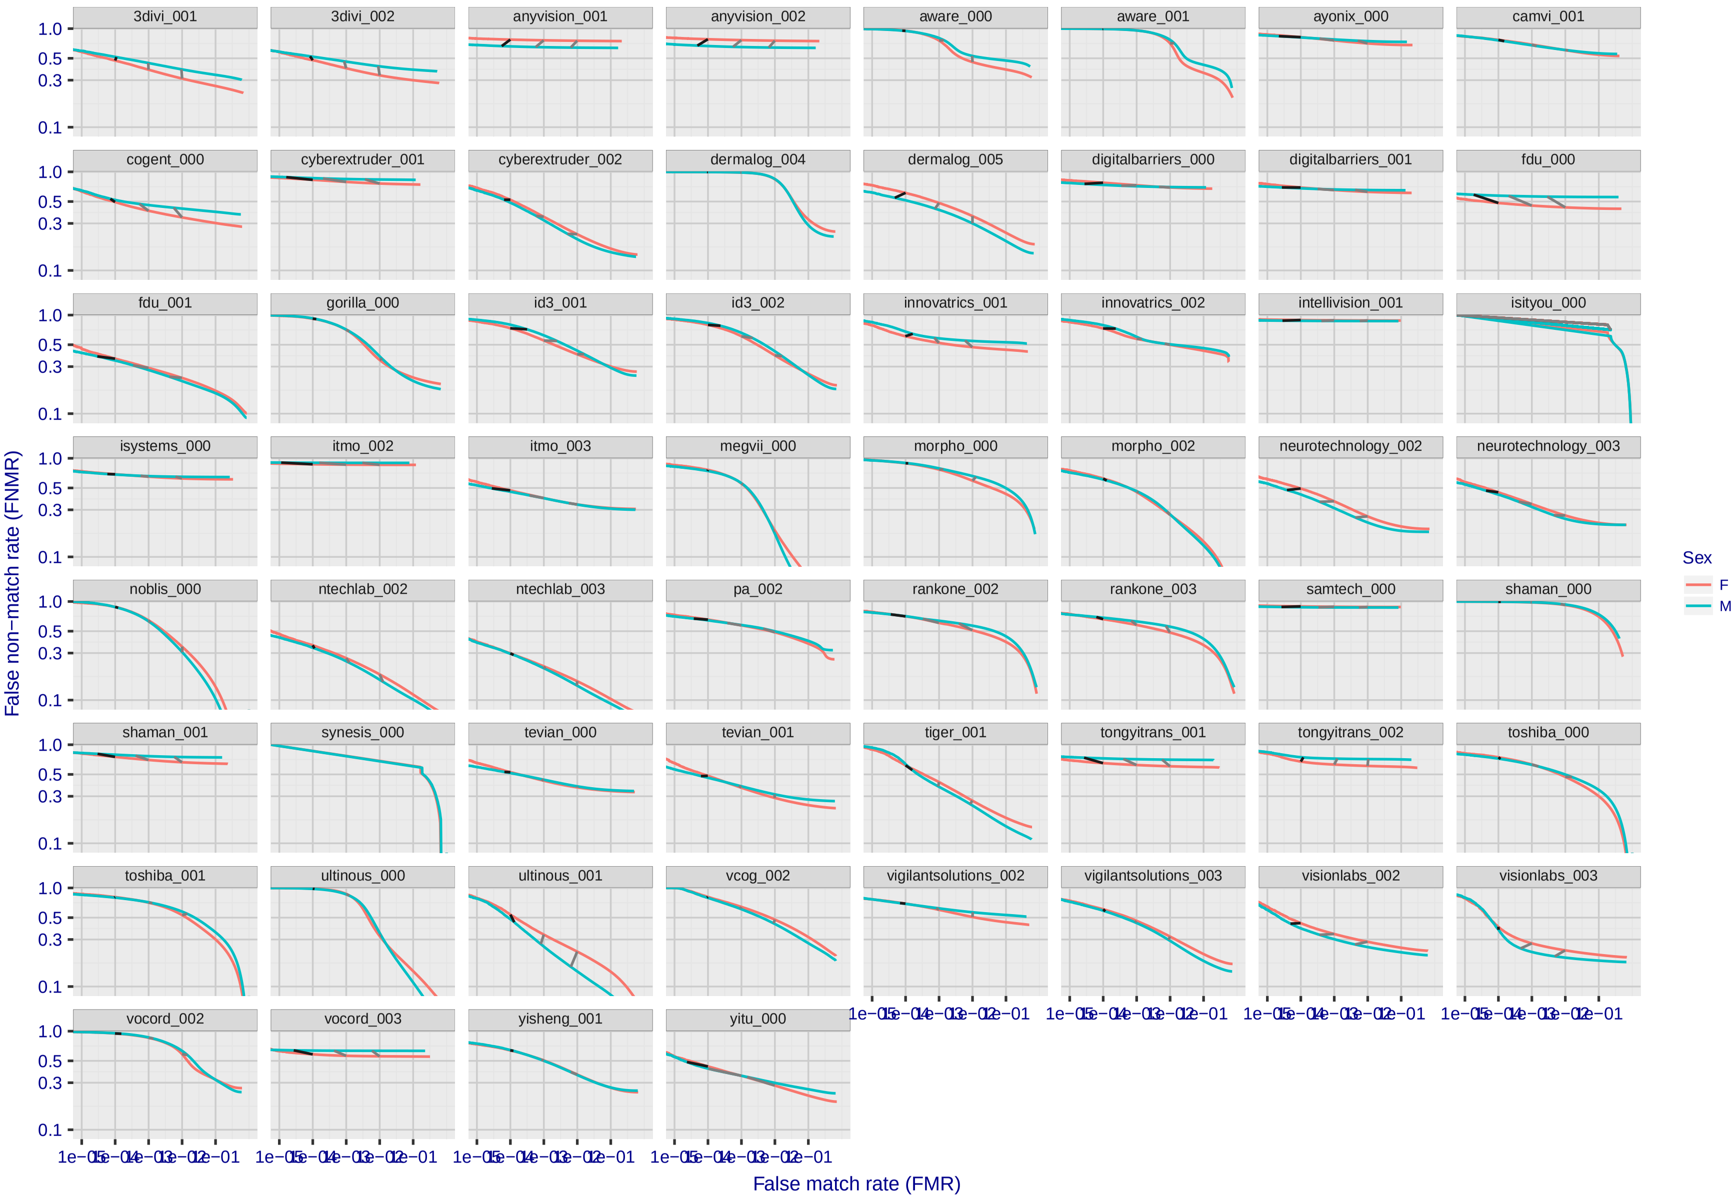

Figure 17: For the wild image comparisons, detection error tradeoff (DET) characteristics showing false non-match rate vs. false match rate plotted parametrically on threshold, T. Error rates are higher here than in the generic wild DET (Fig 9) becuase the impostor pairs here are same-sex only. The scales are logarithmic in order to show several decades of FMR.

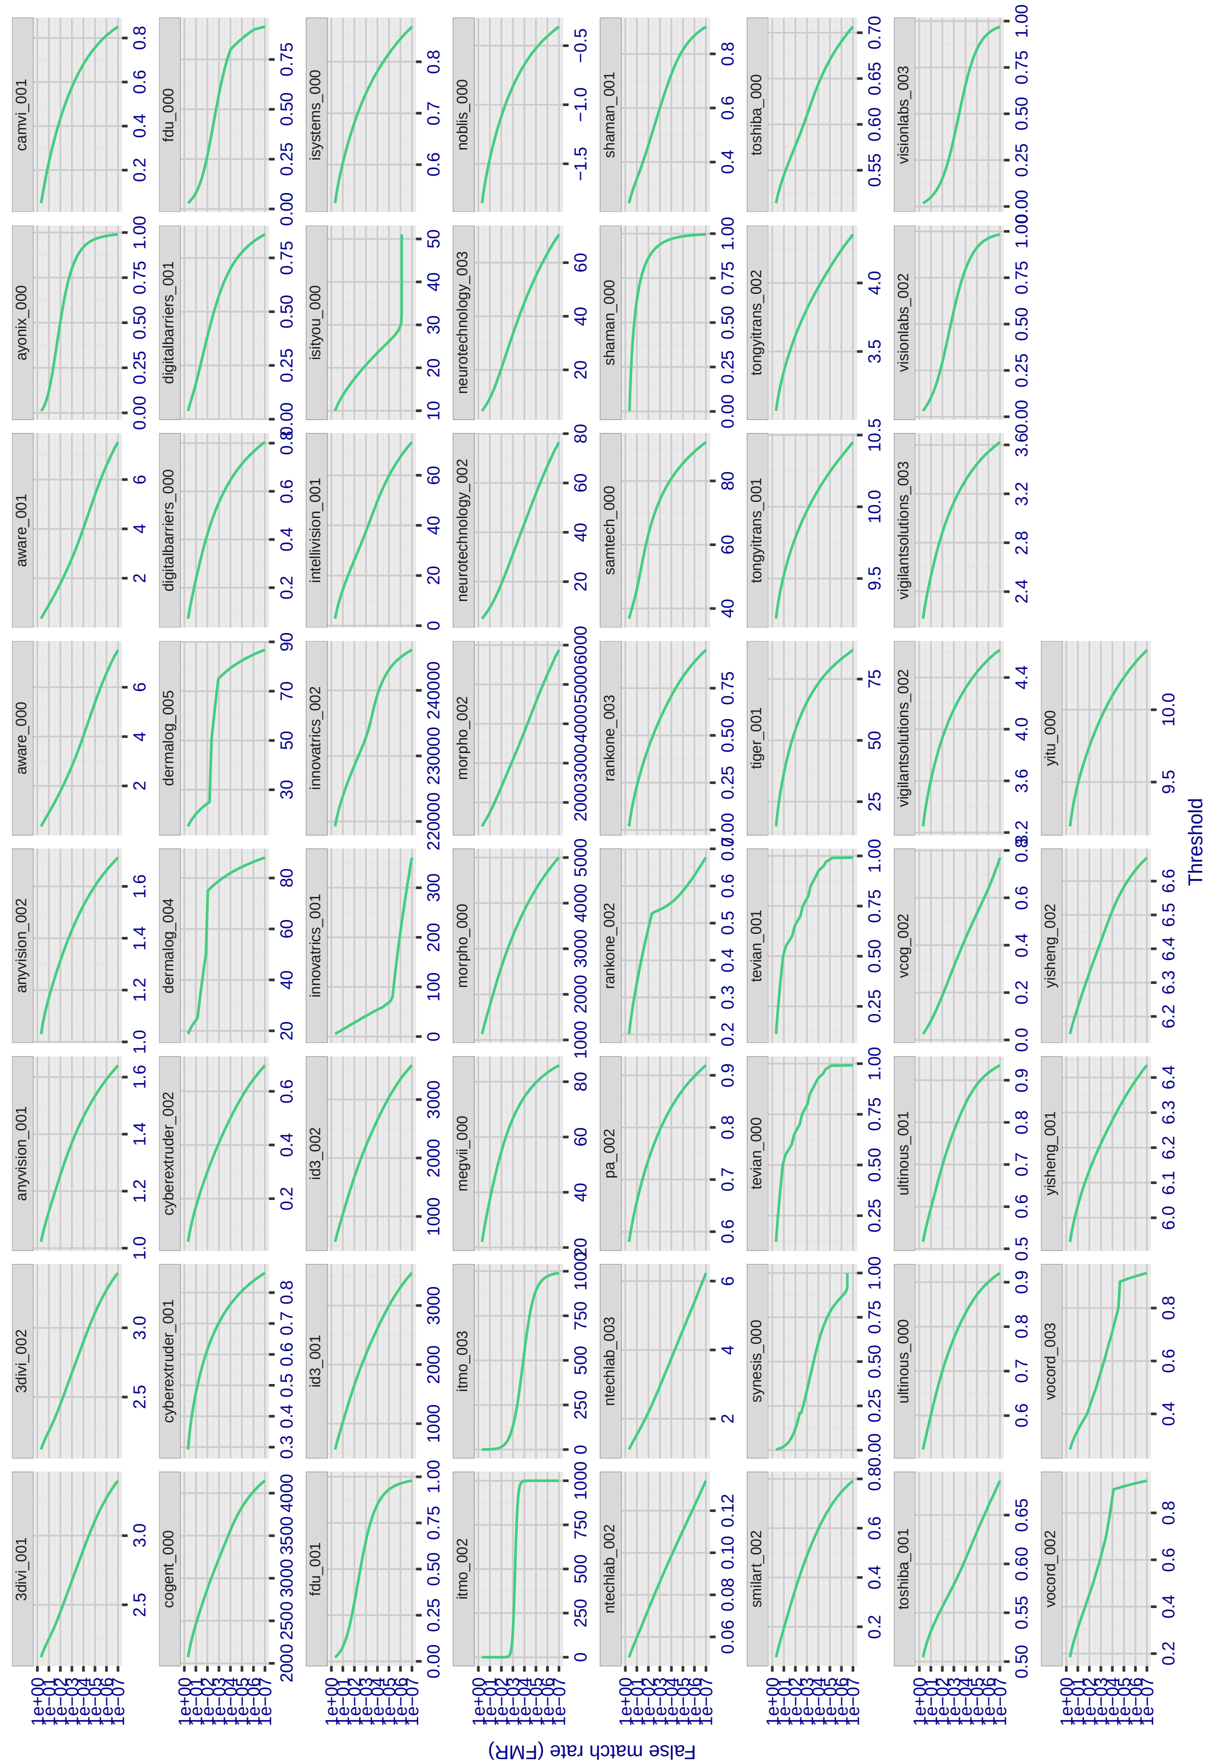

Figure 18: For the visa images, the false match calibration curves show false match rate vs. threshold. These curves apply to zero-effort impostors. As shown later (sec. 5.6), FMR is higher for demographic-matched impostors.

Figure 19: For the visa images, the curves show how false matches are concentrated in certain images. Specifically each line plots FMR(k) with k the number of images rejected in decreasing order of how many false matches that image was involved in. FMR$(0) = 10^{-4}$. In terms of the biometric zoo, the most "wolf-ish" images are rejected first i.e. those enrollment or verification images most often involved in false matches. A flatter response is considered superior. A steeply descending response indicates that certain kinds of images false match against others, e.g. if hypothetically images of men with particular mustaches would falsely match others.

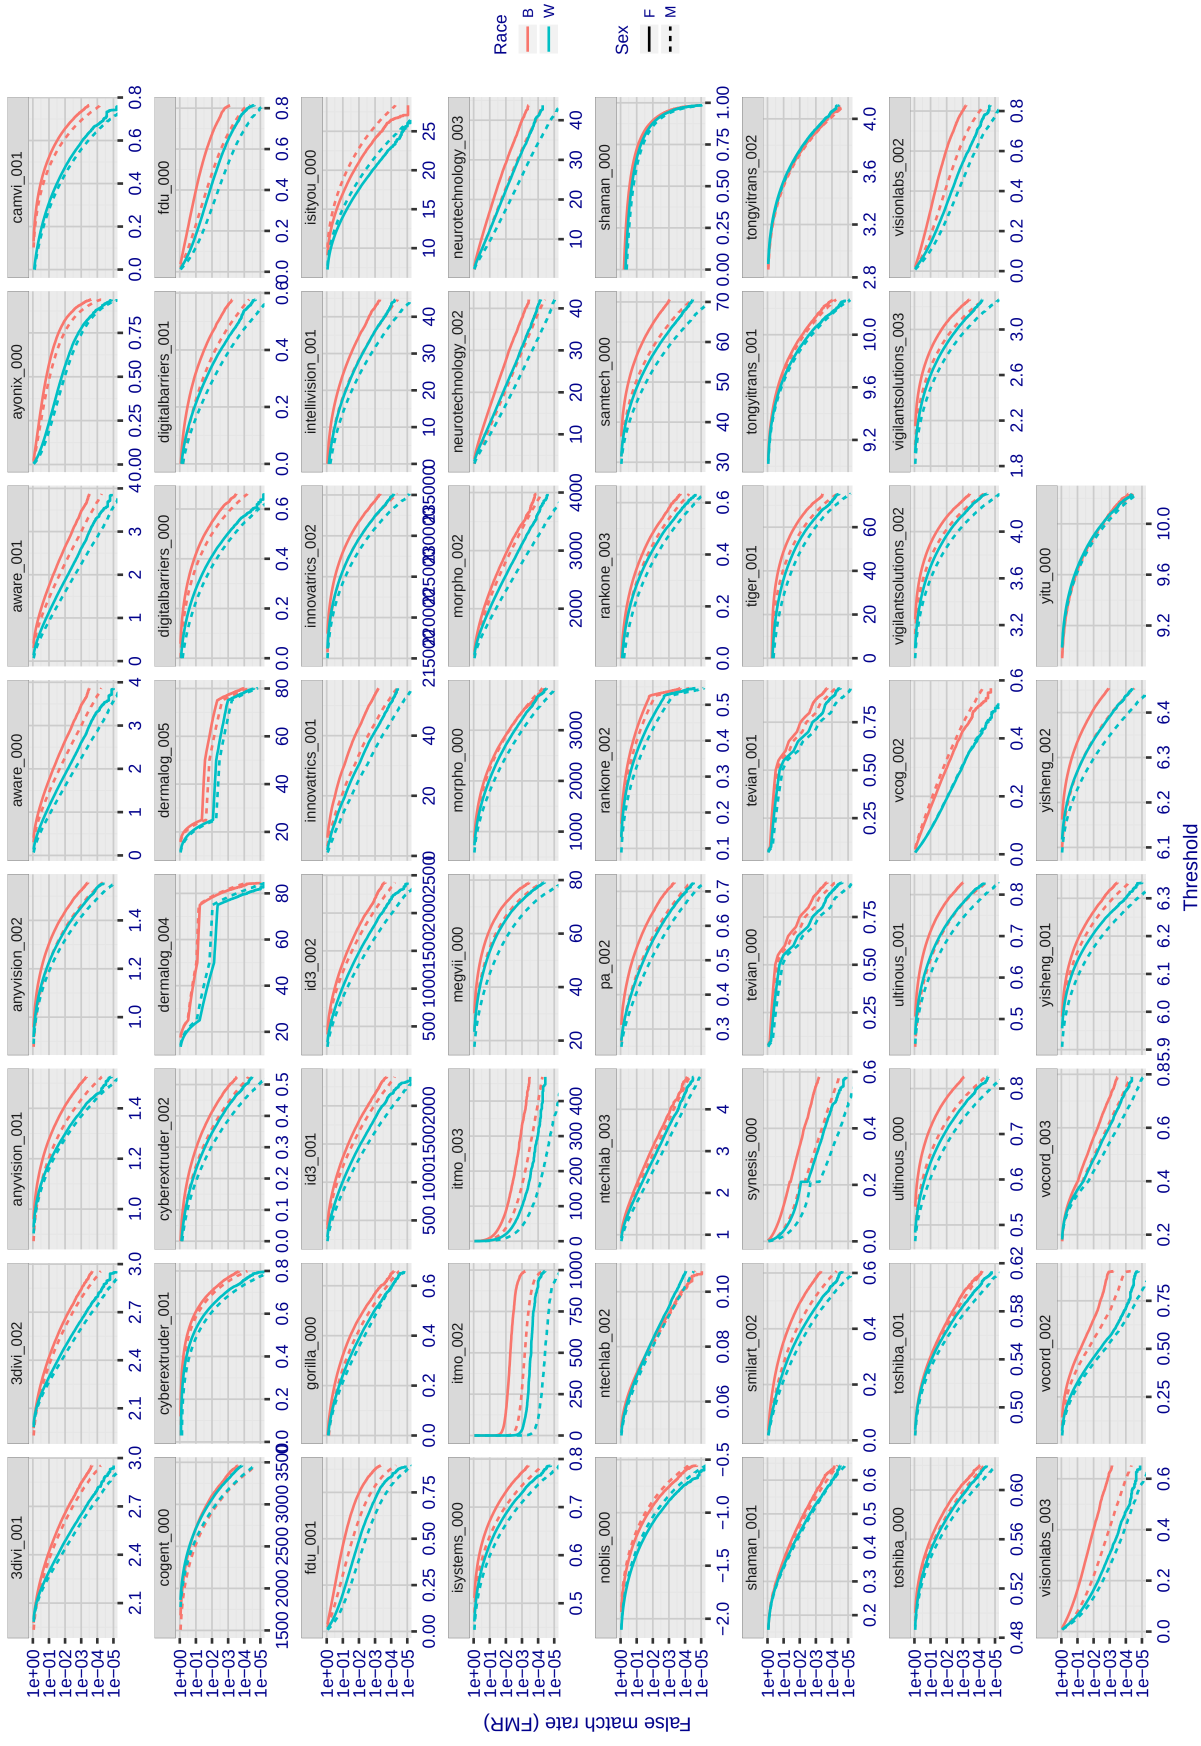

Figure 20: For the mugshot images, the false match calibration curves show false match rate vs. threshold. Separate curves appear for white females, black females, black males and white males.

5.5 Genuine distribution stability

5.5.1 Effect of birth place on the genuine distribution

Background: Both skin tone and bone structure vary geographically. Prior studies have reported variations in FNMR and FMR.

Goal: To measure false non-match rate (FNMR) variation with country of birth.

Methods: Thresholds are determined that give FMR = $\{0.001, 0.0001\}$ over the entire impostor set. Then FNMR is measured over 1000 bootstrap replications of the genuine scores. Only those countries with at least 140 individuals are included in the analysis.

Results: Figure 21 shows FNMR by country of birth for the two thresholds.

Caveats: The results may not relate to subject-specific properties. Instead they could reflect image-specific quality differences, which could occur due to collection protocol or software processing variations.

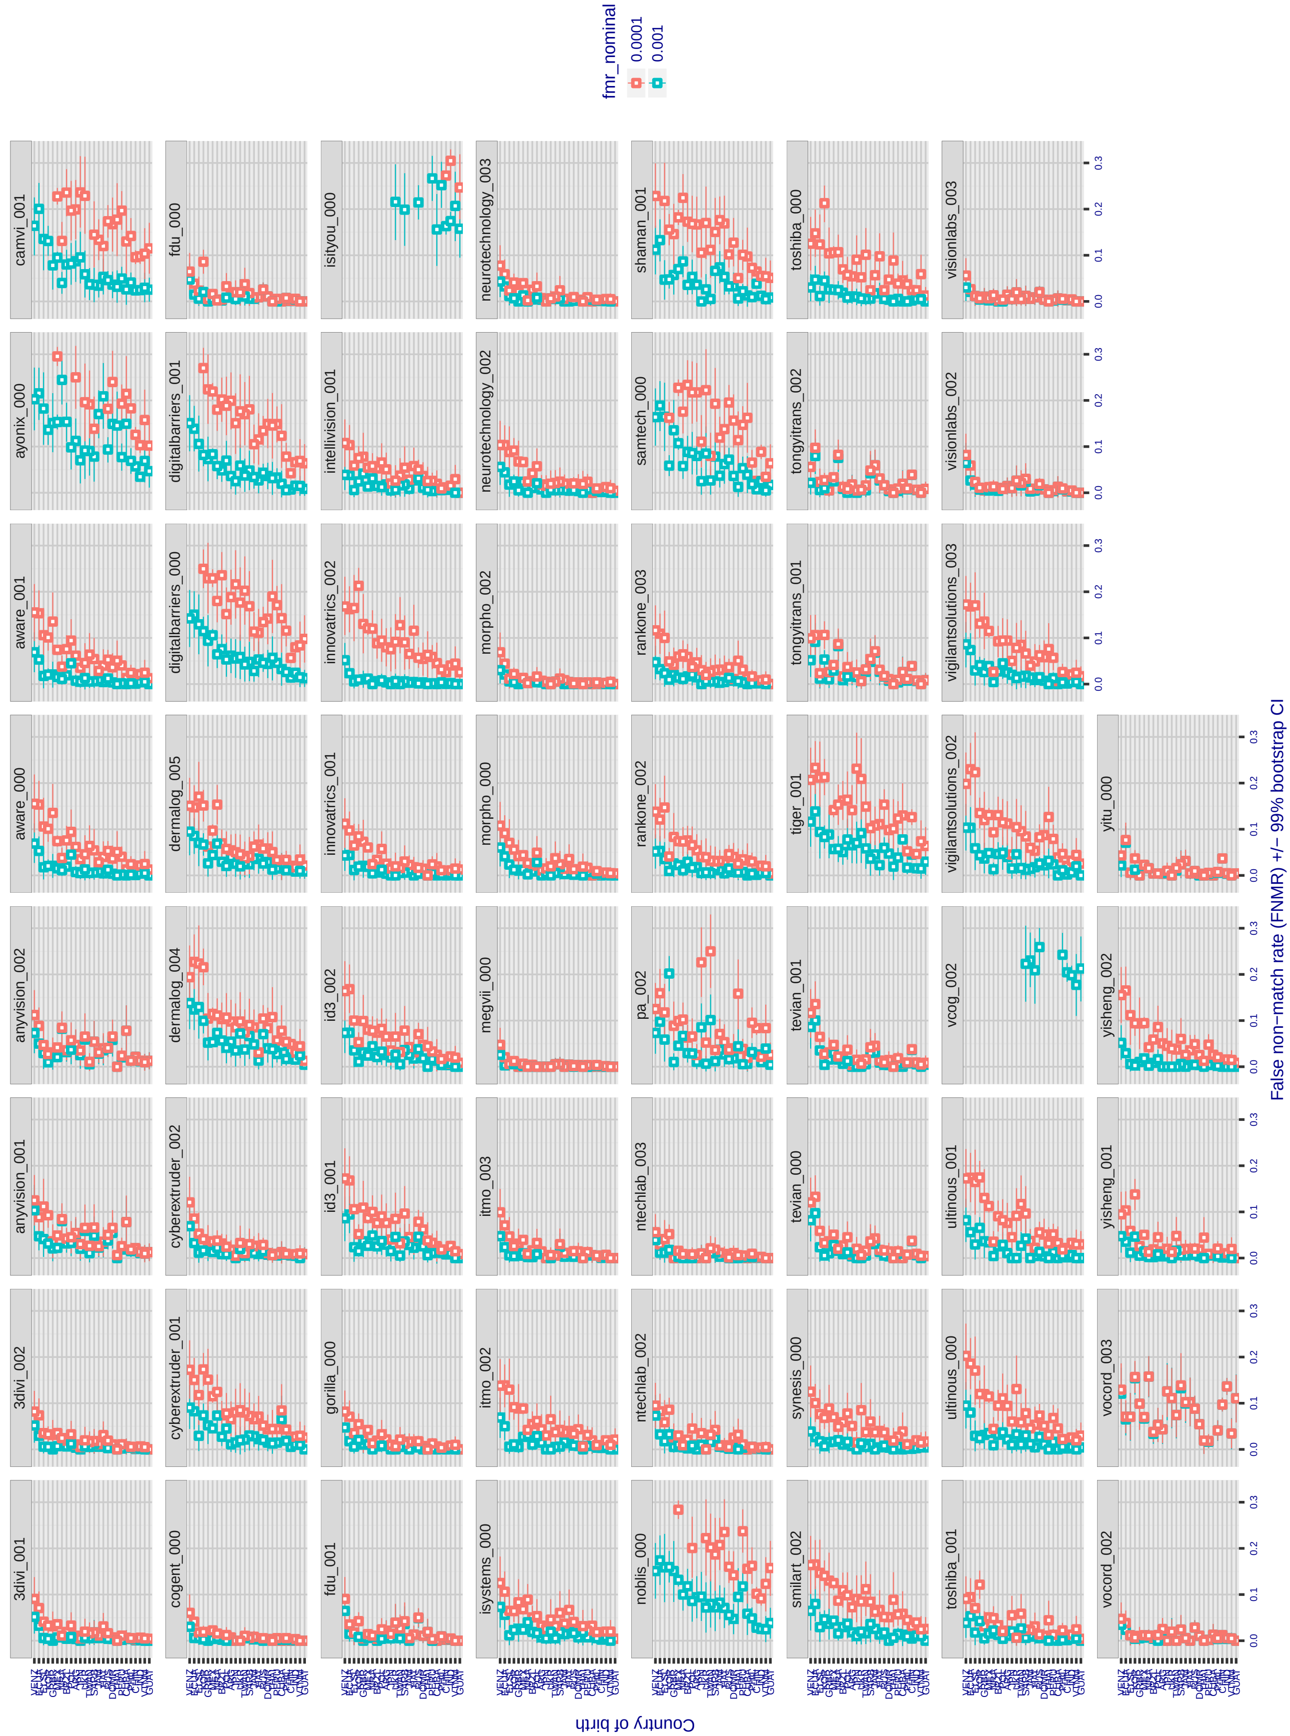

Figure 21: For the visa images, the dots show FNMR by country of birth for two operating thresholds corresponding to FMR = {0.001, 0.0001} computed over all $O(10^{10})$ impostor scores. The figures show an order of magnitude variation in FNMR across country of birth; these effects are due to quality variations. The least accurate countries vary by algorithm.

5.5.2 Effect of age on genuine subjects

Background: Faces change appearance throughout life. Face recognition algorithms have previously beeen reported to give better accuracy on older individuals (See NIST IR 8009).

Goal: To quantify false non-match rates (FNMR) as a function of age. We do not aim to quantify ageing effects here as the separation between two samples is limited to just a few years.

Methods: Using the visa images, thresholds are determined that give FMR = 0.001 and 0.0001 over the entire impostor set. Then FNMR is measured over 1000 bootstrap replications of the genuine scores.

Results: For the visa images, Figure 22 shows how false non-match rates for genuine users, as a function of age group.

The notable aspects are:

▷ Younger subjects give considerably higher FNMR. This is likely due to rapid growth and change in facial appearance.

▷ FNMR trends down throughout life. The last bin, AGE $>$ 72, contains fewer than 140 mated pairs, and may be affected by small sample size.

Caveats: None.

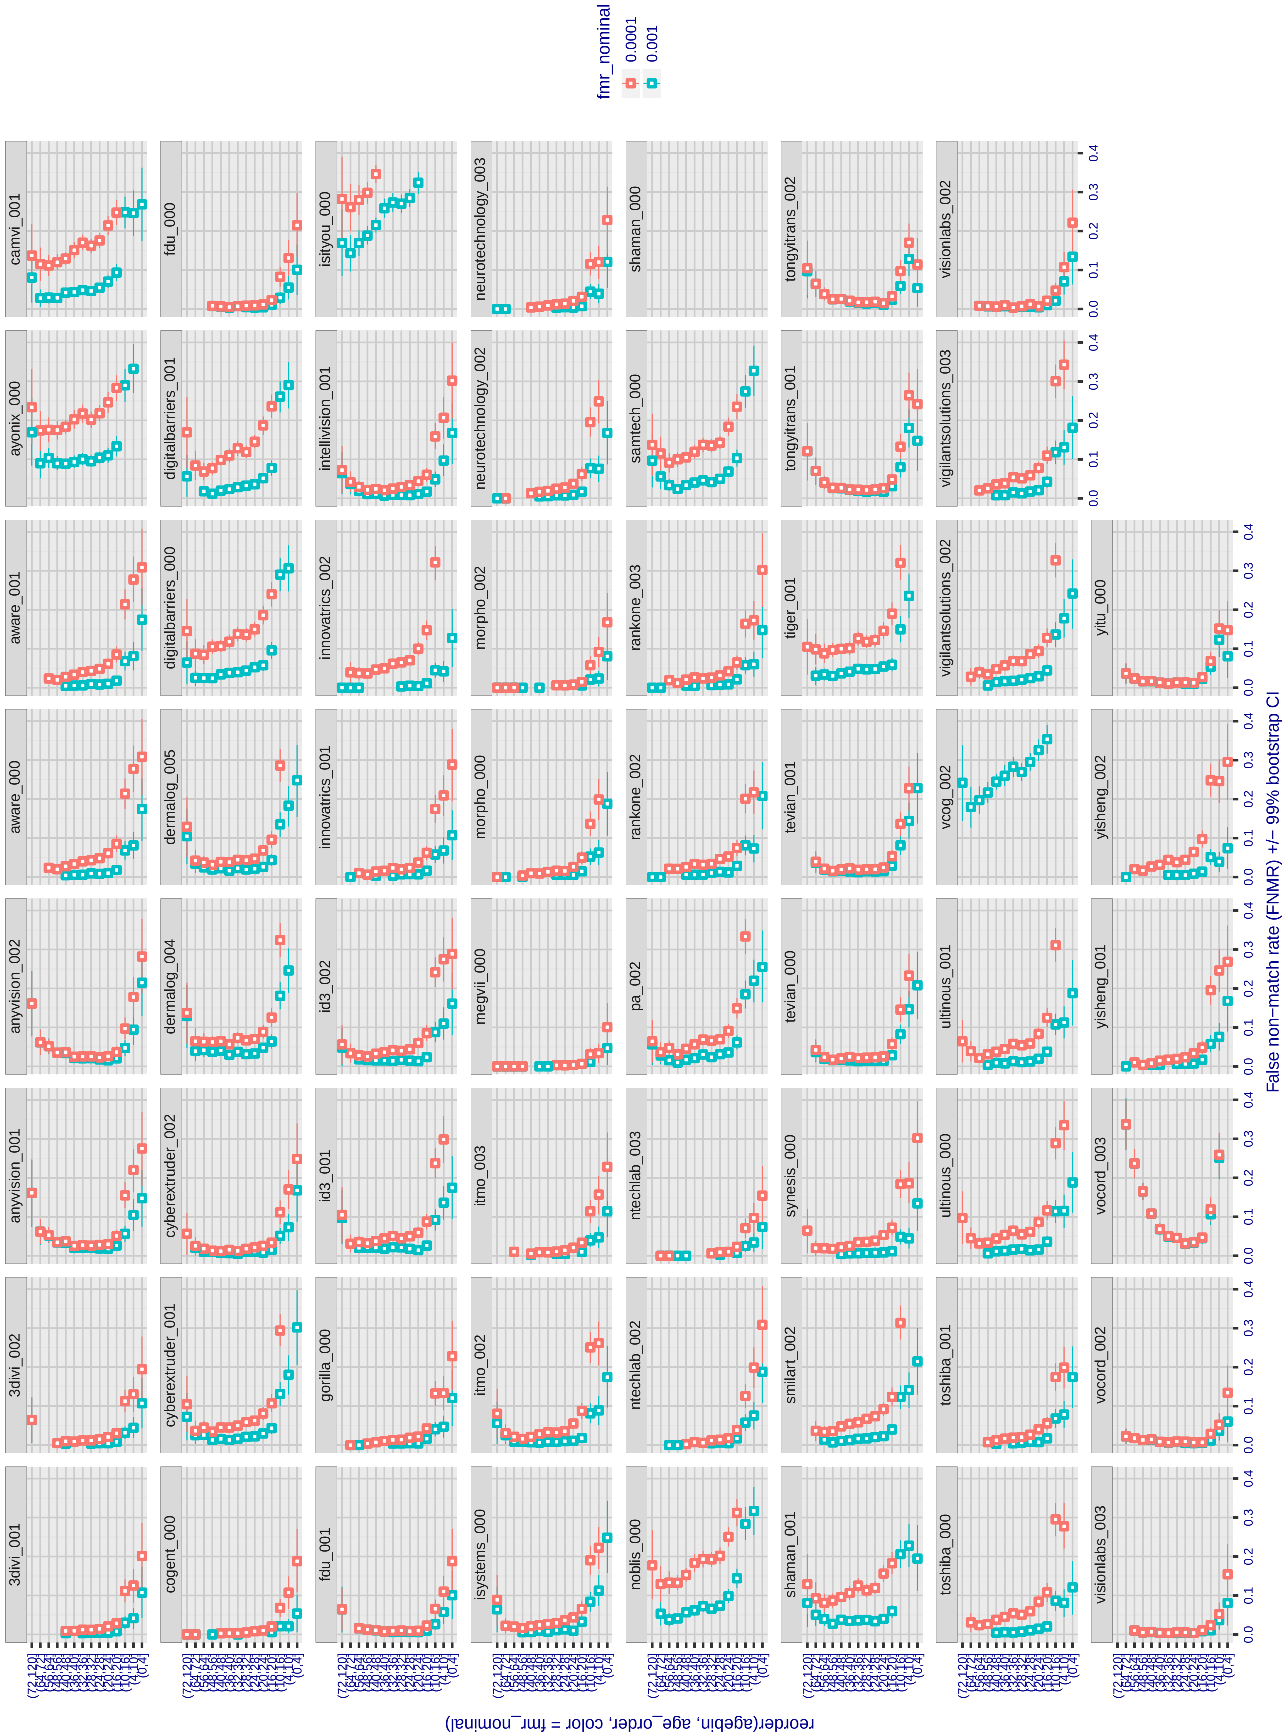

Figure 22: For the visa images, the dots show FNMR by age group for two operating thresholds corresponding to FMR = $\{0.001, 0.0001\}$ computed over all $O(10^{10})$ impostor scores. Given a pair of face images taken at different times, we assign a false non-match to the bin that is the arithmetic average of the subject's ages. This plot shows only the effect of age, not ageing. The number of comparisons in each bin is generally in the thousands. However the FNMR for the first and last bins are each computed over fewer than 150 comparisons.

5.6 Impostor distribution stability

5.6.1 Effect of birth place on the impostor distribution

Background: Facial appearance varies geographically, both in terms of skin tone, cranio-facial structure and size. This section addresses whether false match rates vary intra- and inter-regionally.

Goals:

- ▷ To show the effect of birth region of the impostor and enrollee on false match rates.
- ▷ To determine whether some algorithms give better impostor distribution stability.

Methods:

- ▷ For the visa images, NIST defined 10 regions: Sub-Saharan Africa, South Asia, Polynesia, North Africa, Middle East, Europe, East Asia, Central and South America, Central Asia, and the Caribbean.
- ▷ For the visa images, NIST mapped each country of birth to a region. There is some arbitrariness to this. For example, Egypt could reasonably be assigned to the Middle East instead of North Africa. An alternative methodology could, for example, assign the Philippines to *both* Polynesia and East Asia.
- ▷ FMR is computed for cases where all face images of impostors born in region r_2 are compared with enrolled face images of persons born in region r_1.

$$\text{FMR}(r_1, r_2, T) = \frac{\sum_{i=1}^{N_{r_1,r_2}} H(s_i - T)}{N_{r_1,r_2}} \quad (5)$$

  where the same threshold, T, is used in all cells, and H is the unit step function. The threshold is set to give FMR(T) = 0.001 over the entire set of visa image impostor comparisons.
- ▷ This analysis is then repeated by country-pair, but only for those country pairs where both have at least 1000 images available. The countries[1] appear in the axes of graphs that follow.
- ▷ The mean number of impostor scores in any cross-region bin is 33 million. The smallest number of impostor scores in any bin is 135000, for Central Asia - North Africa. While these counts are large enough to support reasonable significance, the number of individual faces is much smaller, O($N^{0.5}$).
- ▷ The numbers of impostor scores in any cross-country bin is shown in Figure 136.

Results: Subsequent figures show heatmaps that use color to represent the base-10 logarithm of the false match rate. Red colors indicate high (bad) false match rates. Dark colors indicate benign false match rates. There are two series of graphs corresponding to aggregated geographical regions, and to countries. The notable observations are:

- ▷ The on-diagonal elements correspond to within-region impostors. FMR is generally above the nominal value of FMR = 0.001. Particularly there is usually higher FMR in, Sub-Saharan Africa, South Asia, and the Caribbean. Europe and Central Asia, on the other hand, usually give FMR closer to the nominal value.
- ▷ The off-diagonal elements correspond to across-region impostors. The highest FMR is produced between the Caribbean and Sub-Saharan Africa.
- ▷ Algorithms vary.

[1]These are Argentina, Australia, Brazil, Chile, China, Costa Rica, Cuba, Czech Republic, Dominican Republic, Ecuador, Egypt, El Salvador, Germany, Ghana, Great Britain, Greece, Guatamala, Haiti, Hong Kong, Honduras, Indonesia, India, Israel, Jamaica, Japan, Kenya, Korea, Lebanon, Mexico, Malaysia, Nepal, Nigeria, Peru, Philippines, Pakistan, Poland, Romania, Russia, South Africa, Saudi Arabia, Thailand, Trinidad, Turkey, Taiwan, Ukraine, Venezuela, and Vietnam.

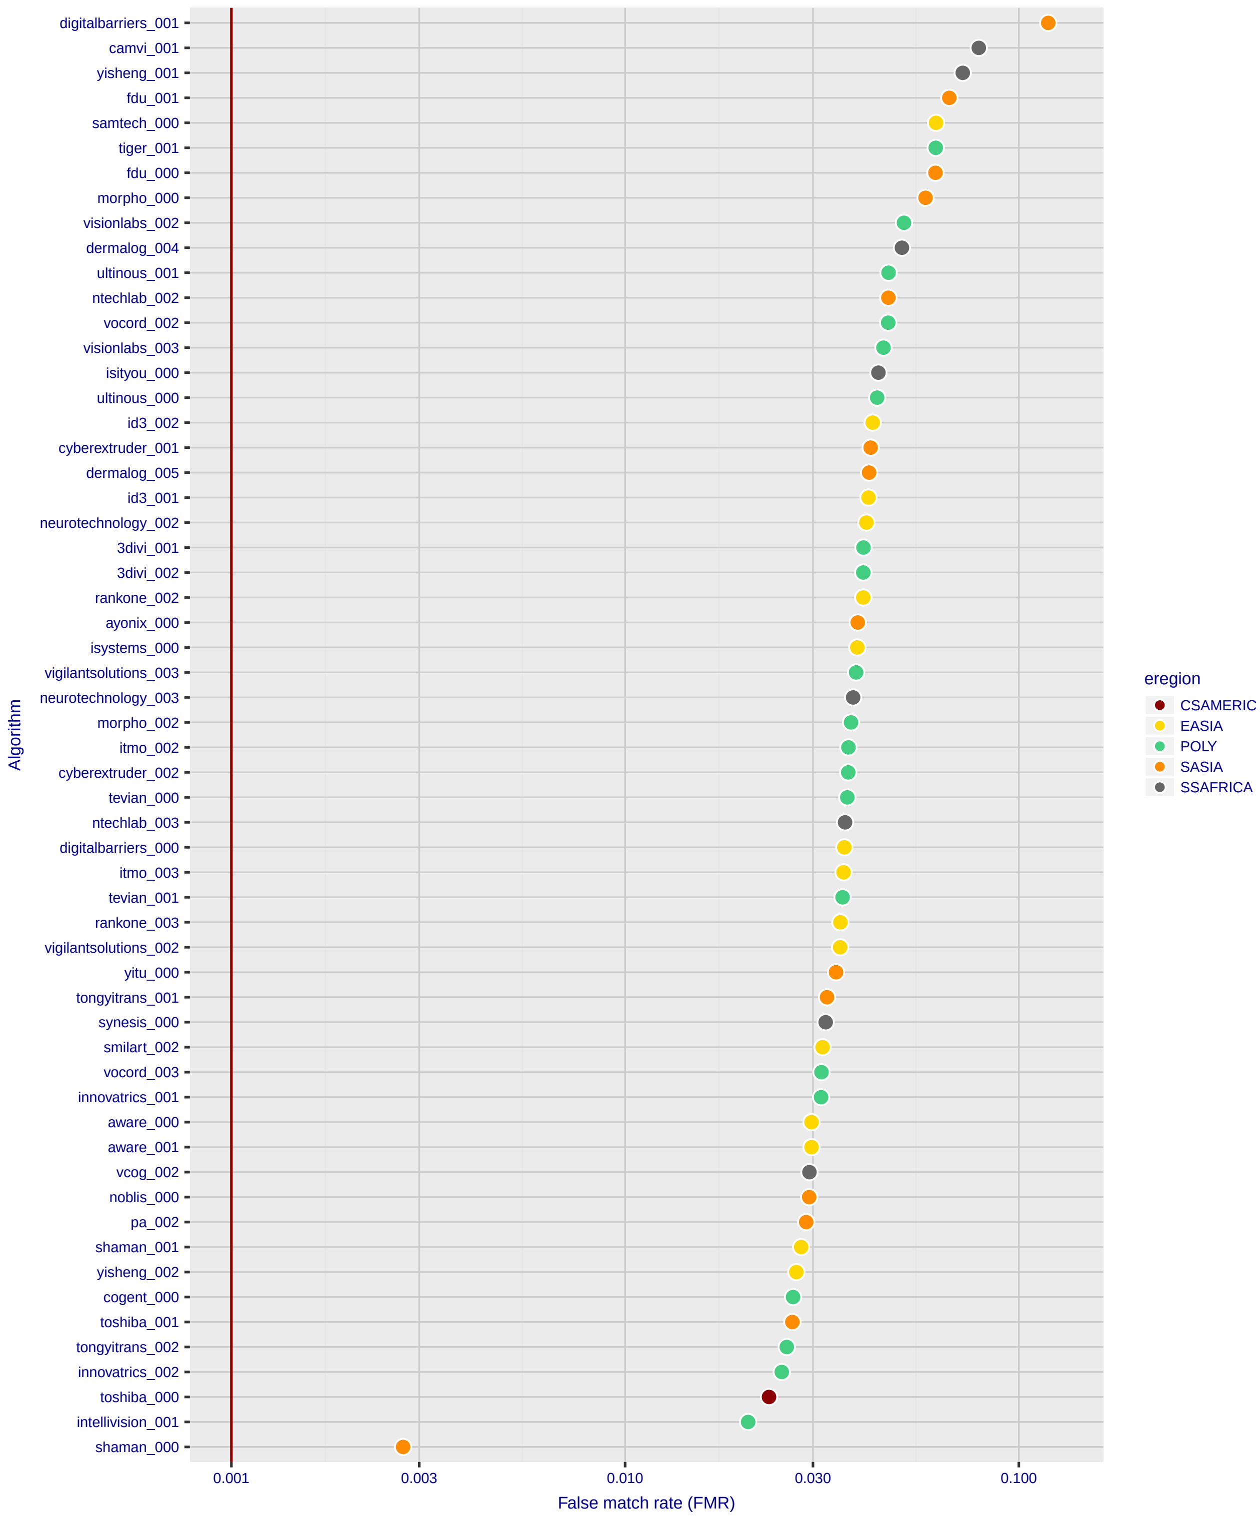

Figure 23: For the visa images, the dots show FMR for impostor comparisons of individuals of the same sex and same age group for the region of the world that gives the worst (highest) FMR when the threshold is set to give FMR = 0.001 (red vertical line) over all $O(10^{10})$ impostor scores i.e. zero-effort. The shift of the dots to right shows massive increases in FMR when impostors have the same sex, age, and region of birth. The color code indicates which region gives the worst case FMR. If the observed variation is due to the prevalence of one kind of images in the training imagery, then algorithms developed on one kind of data might be expected to give higher FMR on other kinds.

▷ We computed the same quantities for a global FMR = 0.0001. The effects are similar.

Caveats:

▷ The effects of variable impostor rates on one-to-many identification systems may well differ from what's implied by these one-to-one verification results. Two reasons for this are a) the enrollment galleries are usually imbalanced across countries of birth, age and sex; b) one-to-many identification algorithms often implement techniques aimed at stabilizing the impostor distribution. Further research is necessary.

▷ In principle, the effects seen in this subsection could be due to differences in the image capture process. We consider this unlikely since the effects are maintained across geography - e.g. Caribbean vs. Africa, or Japan vs. China.

Figure 24: *For visa images, the heatmap shows how the mean of the impostor distribution for the country pair (a,b) is shifted relative to the mean of the global impostor distribution, expressed as a number of standard deviations of the global impostor distribution. This statistic is designed to show shifts in the entire impostor distribution, not just tail effects that manifest as the anomalously high (or low) false match rates that appear in the subsequent figures. The countries are chosen to show that skin tone alone does not explain impostor distribution shifts. The reduced shift in Asian populations with the Yitu and TongYiTrans algorithms, is accompanied by positive shifts in the European populations. This reversal relative to most other algorithms, may derive from use of nationally weighted training sets. The Visionlabs algorithm appears most insensitive to country effects. The figure is computed from same-sex and same-age impostor pairs.*

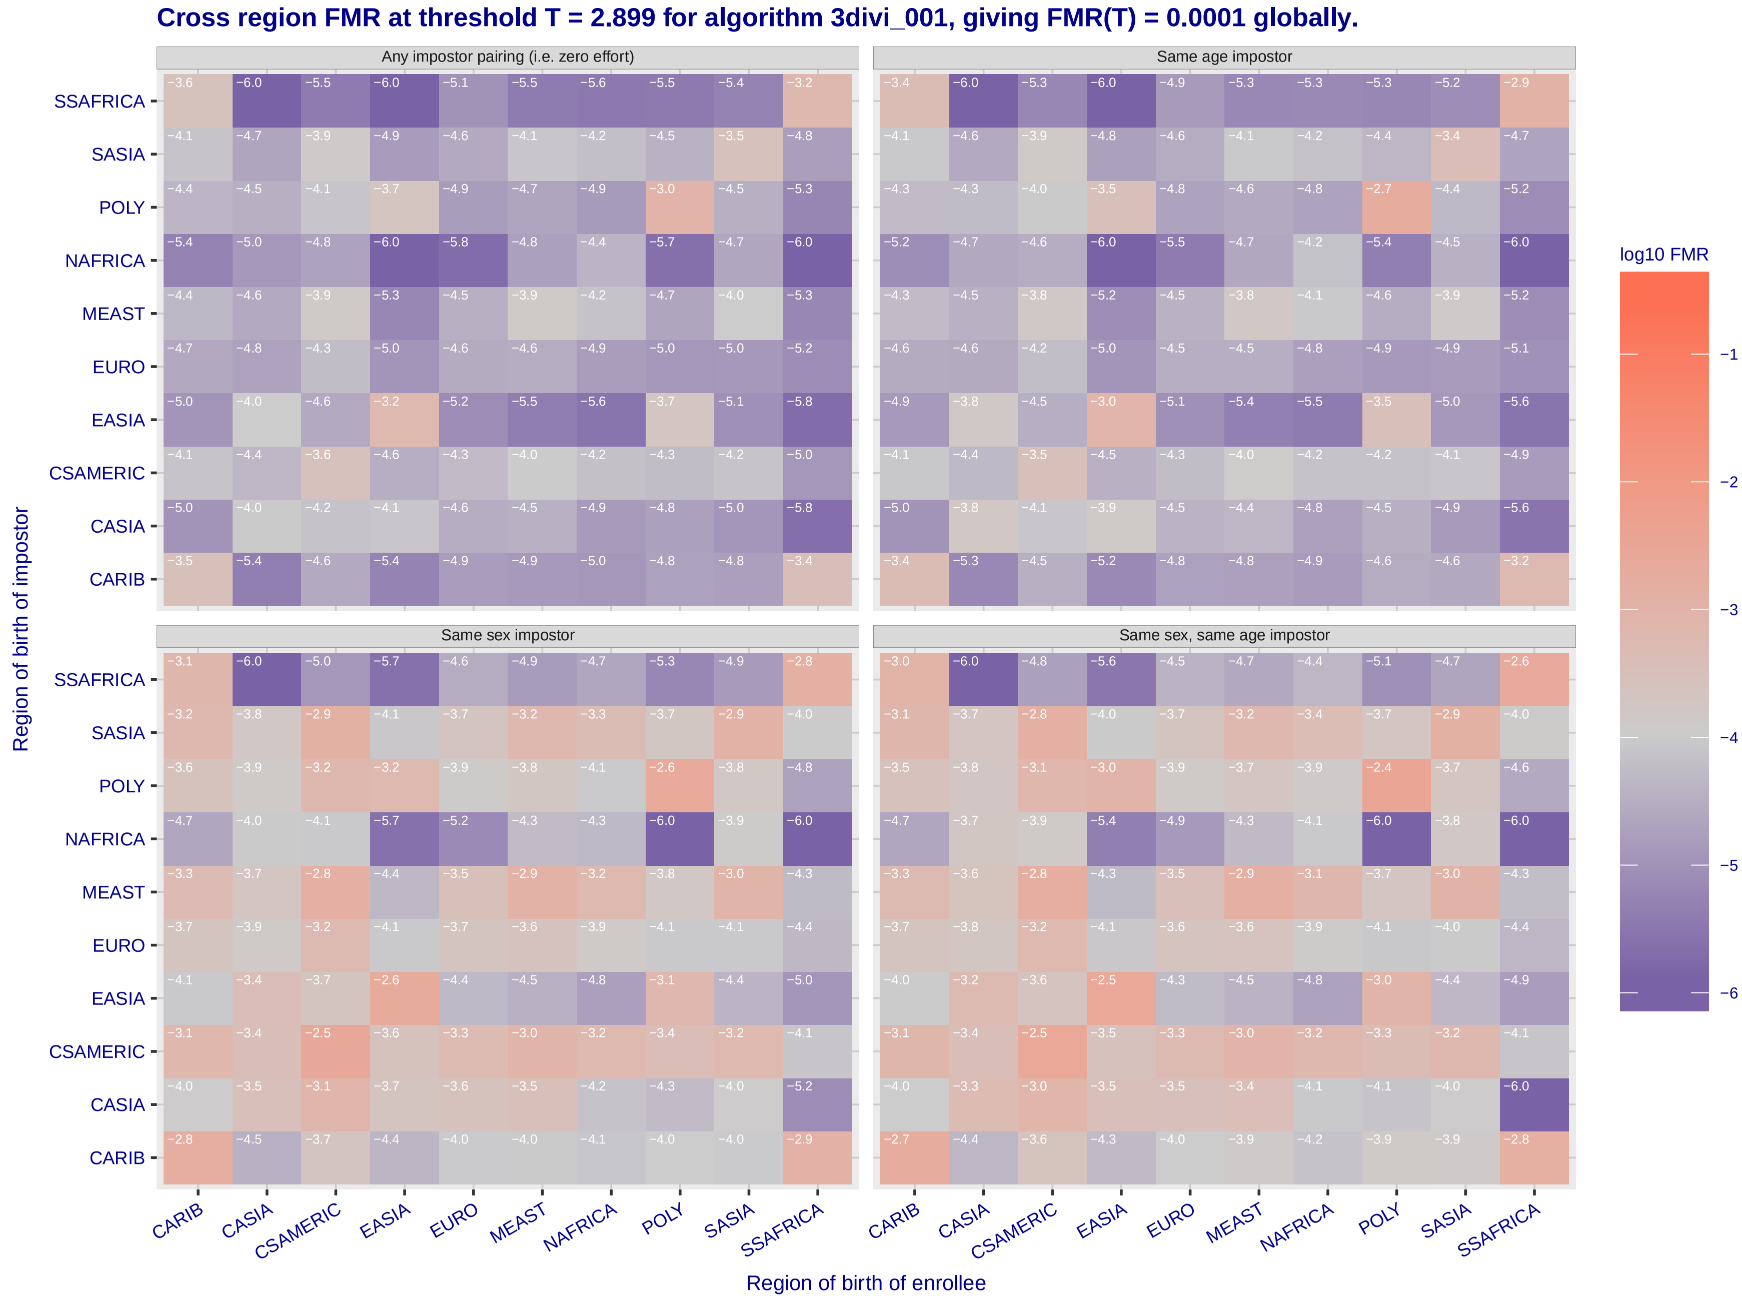

Figure 25: For algorithm 3divi-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

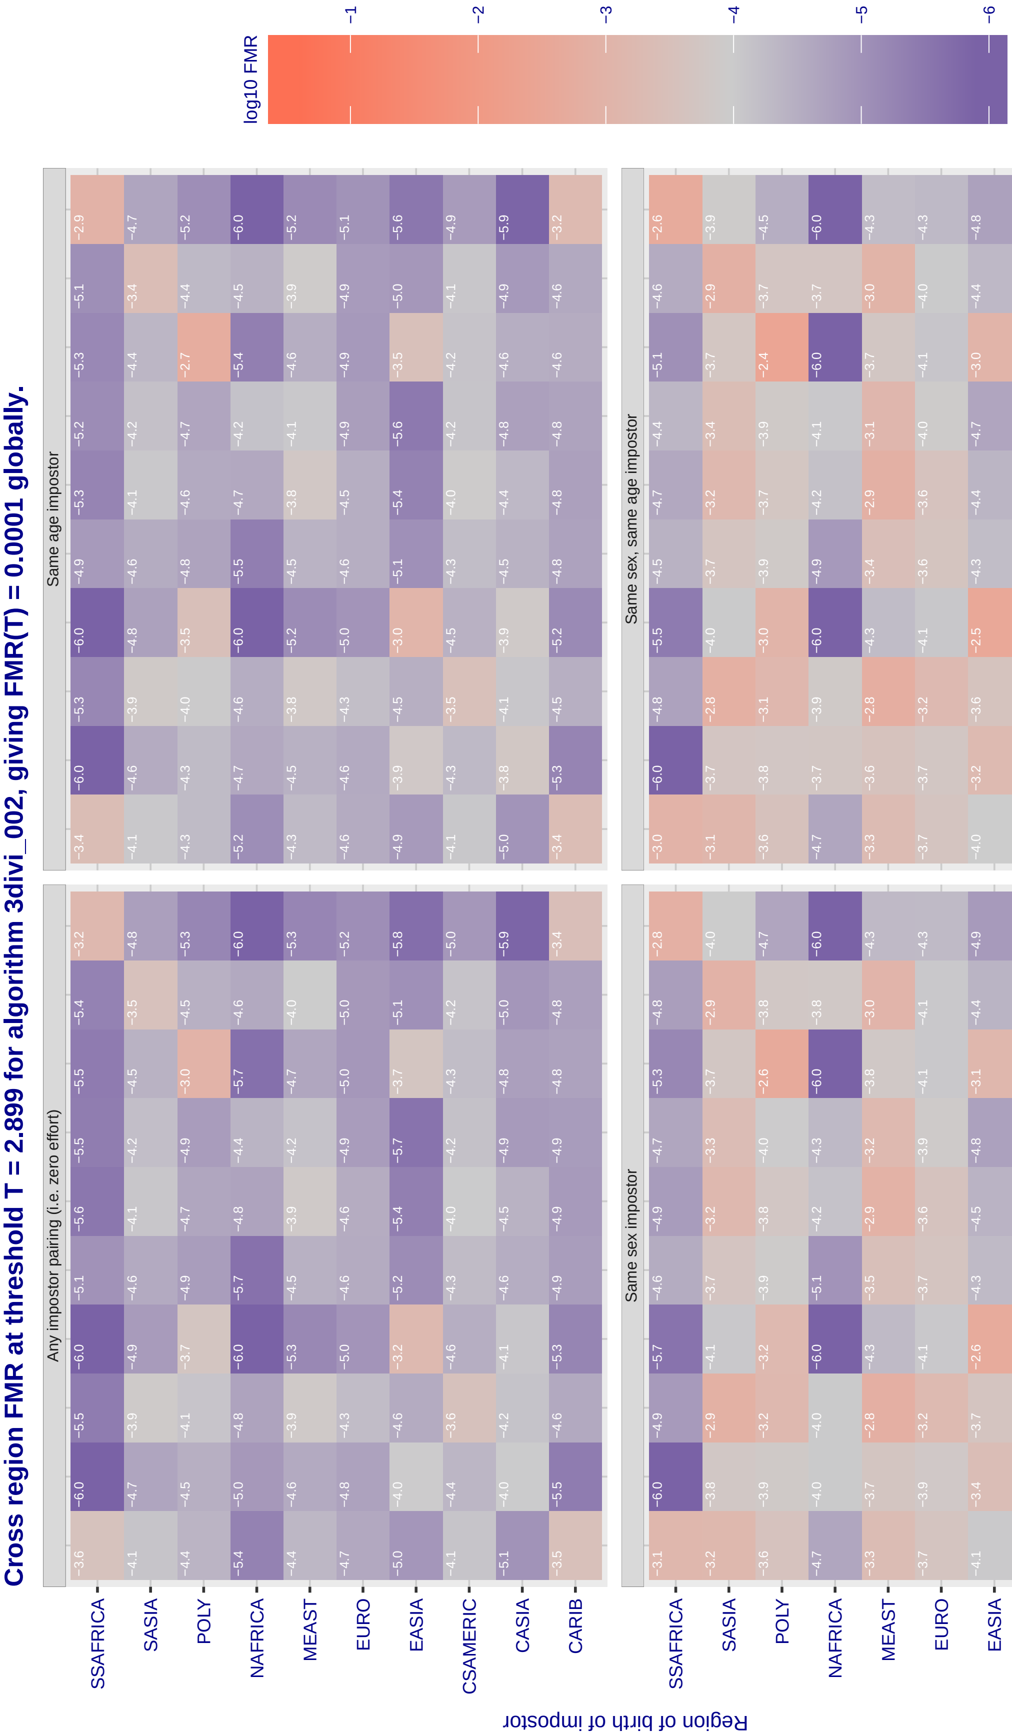

Figure 26: For algorithm 3divi-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

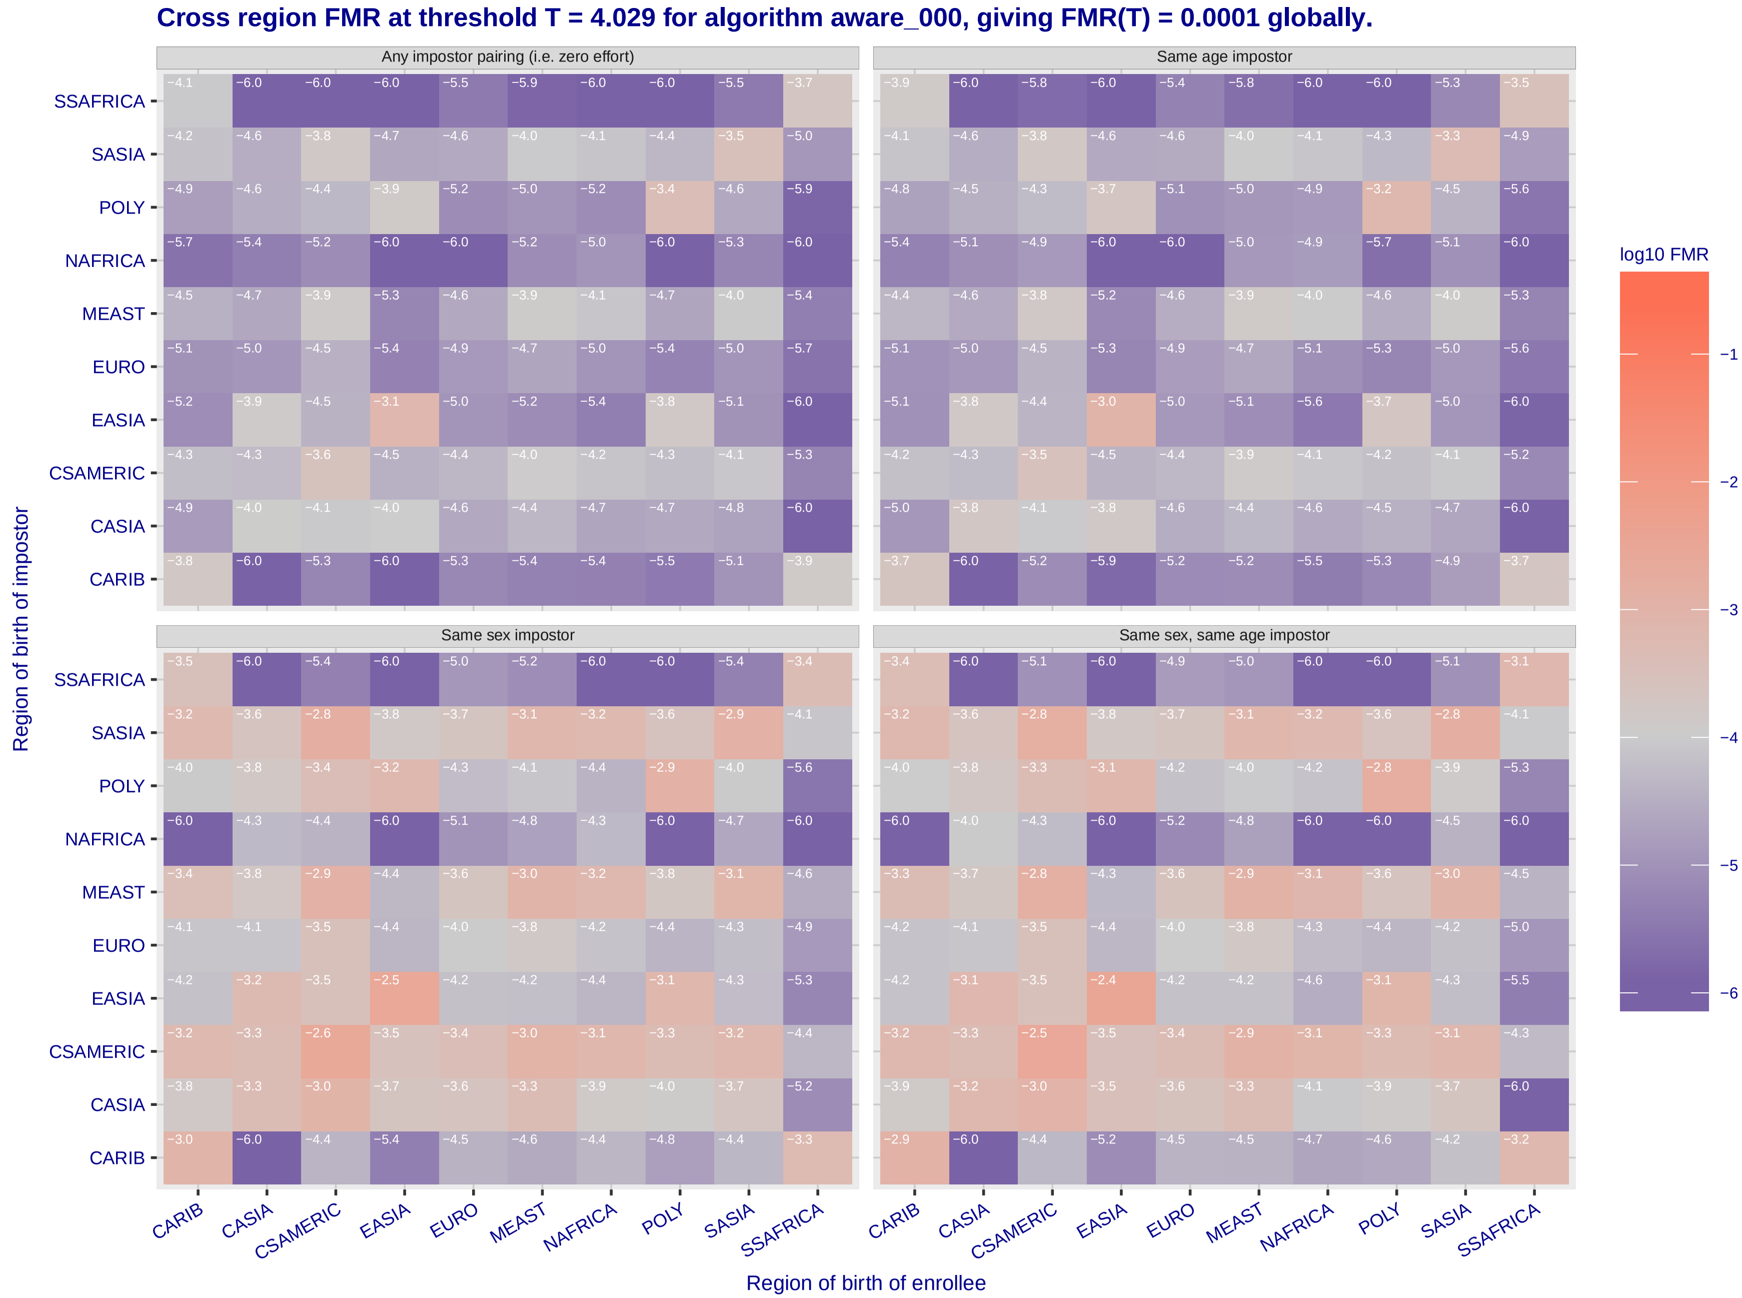

Figure 27: For algorithm aware-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

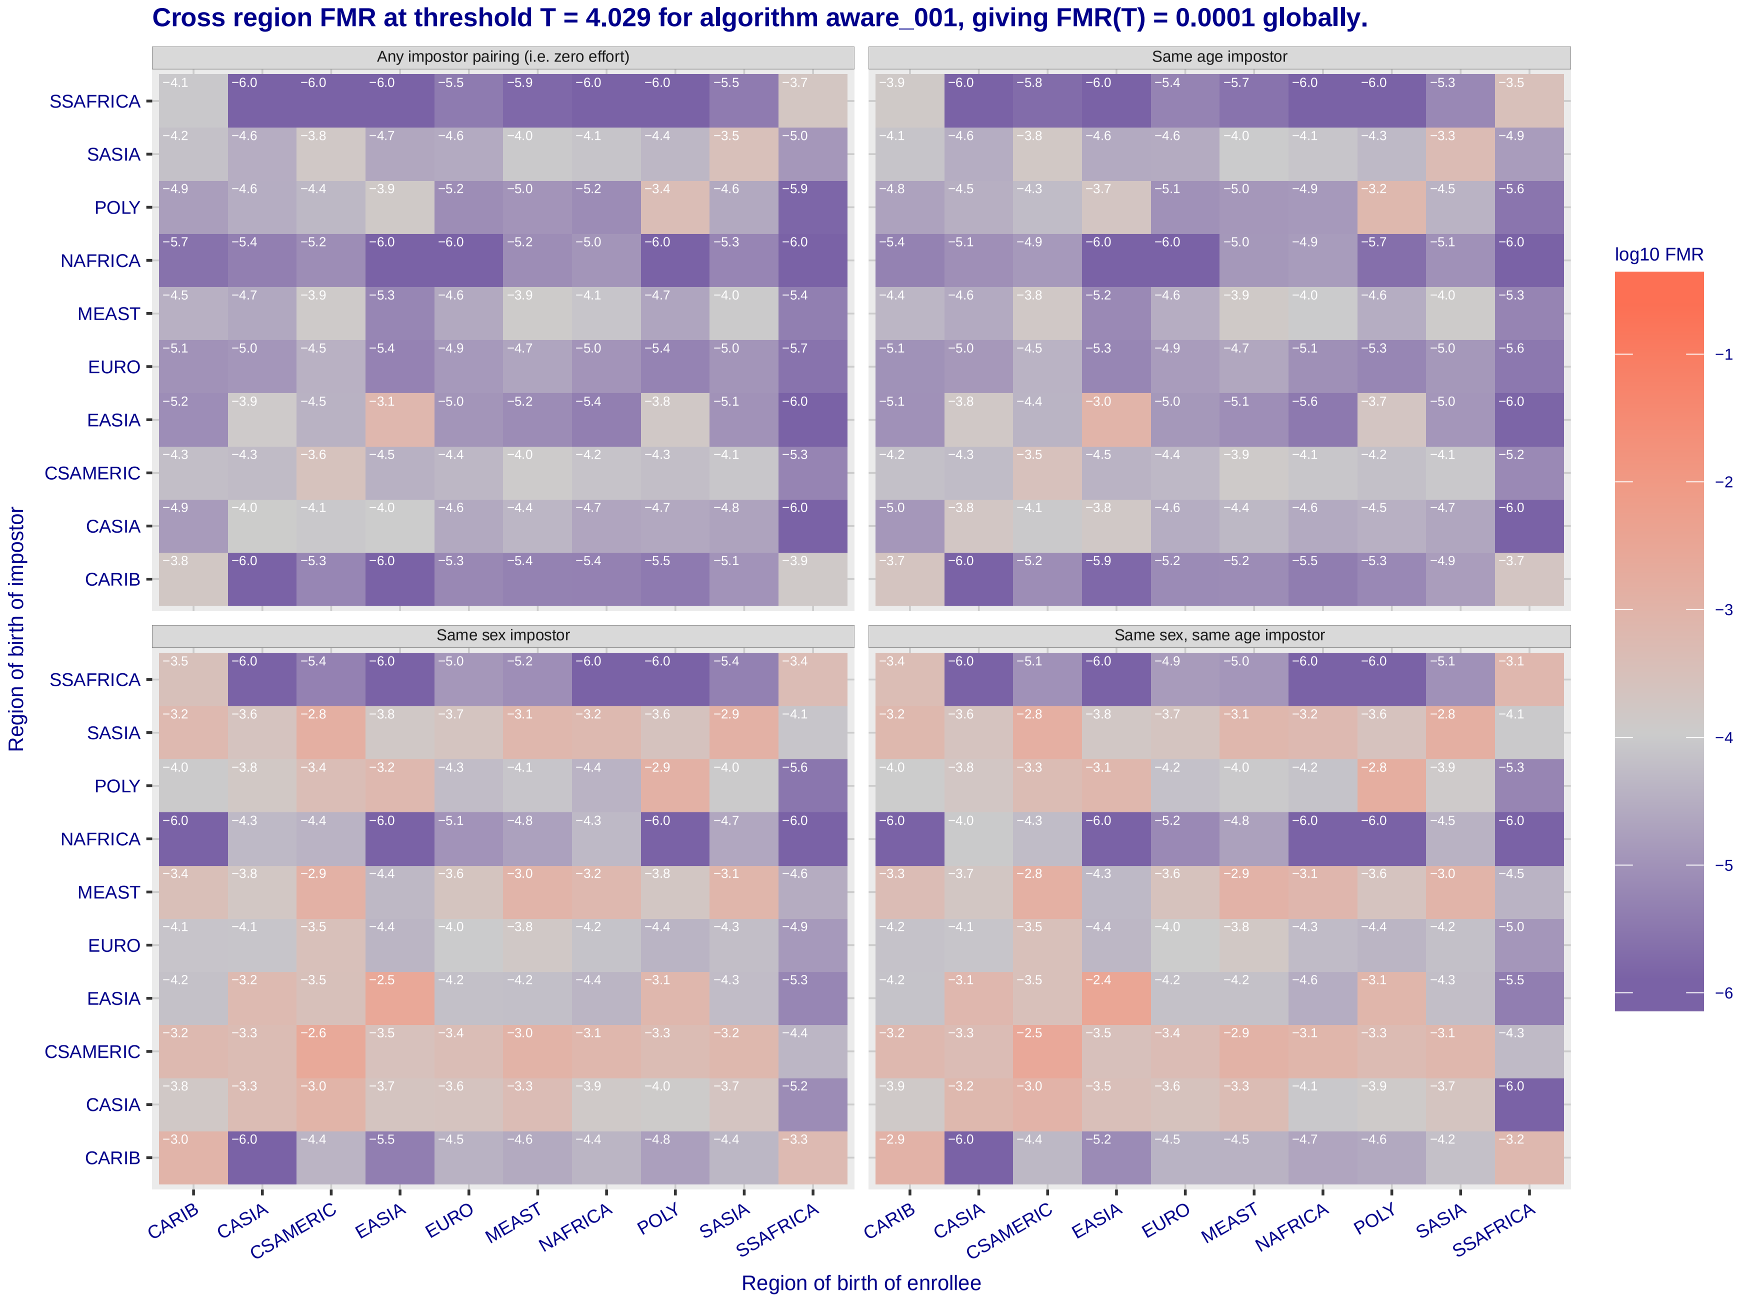

Figure 28: For algorithm aware-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

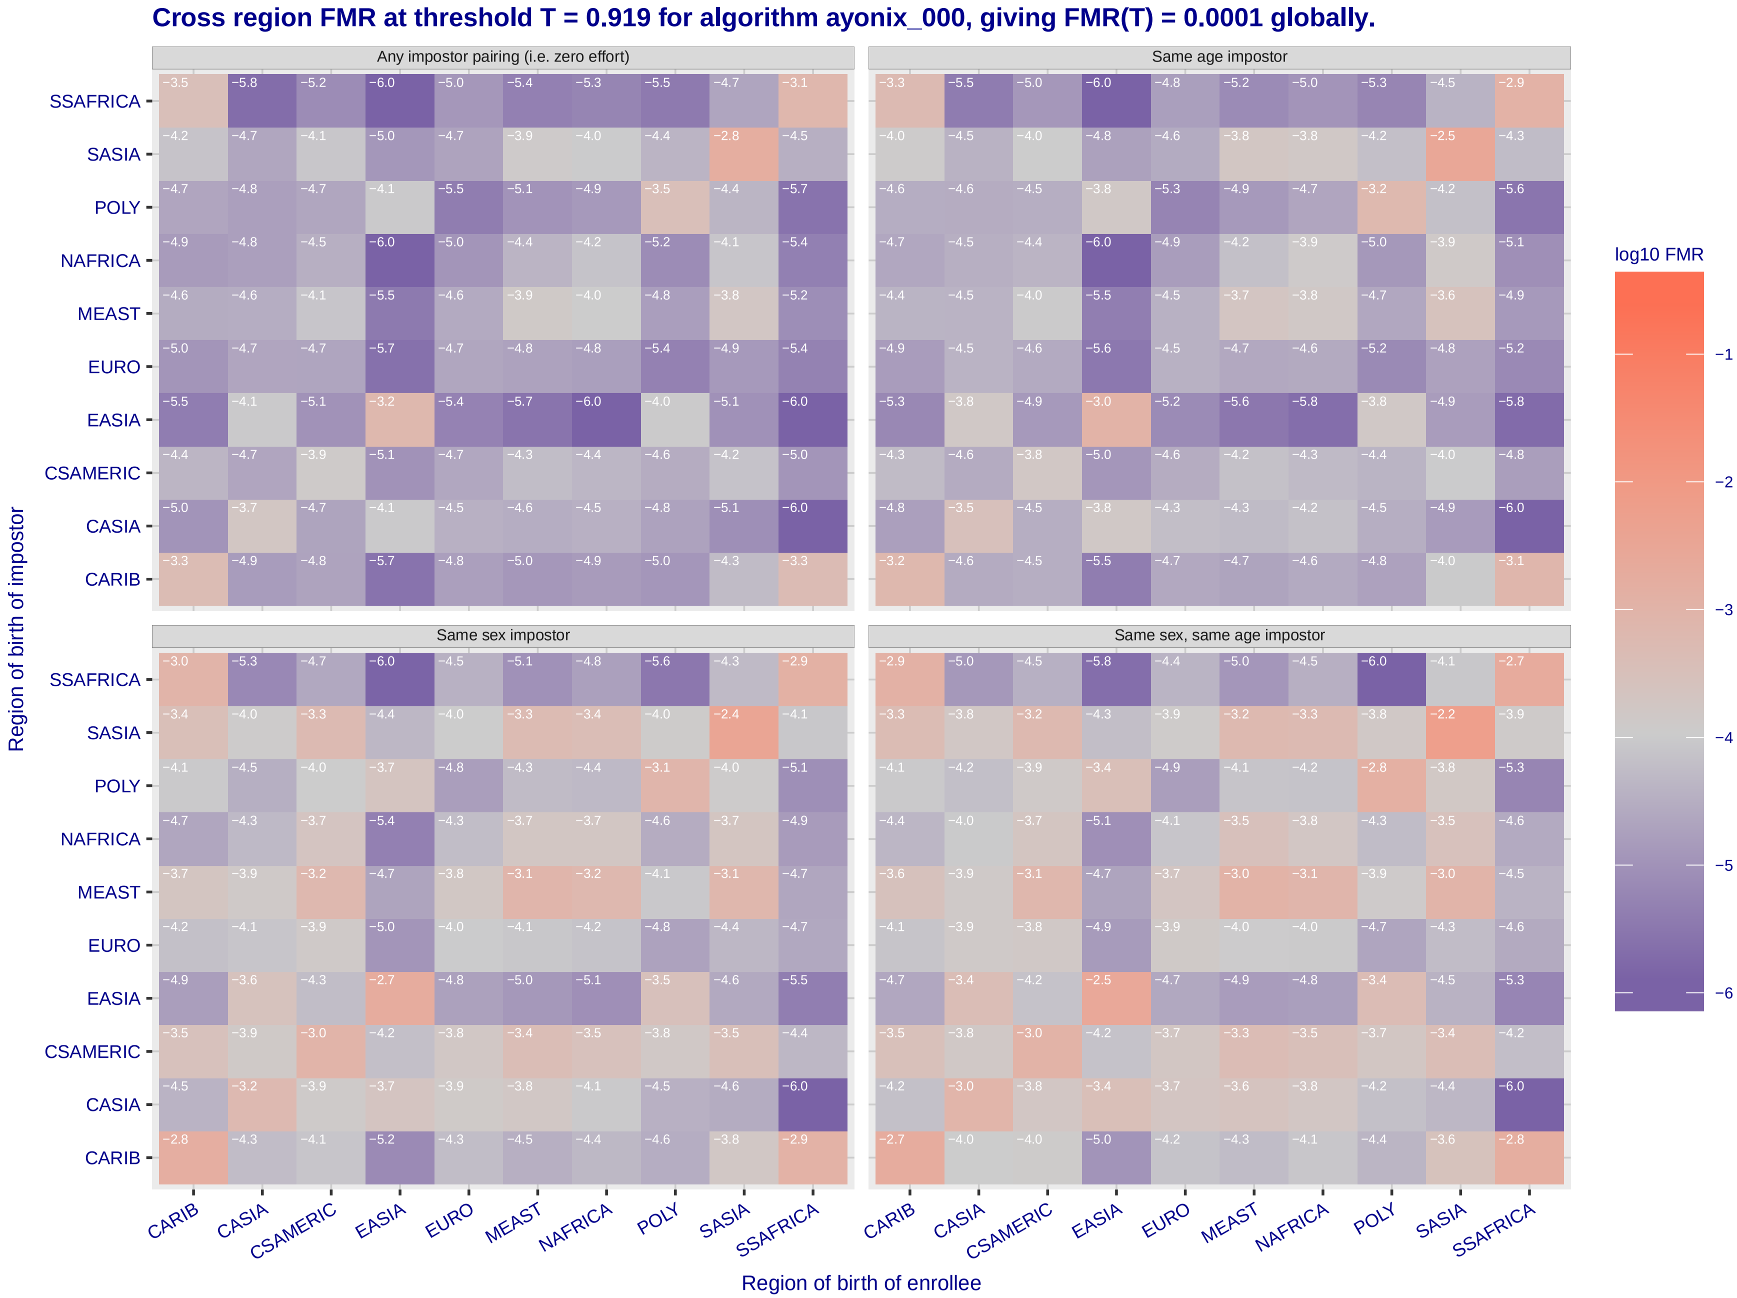

Figure 29: For algorithm ayonix-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

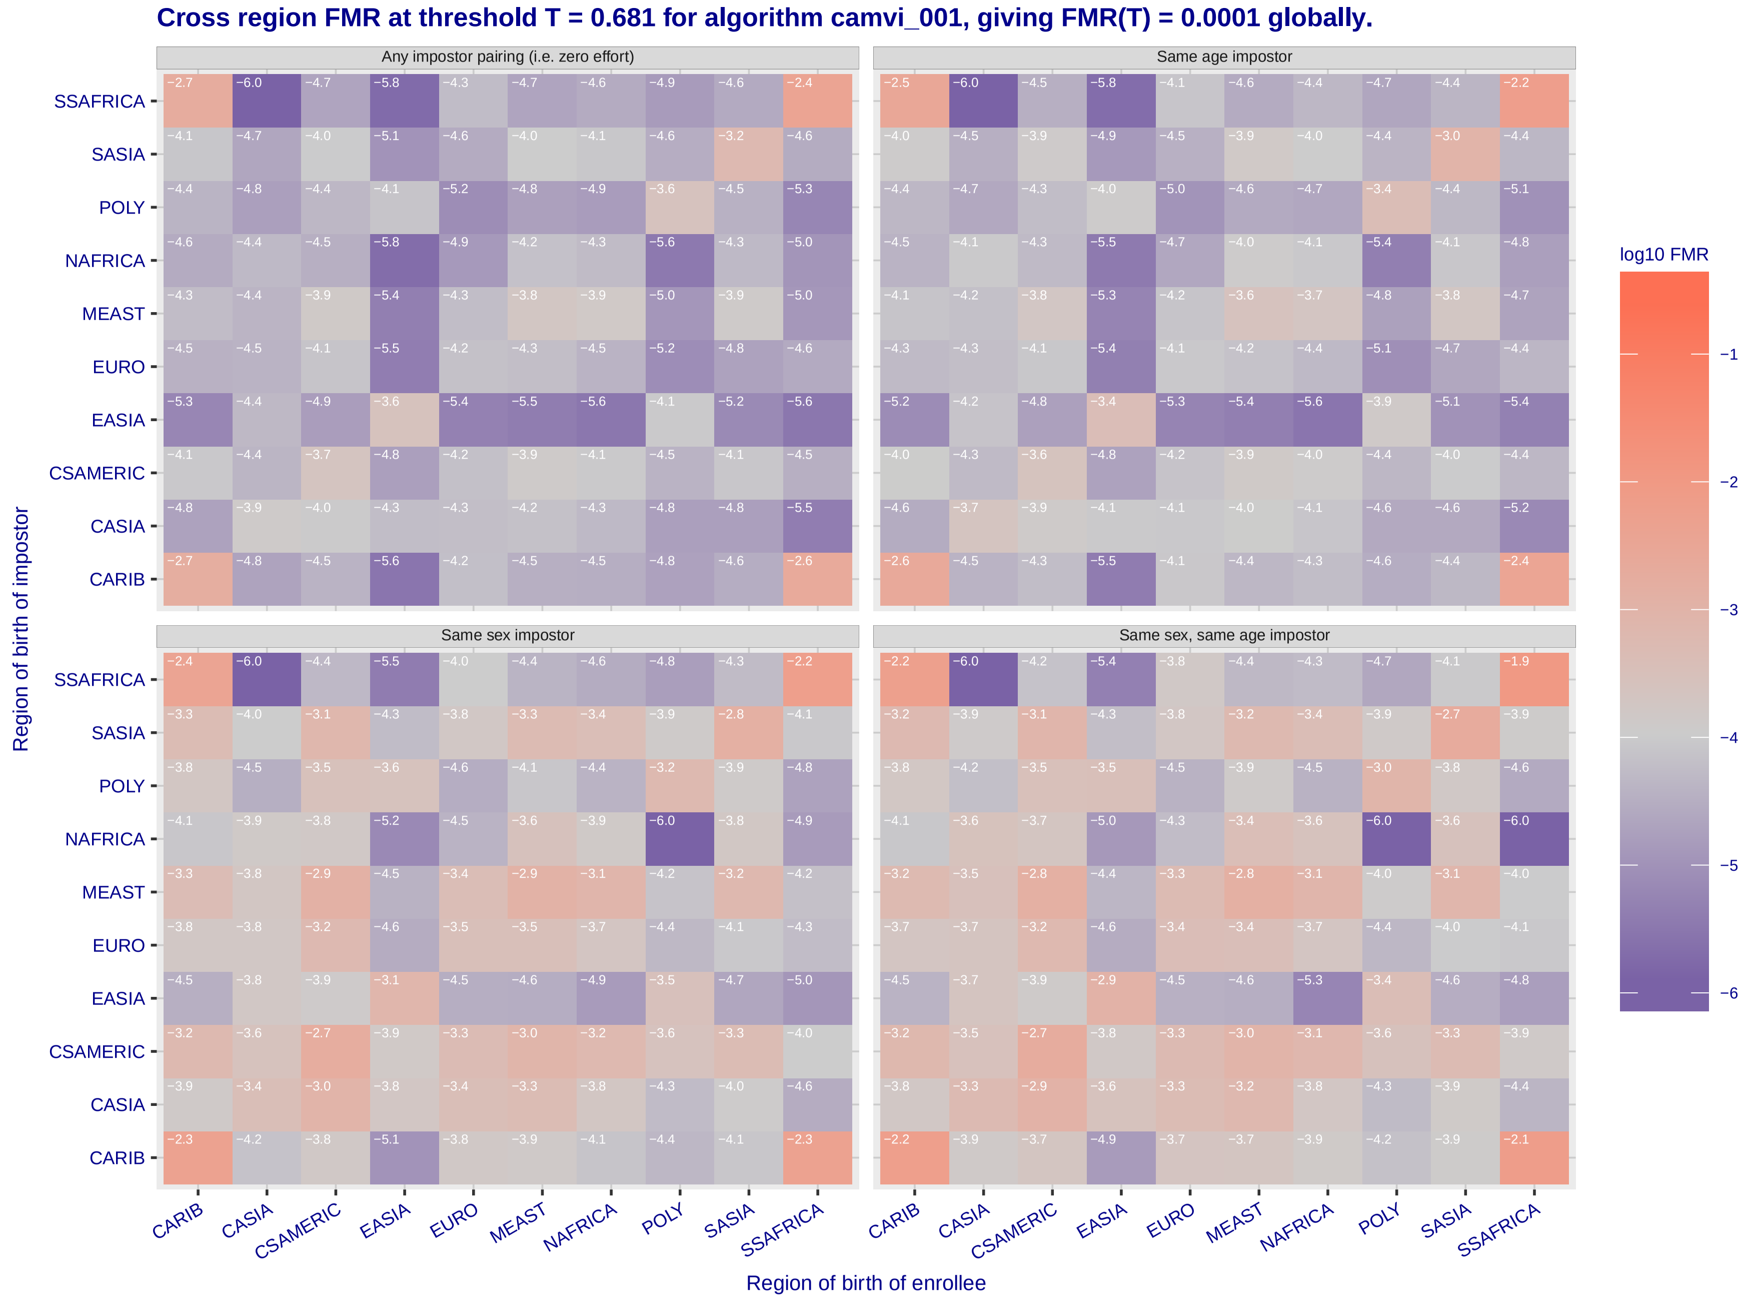

Figure 30: For algorithm camvi-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

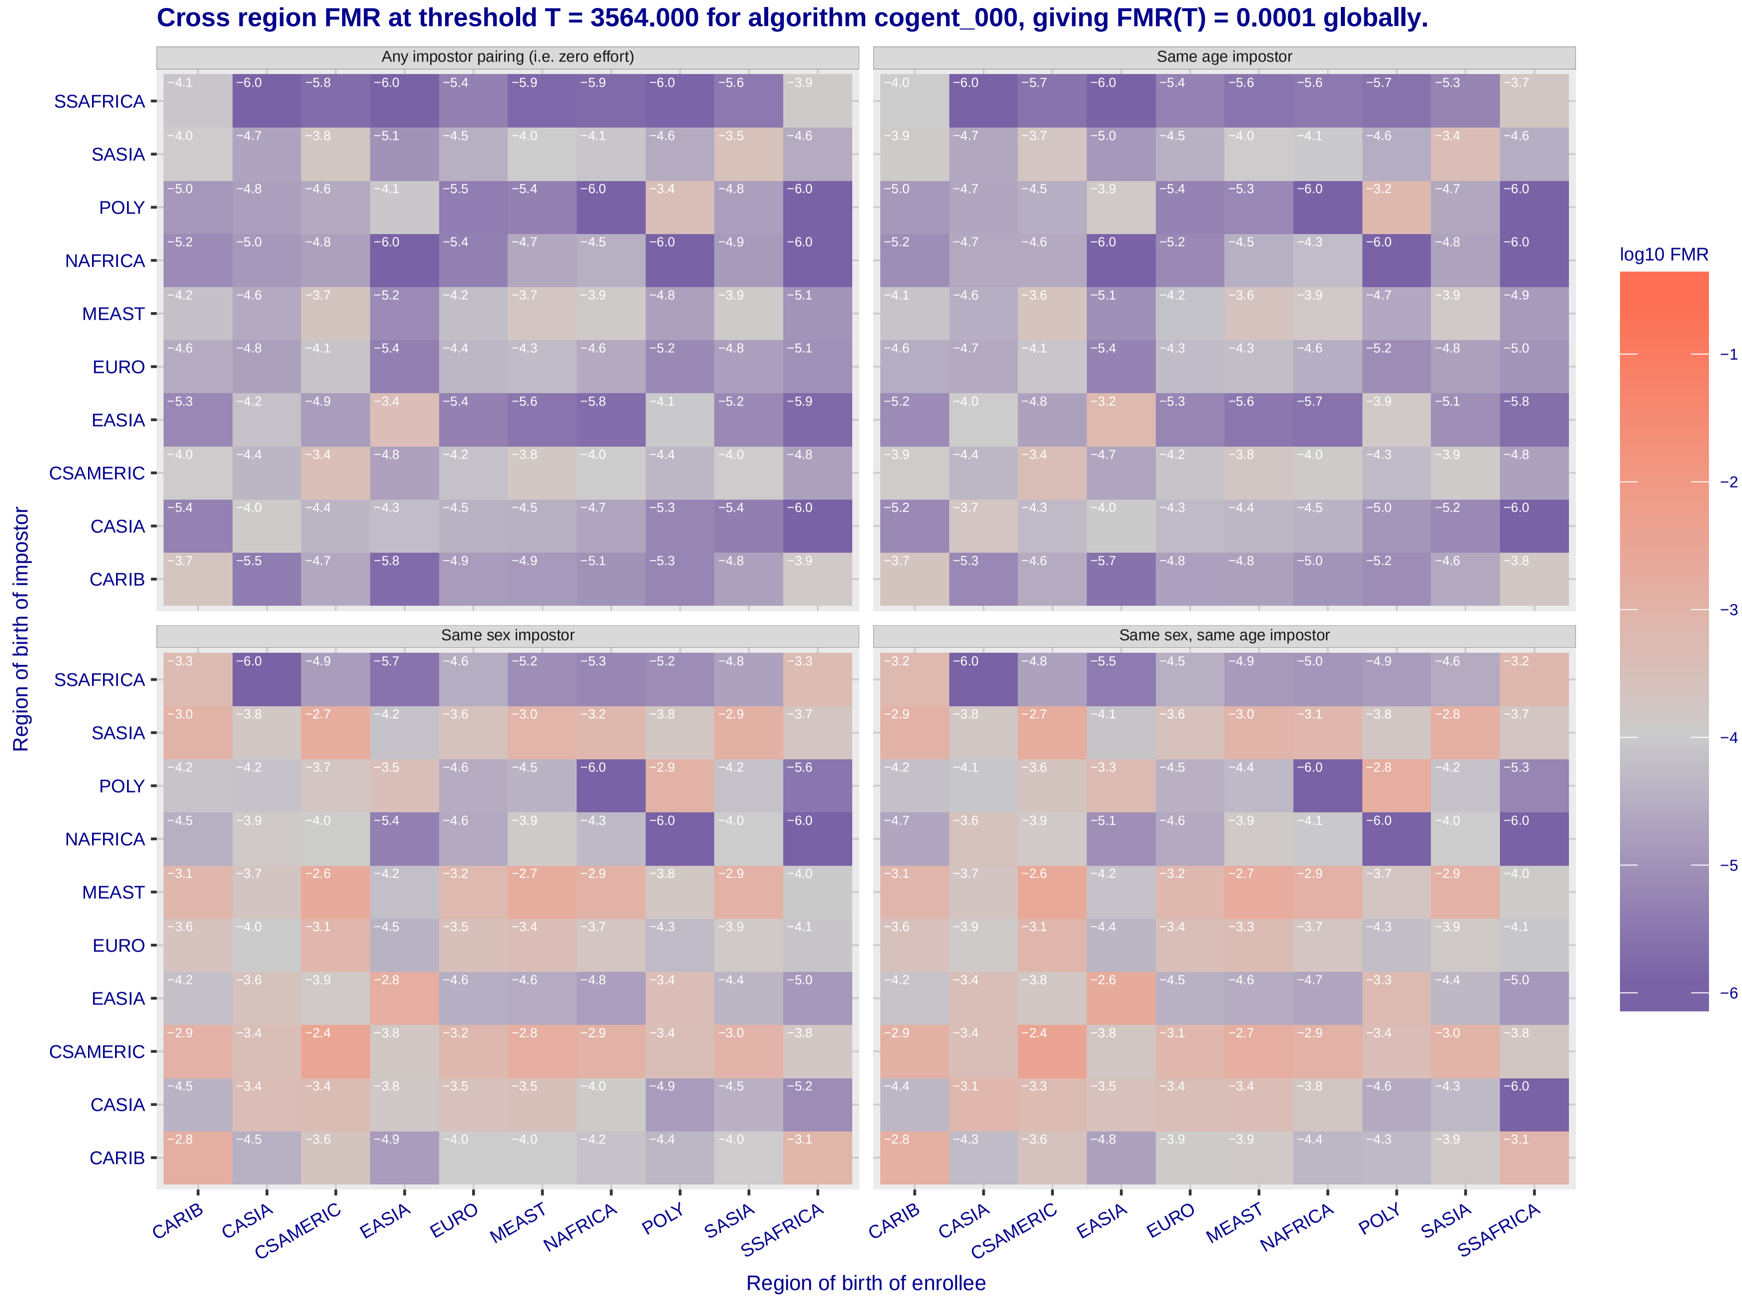

Figure 31: For algorithm cogent-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

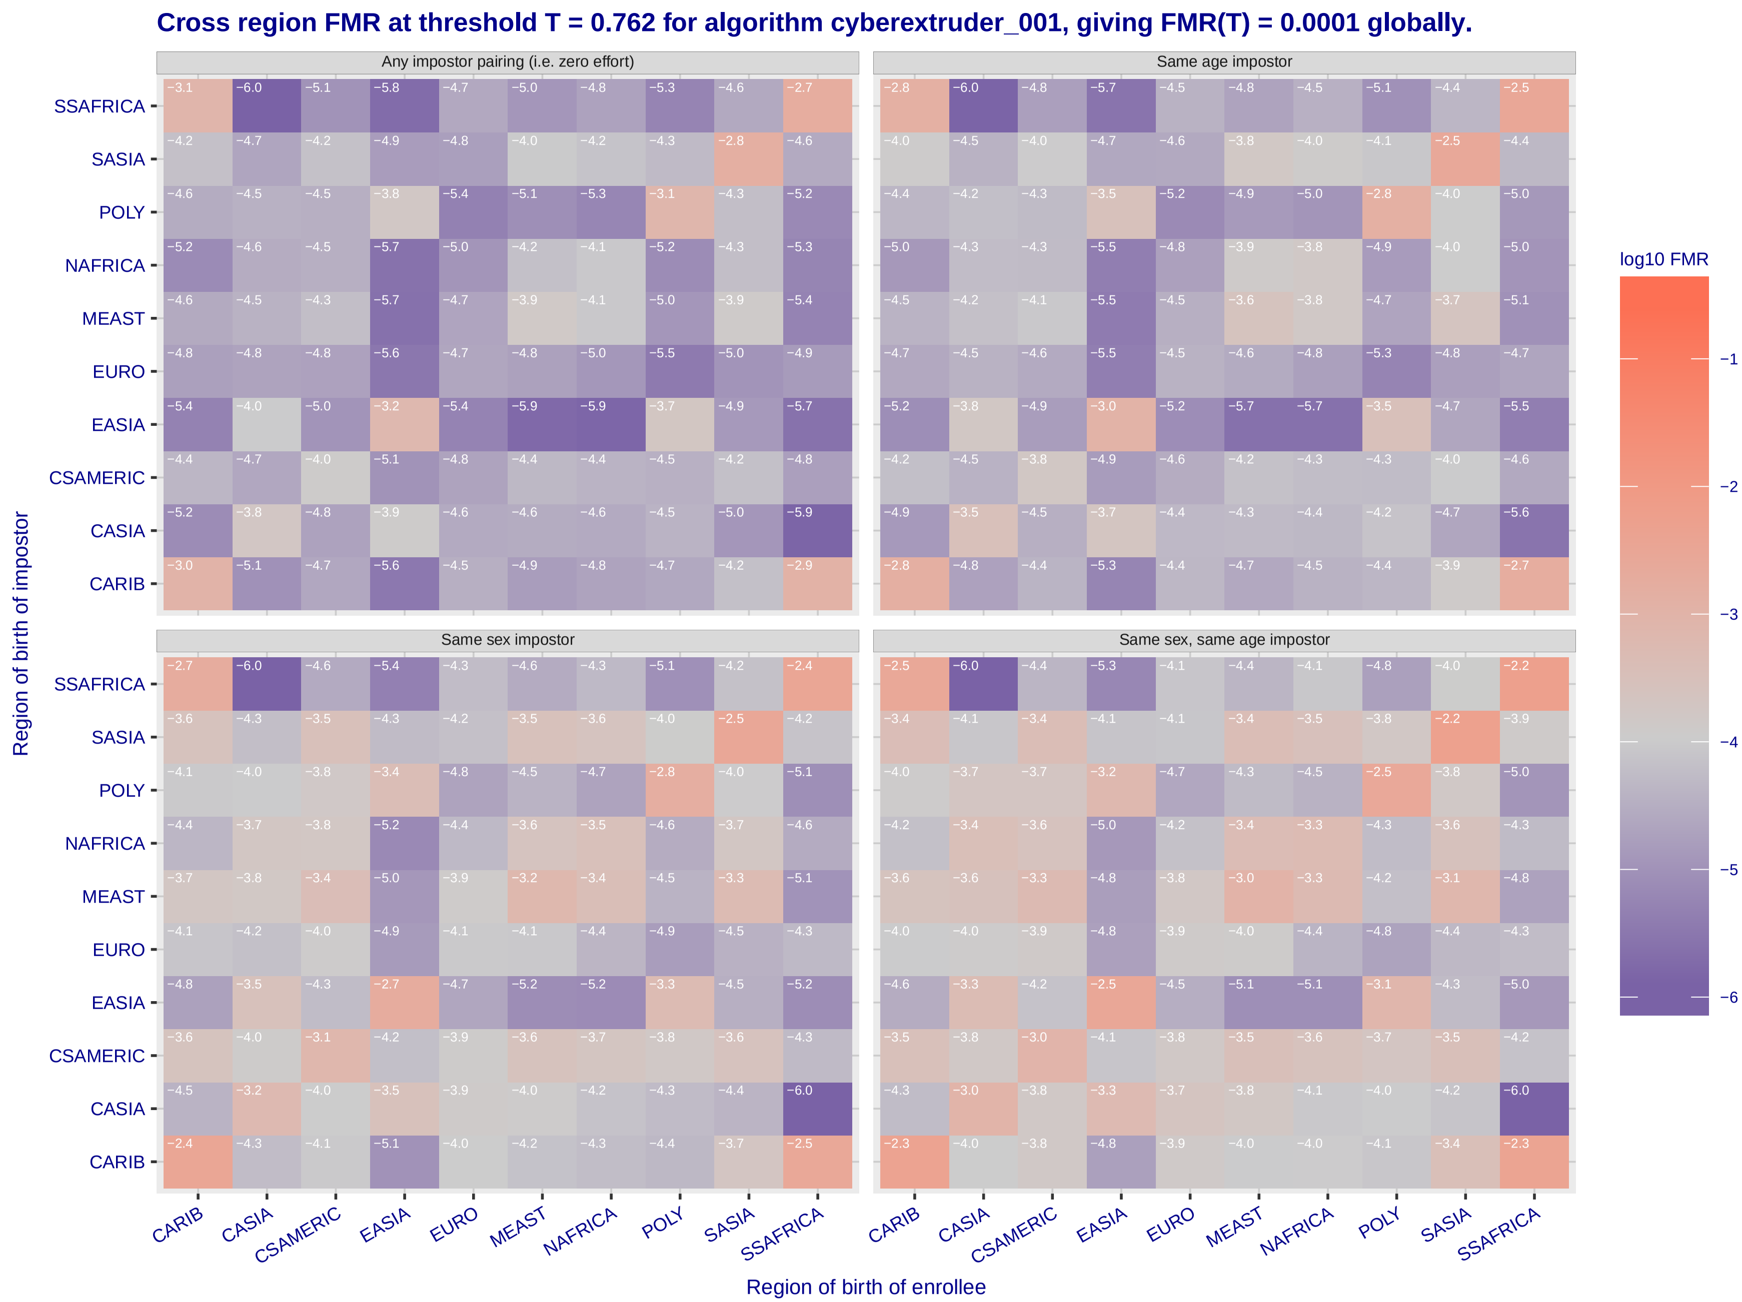

Figure 32: *For algorithm cyberextruder-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.*

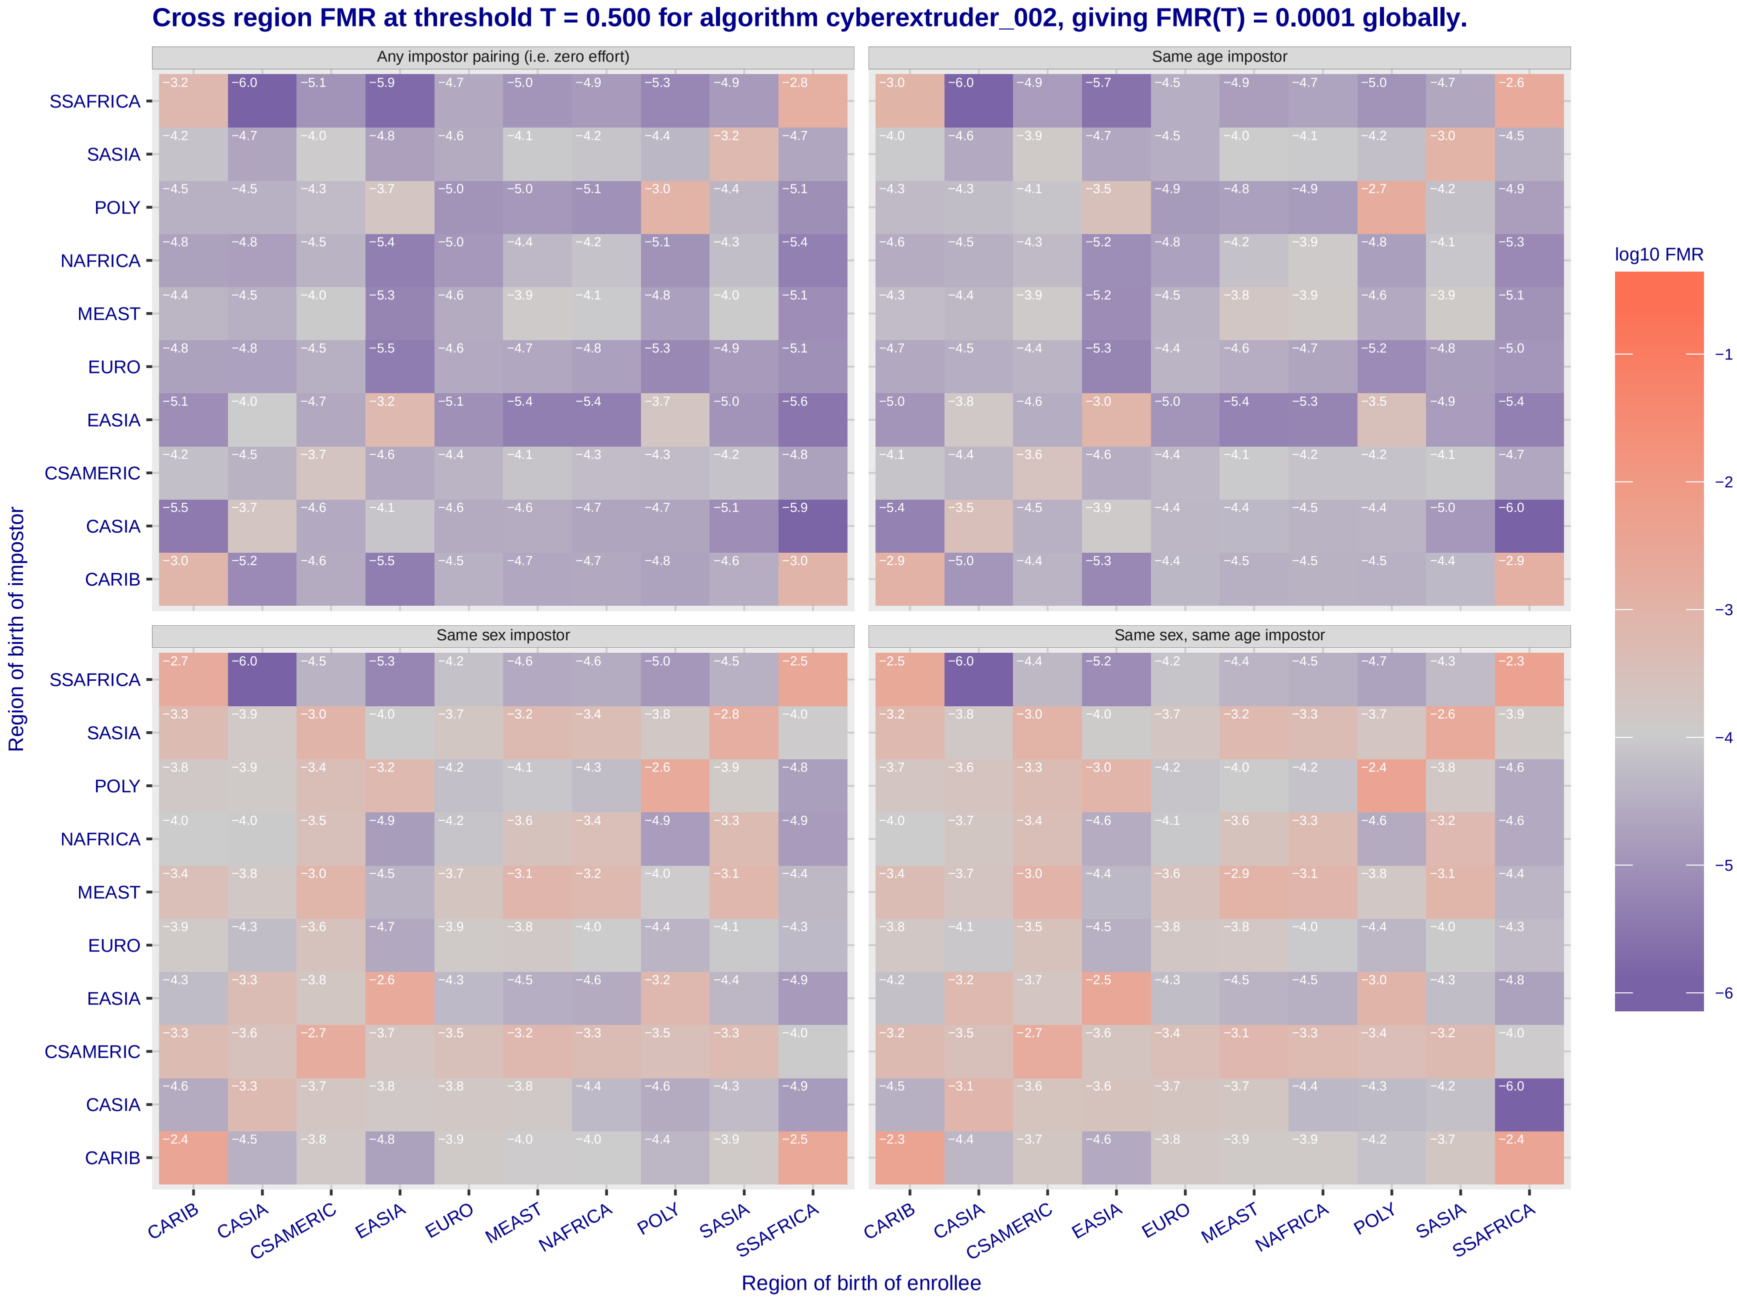

Figure 33: For algorithm cyberextruder-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

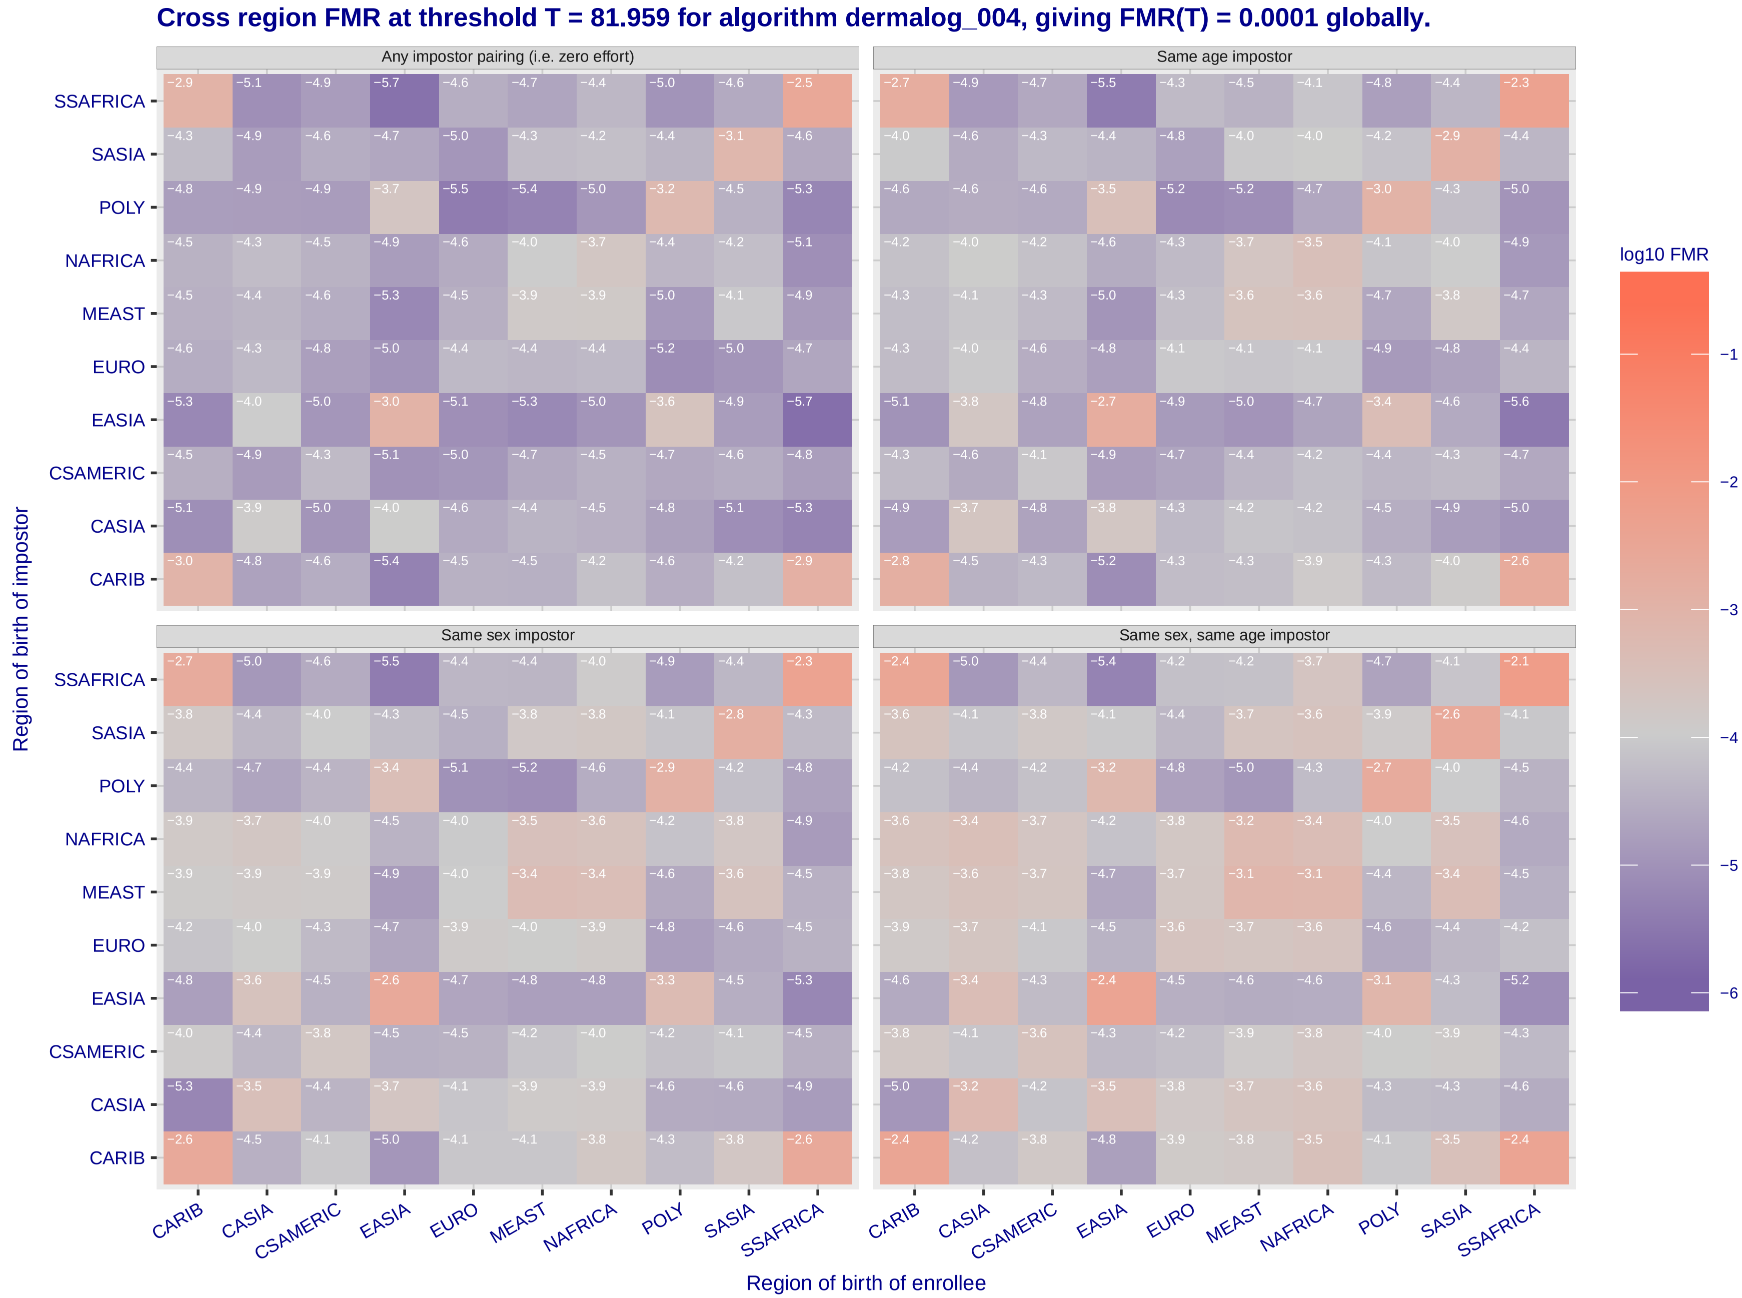

Figure 34: For algorithm dermalog-004 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

Figure 35: For algorithm dermalog-005 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

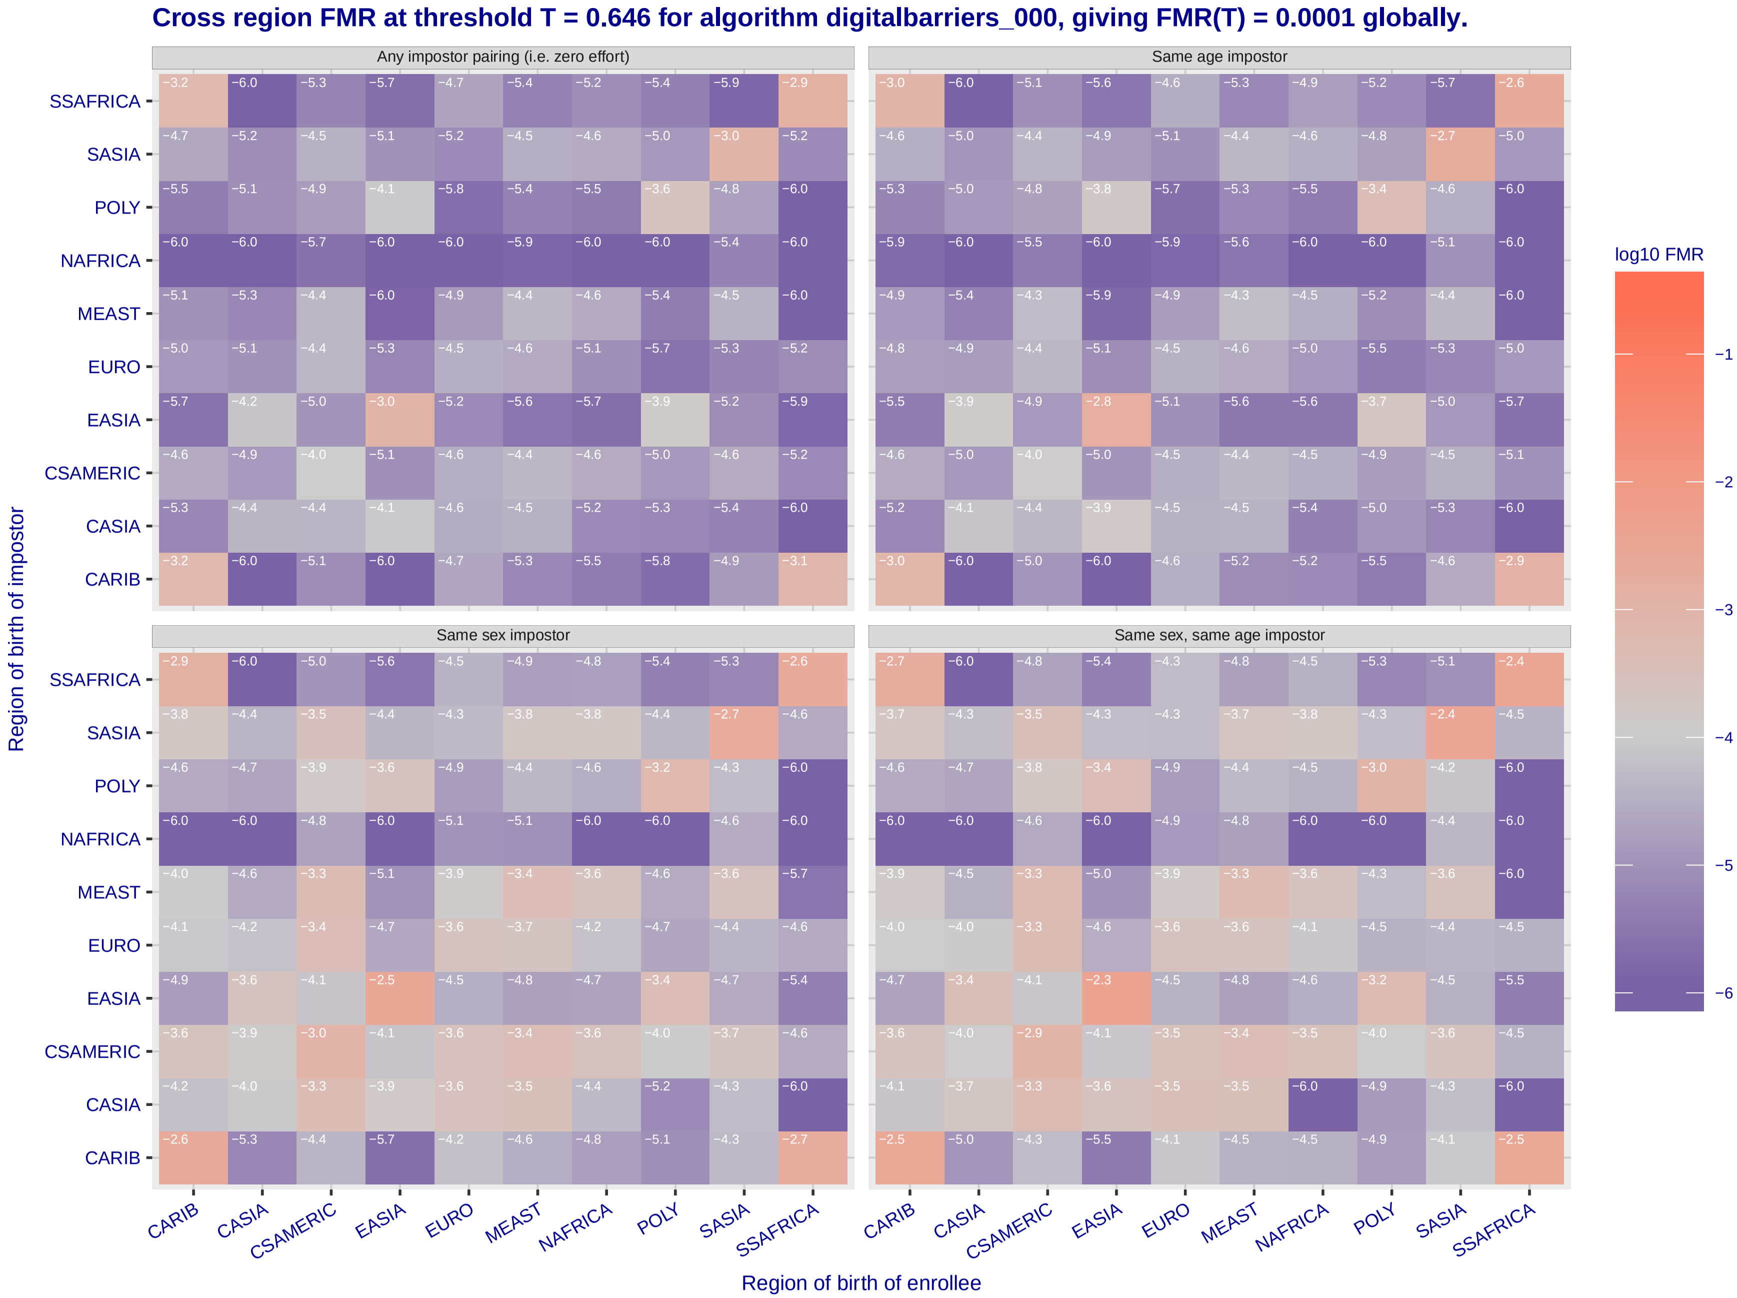

Figure 36: For algorithm digitalbarriers-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

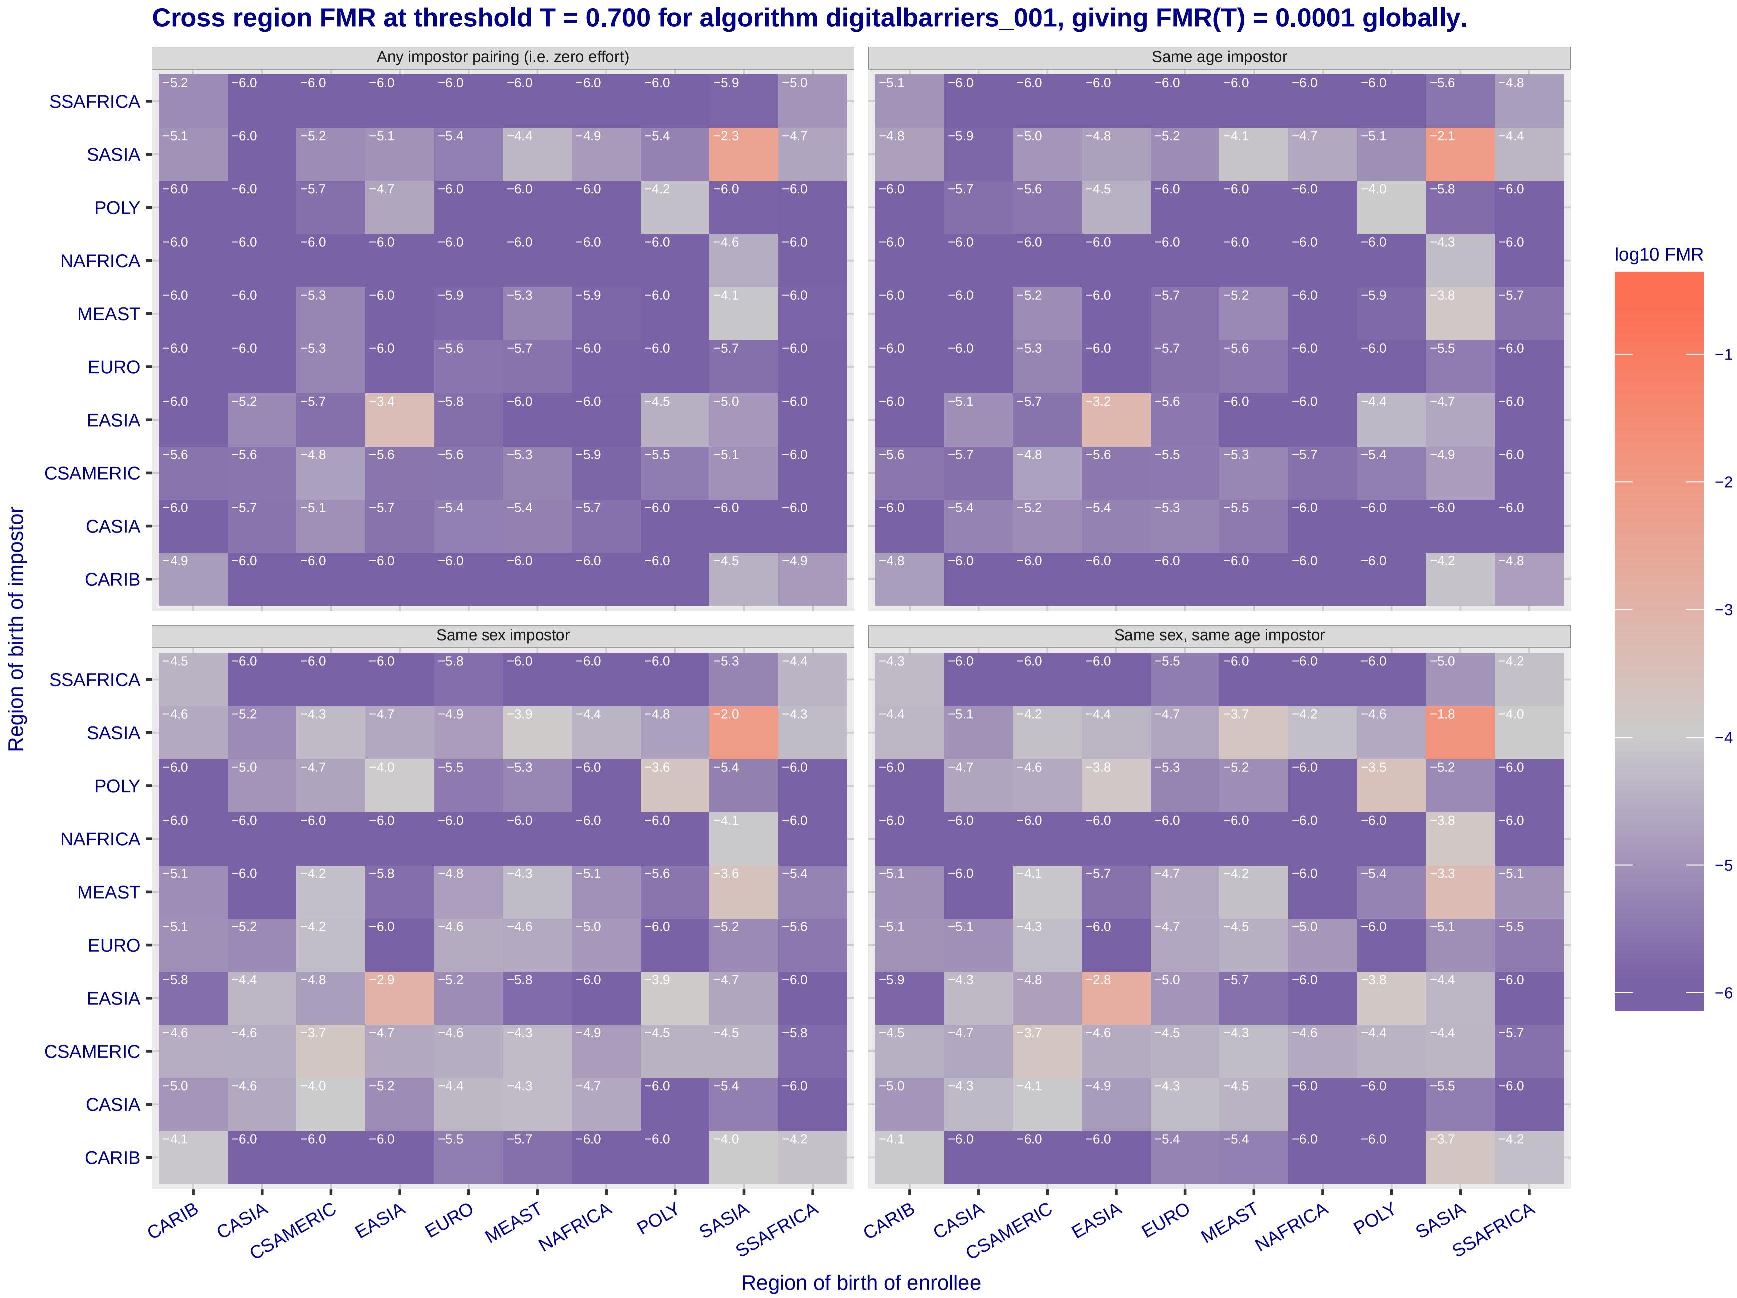

Figure 37: For algorithm digitalbarriers-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

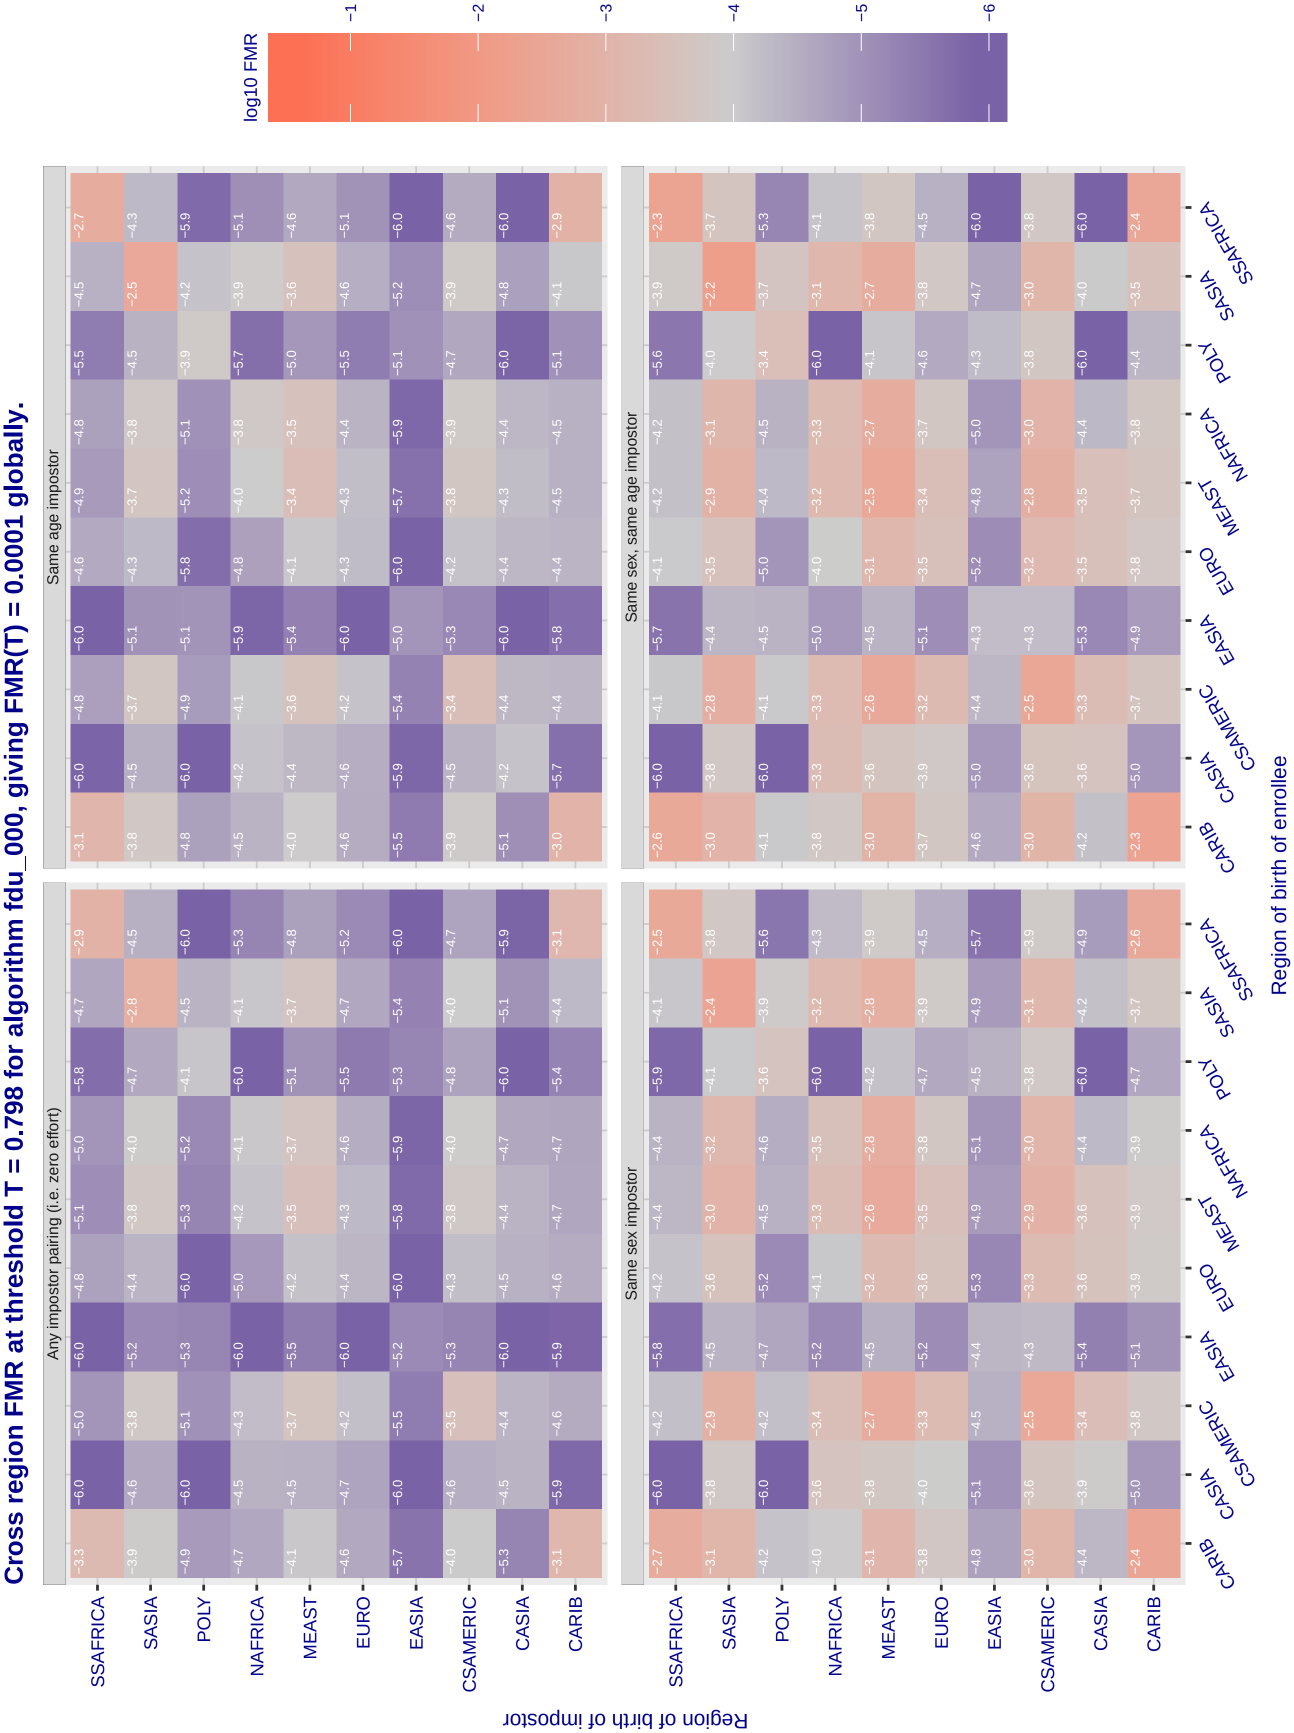

Figure 38: For algorithm fdu-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

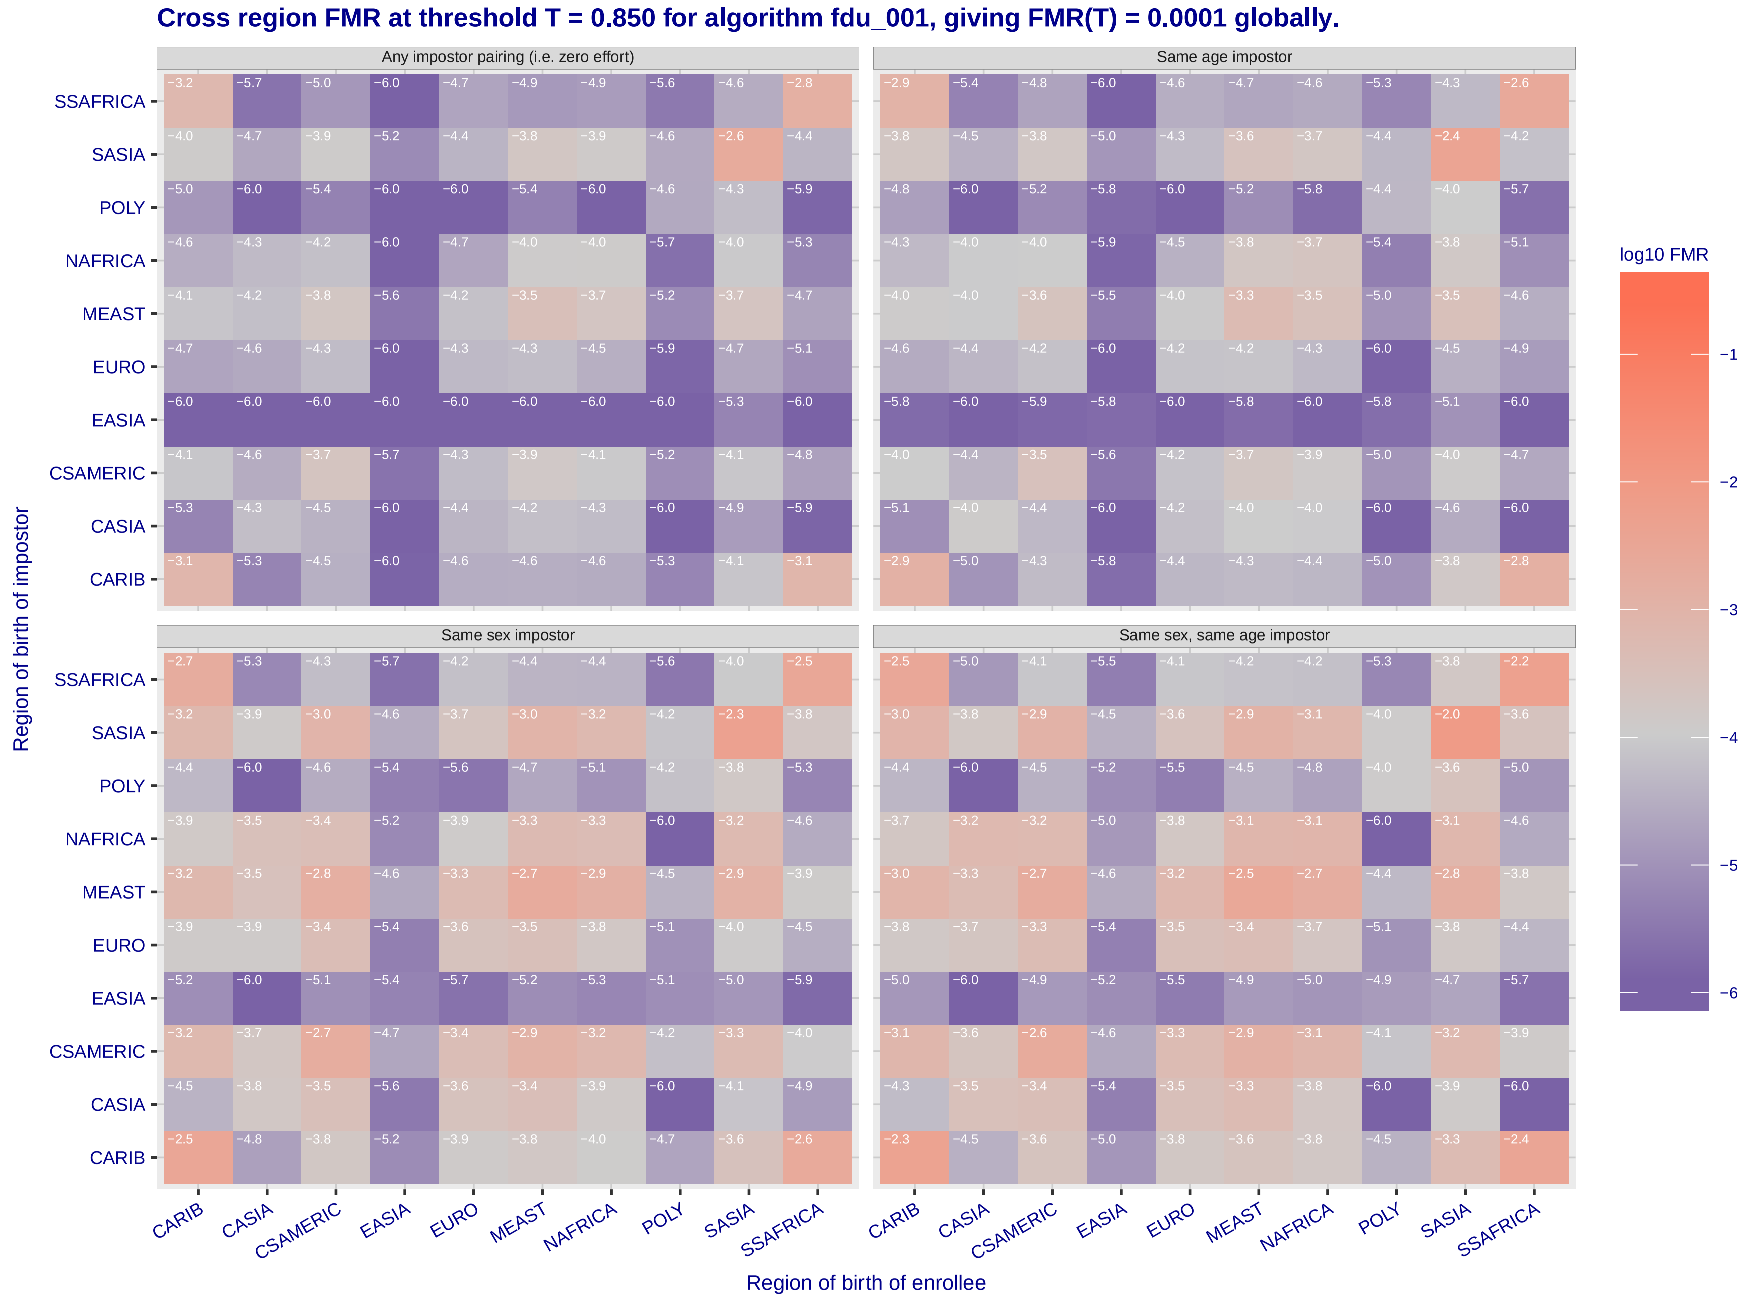

Figure 39: For algorithm fdu-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

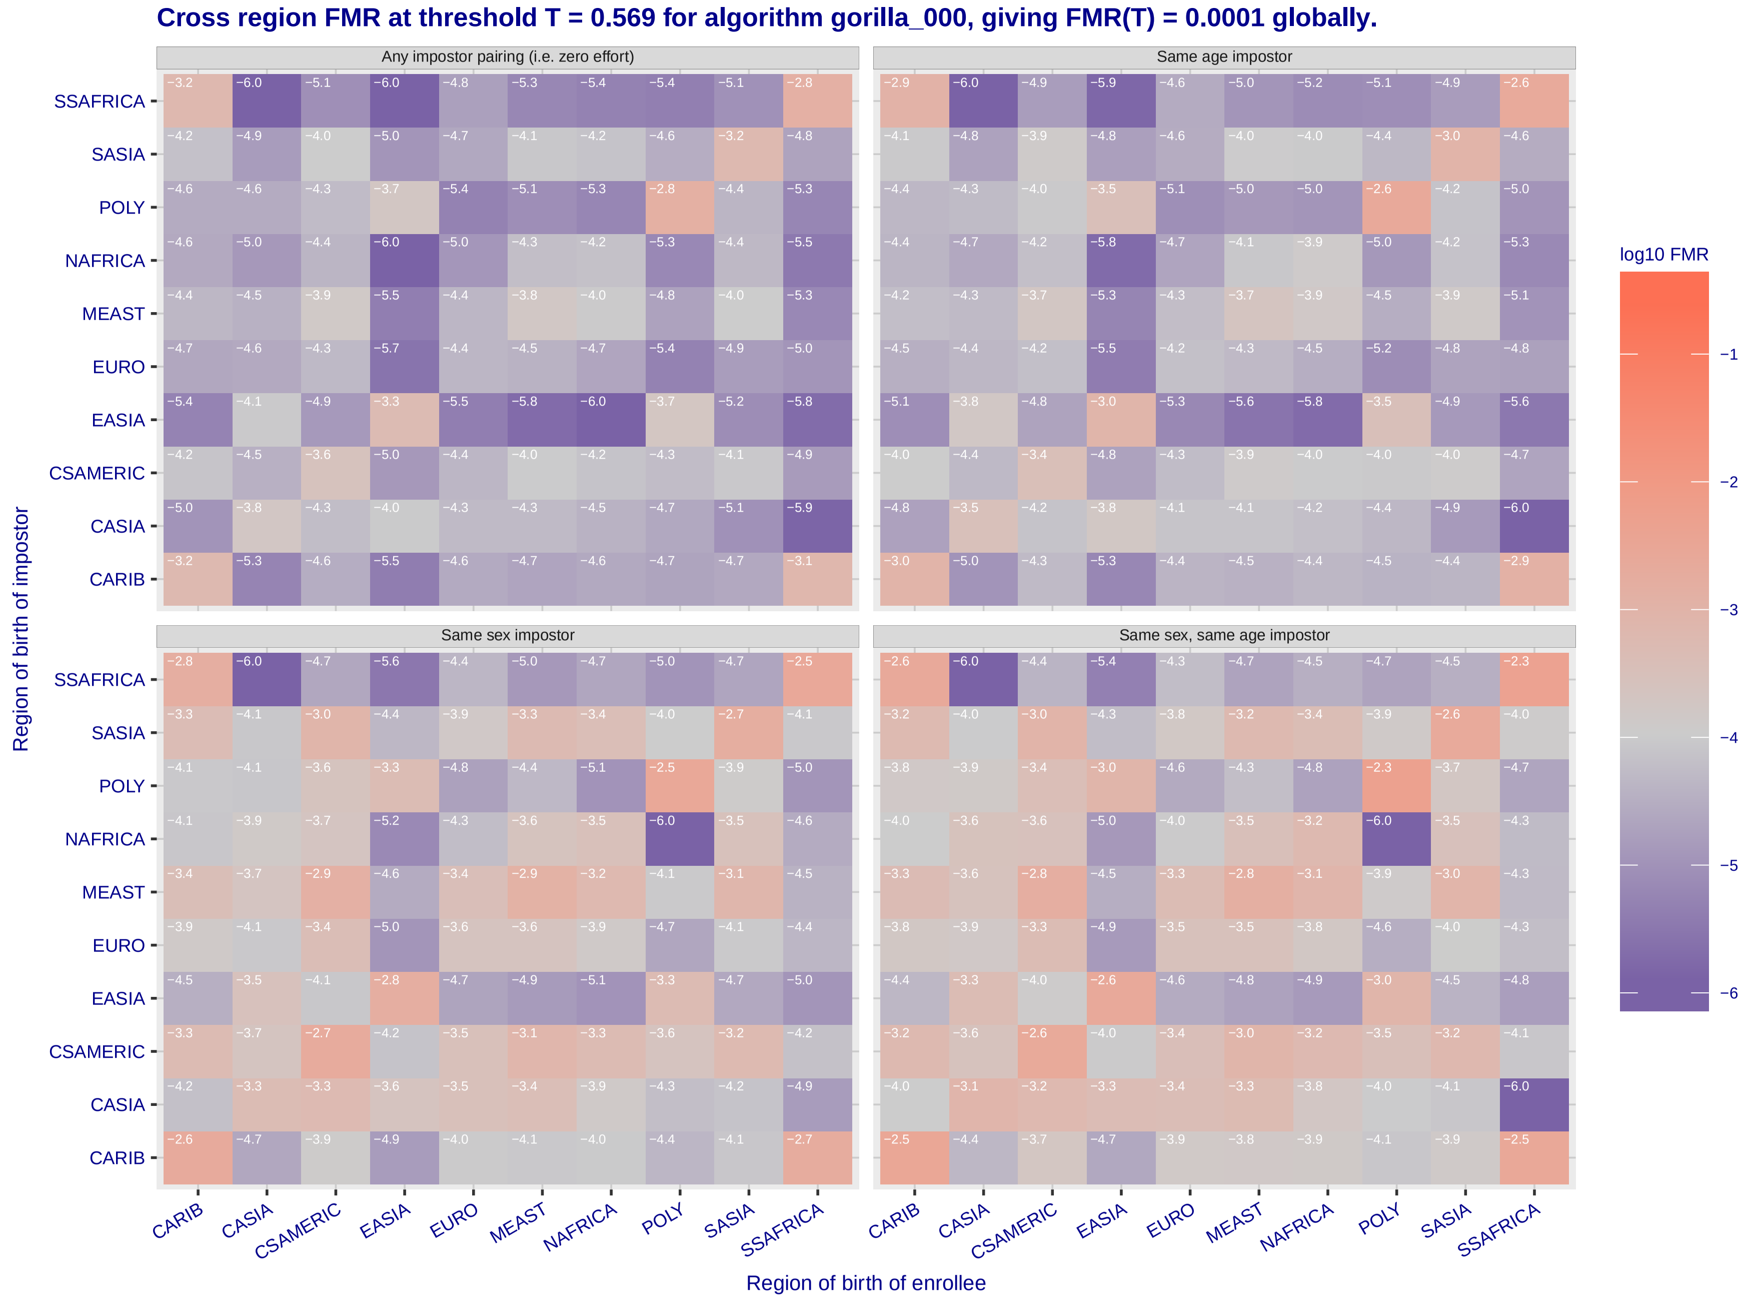

Figure 40: For algorithm gorilla-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

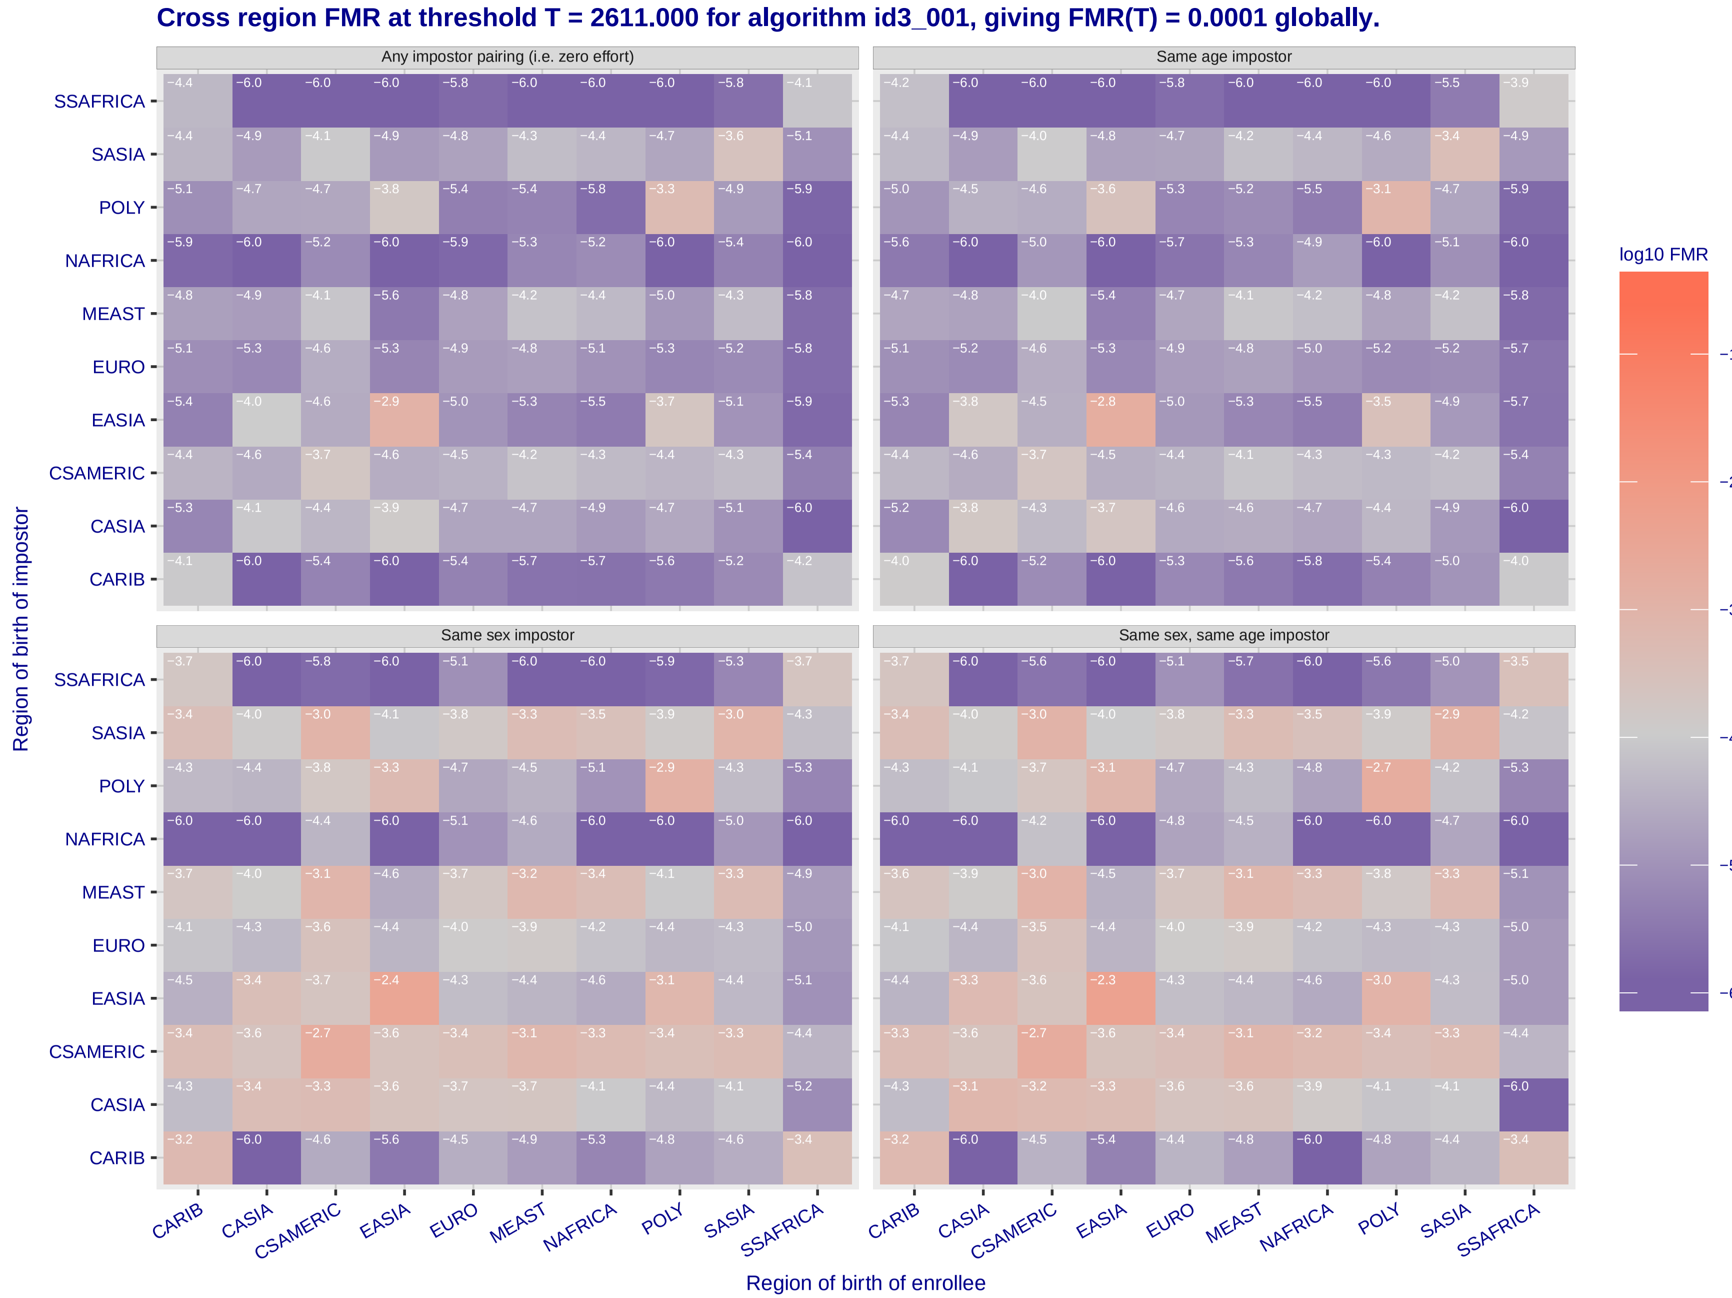

Figure 41: For algorithm id3-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

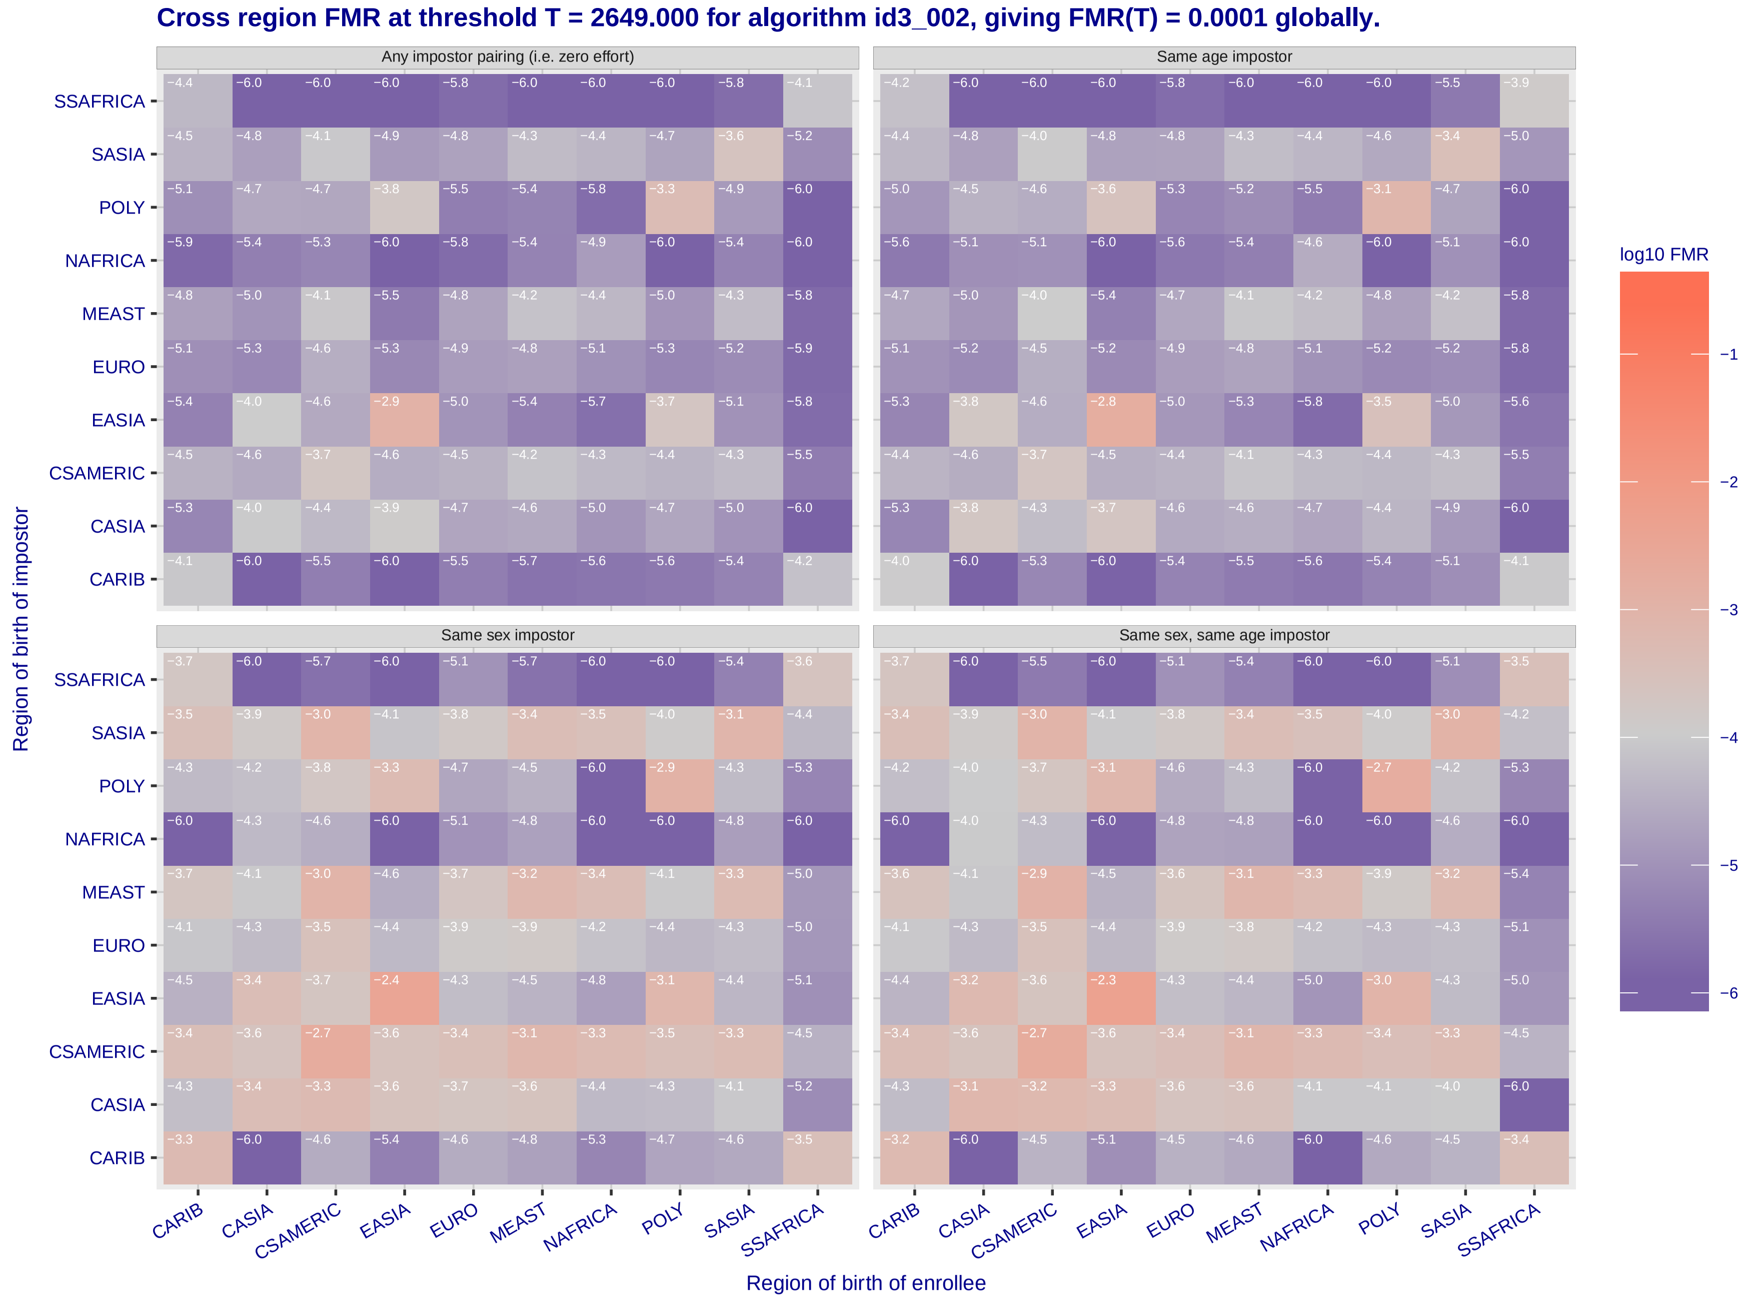

Figure 42: For algorithm id3-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

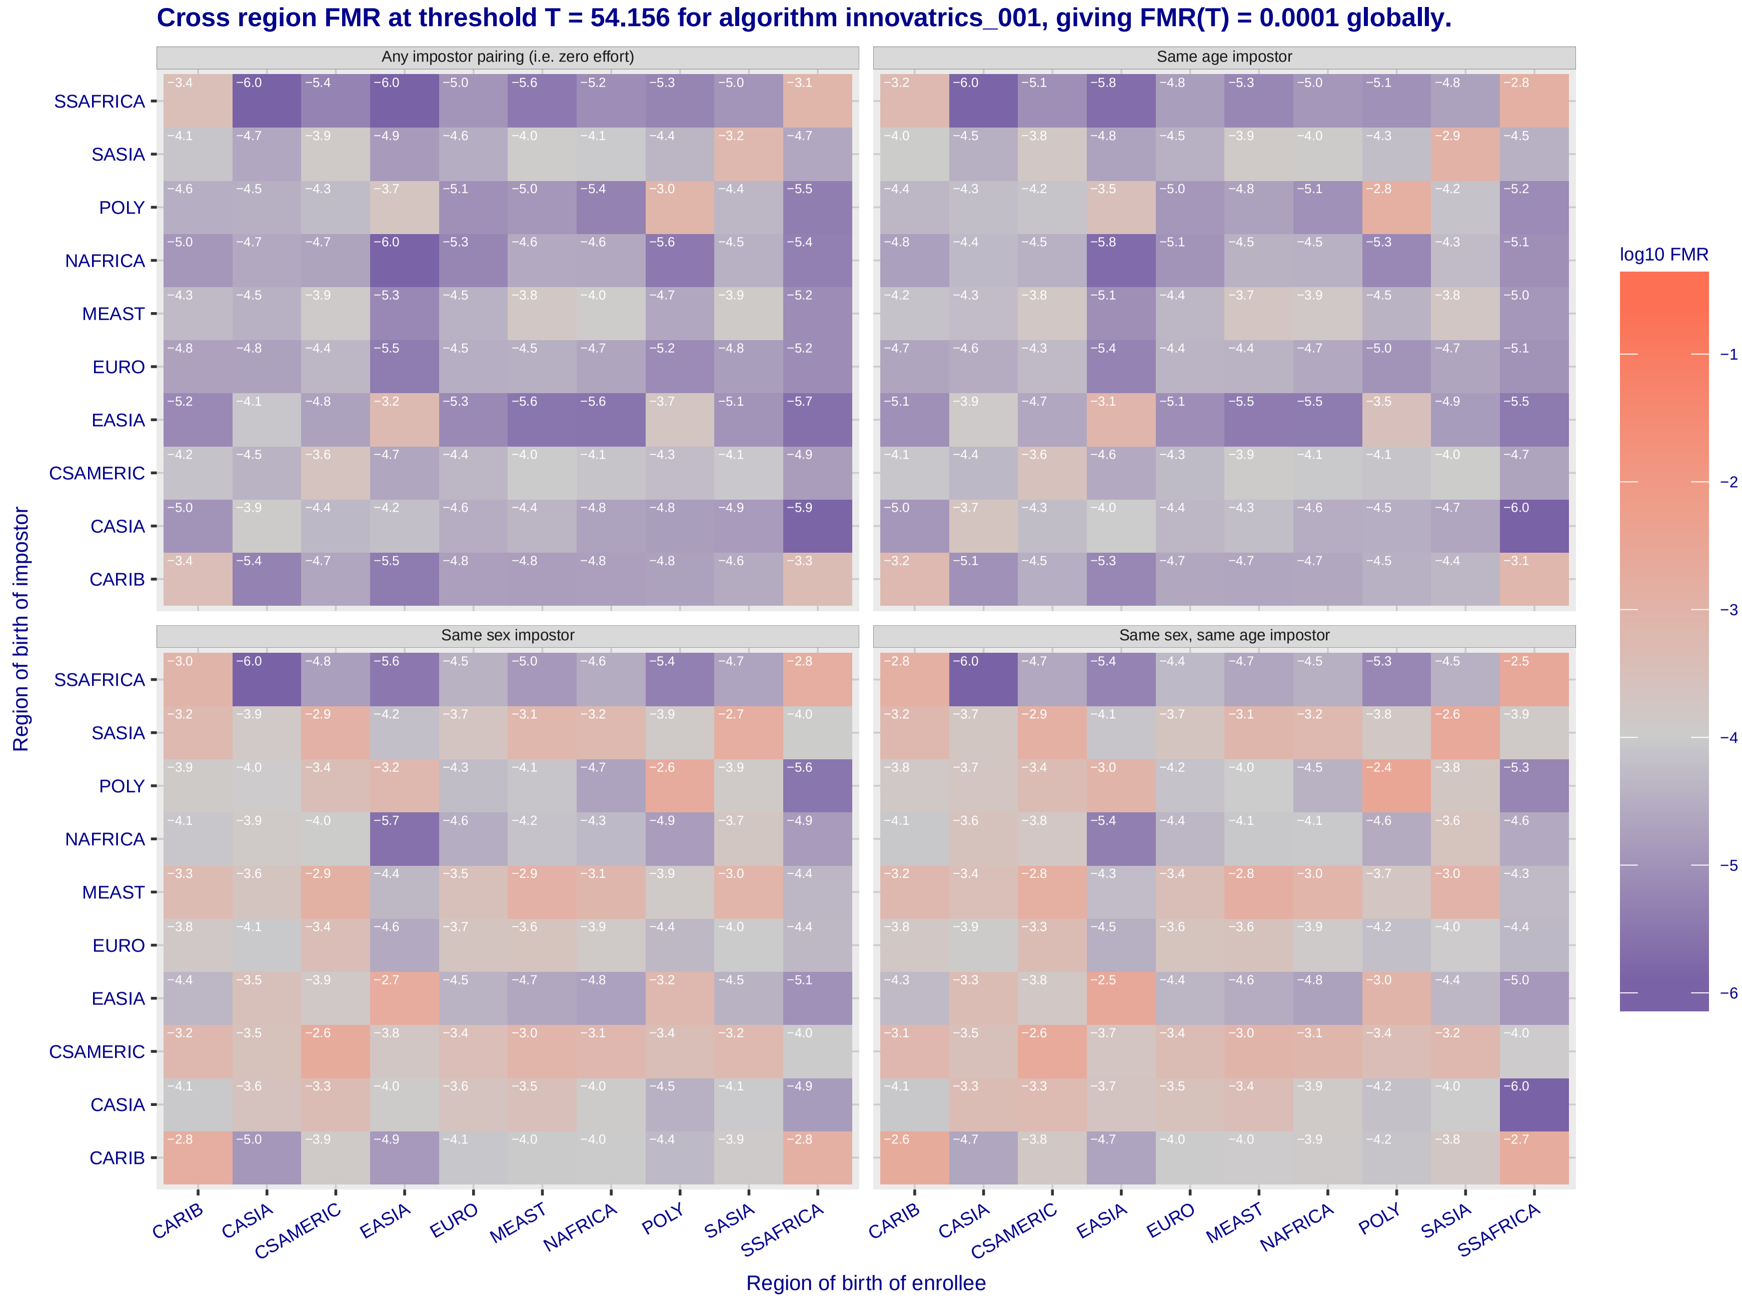

Figure 43: For algorithm innovatrics-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

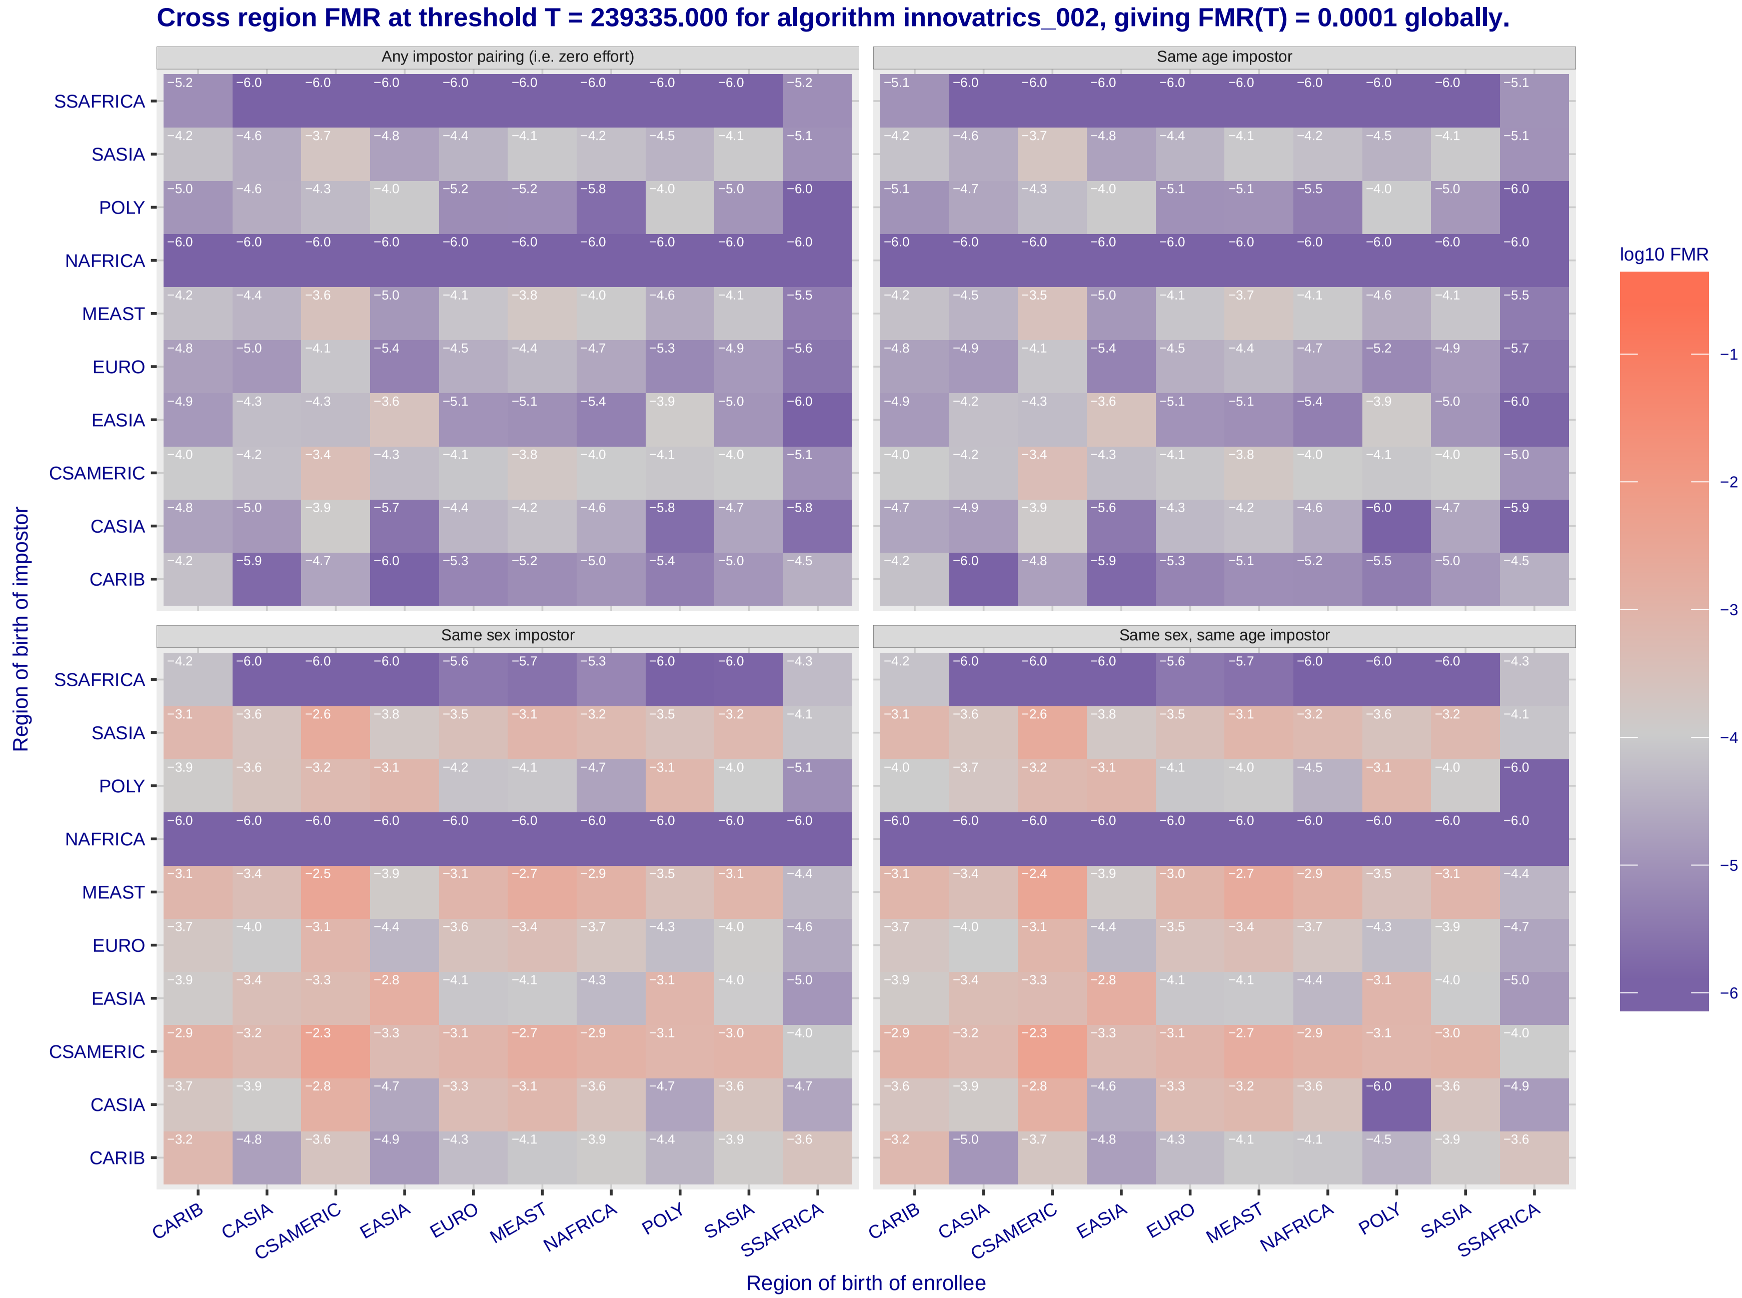

Figure 44: For algorithm innovatrics-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

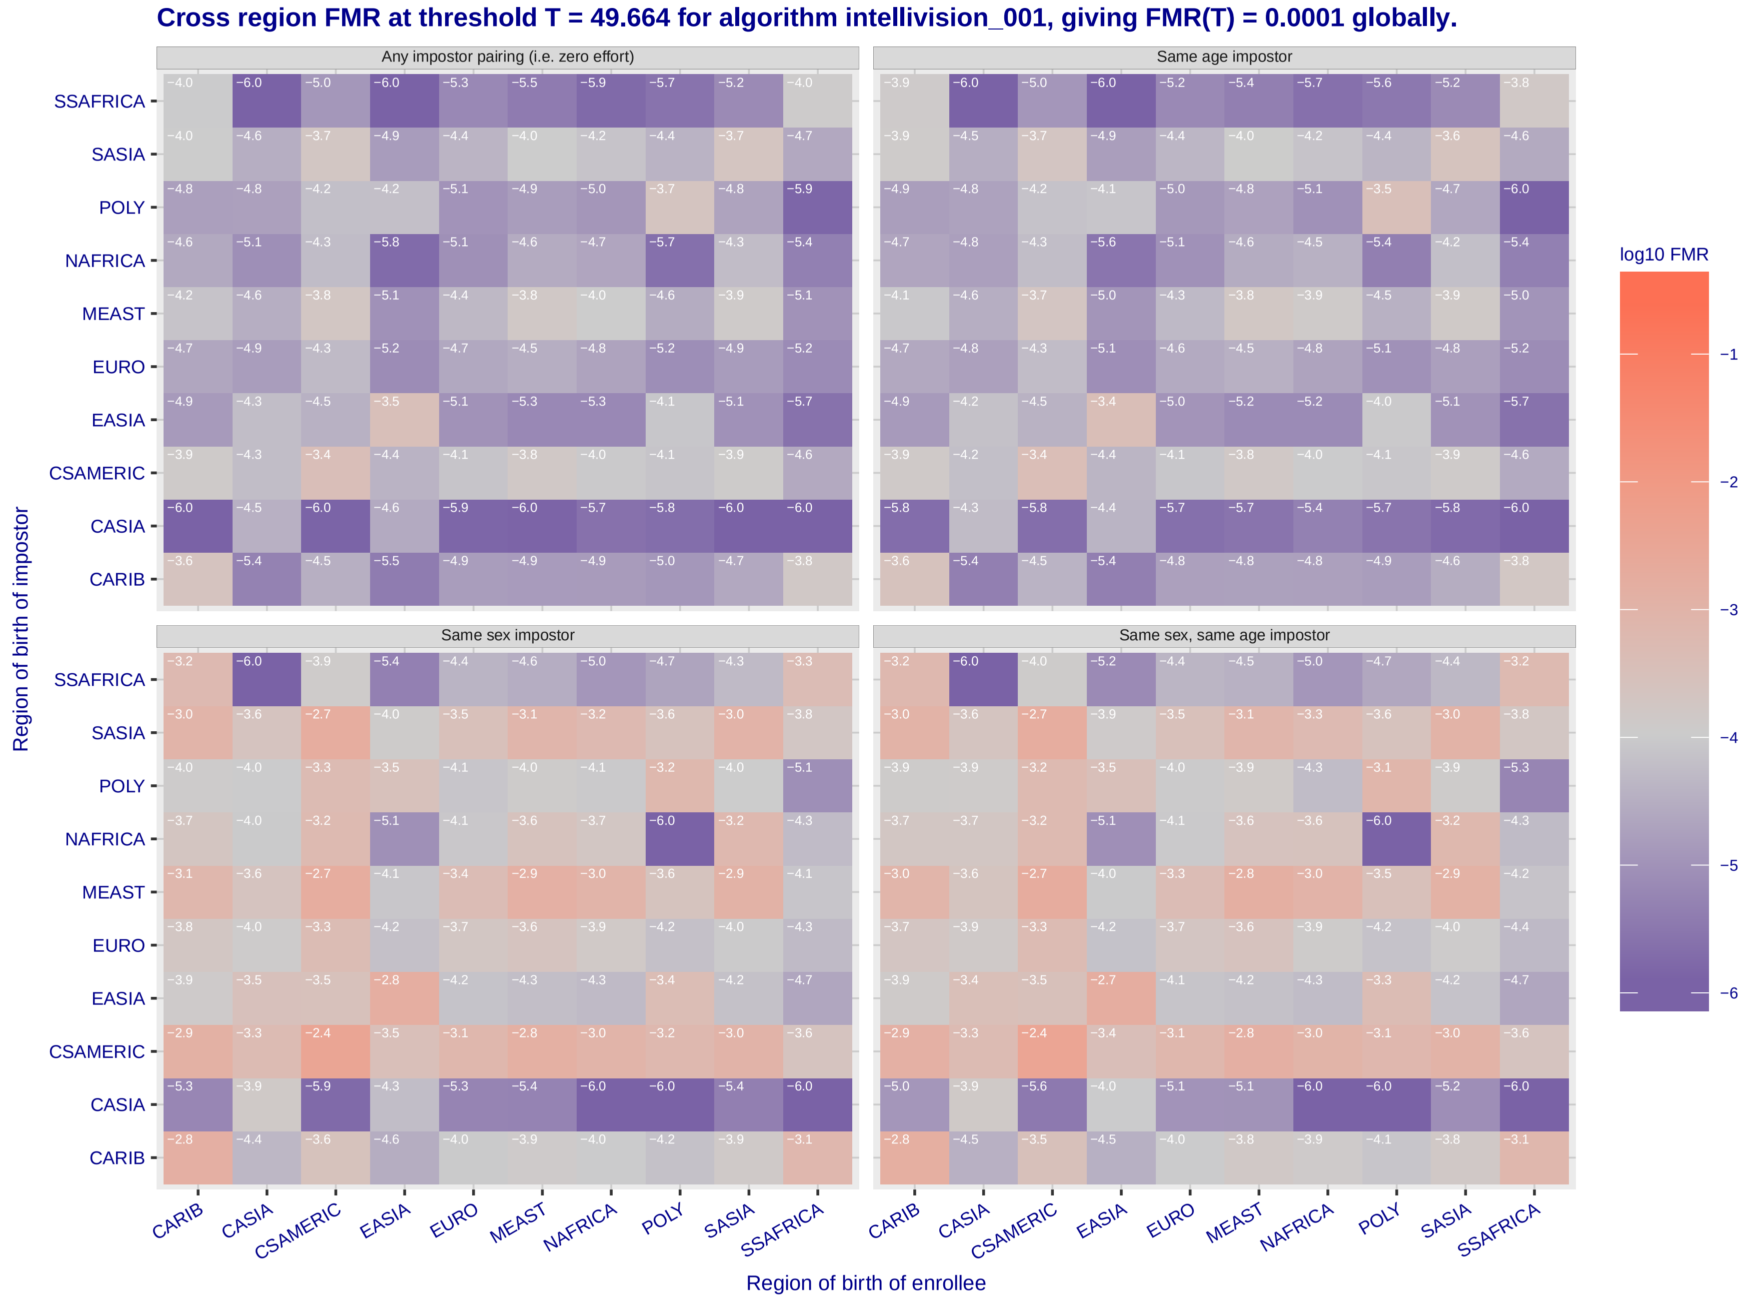

Figure 45: For algorithm intellivision-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

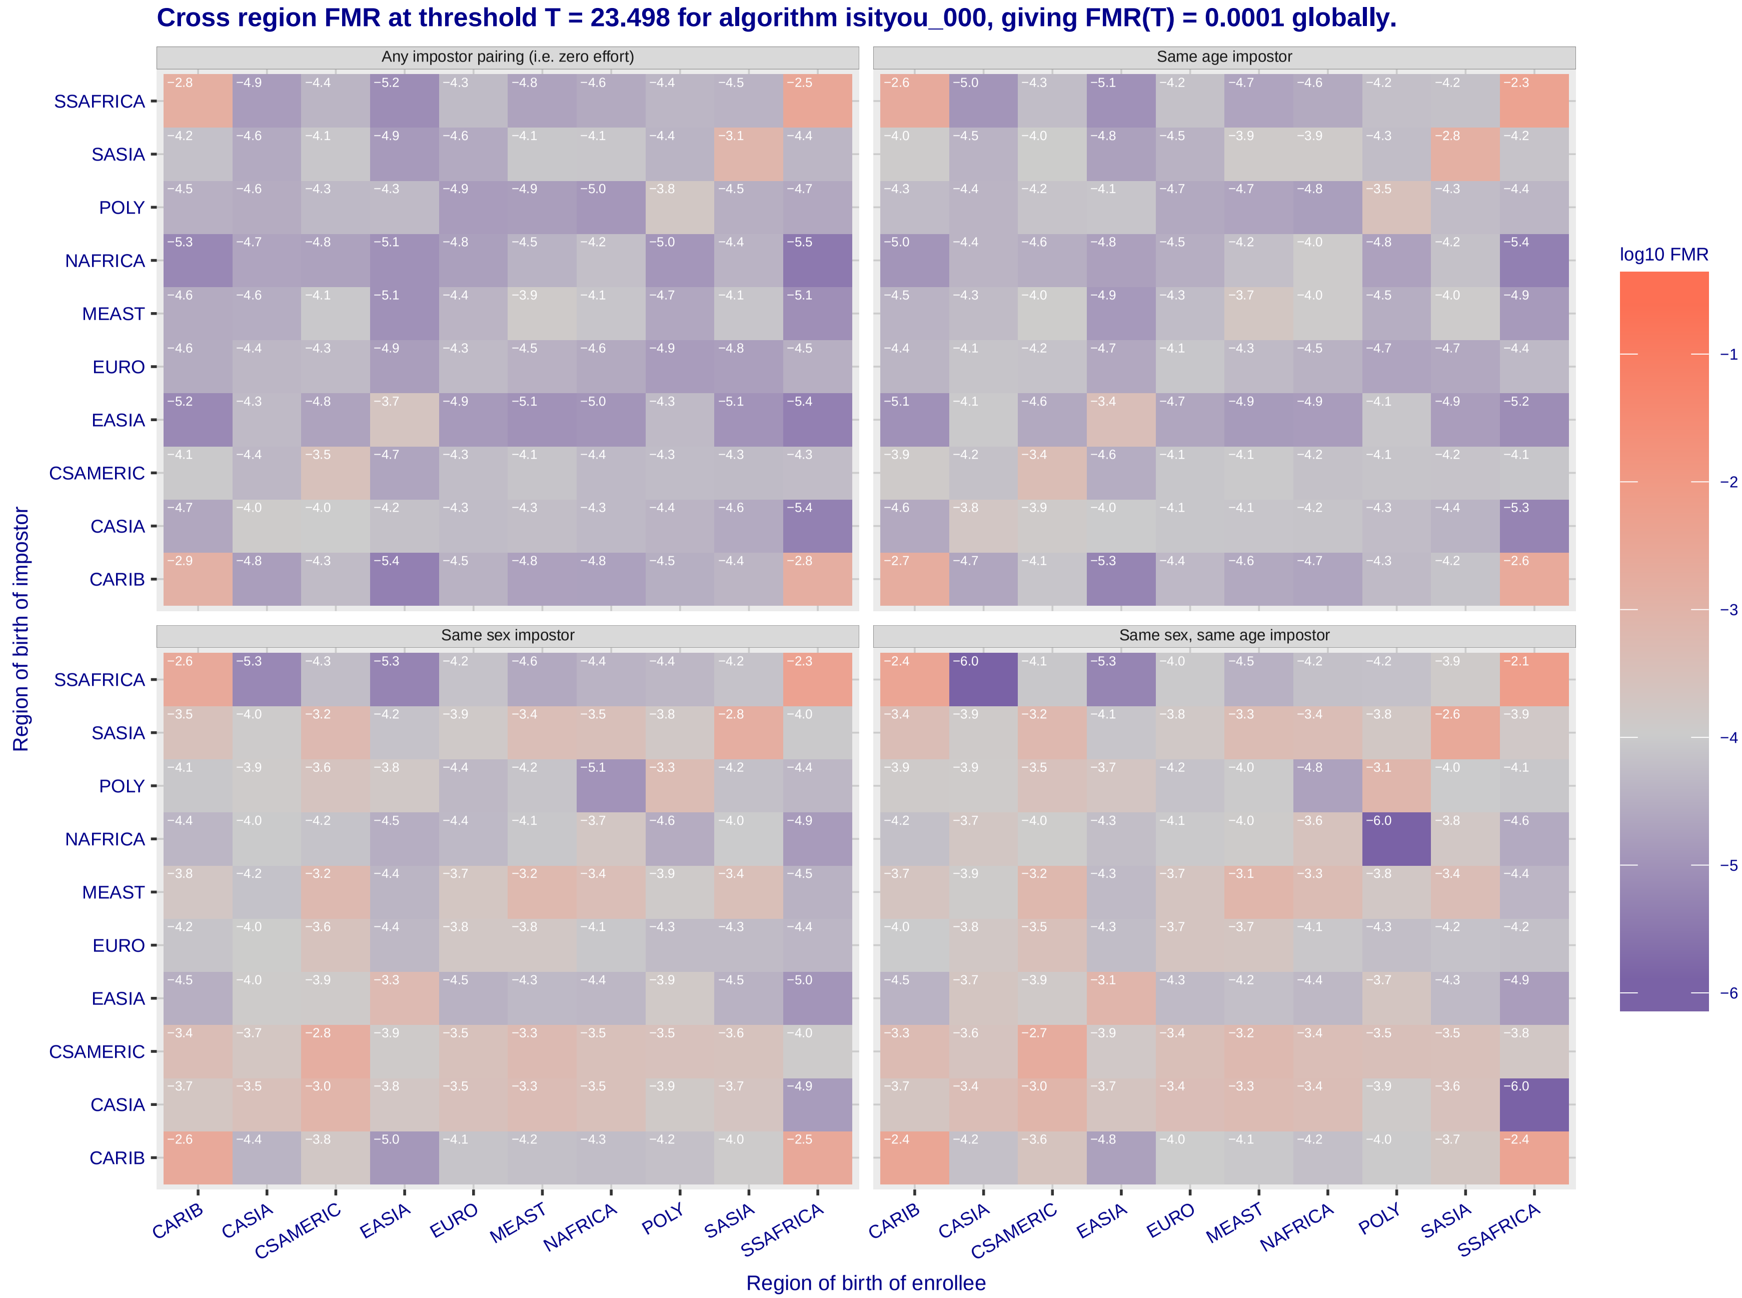

Figure 46: For algorithm isityou-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

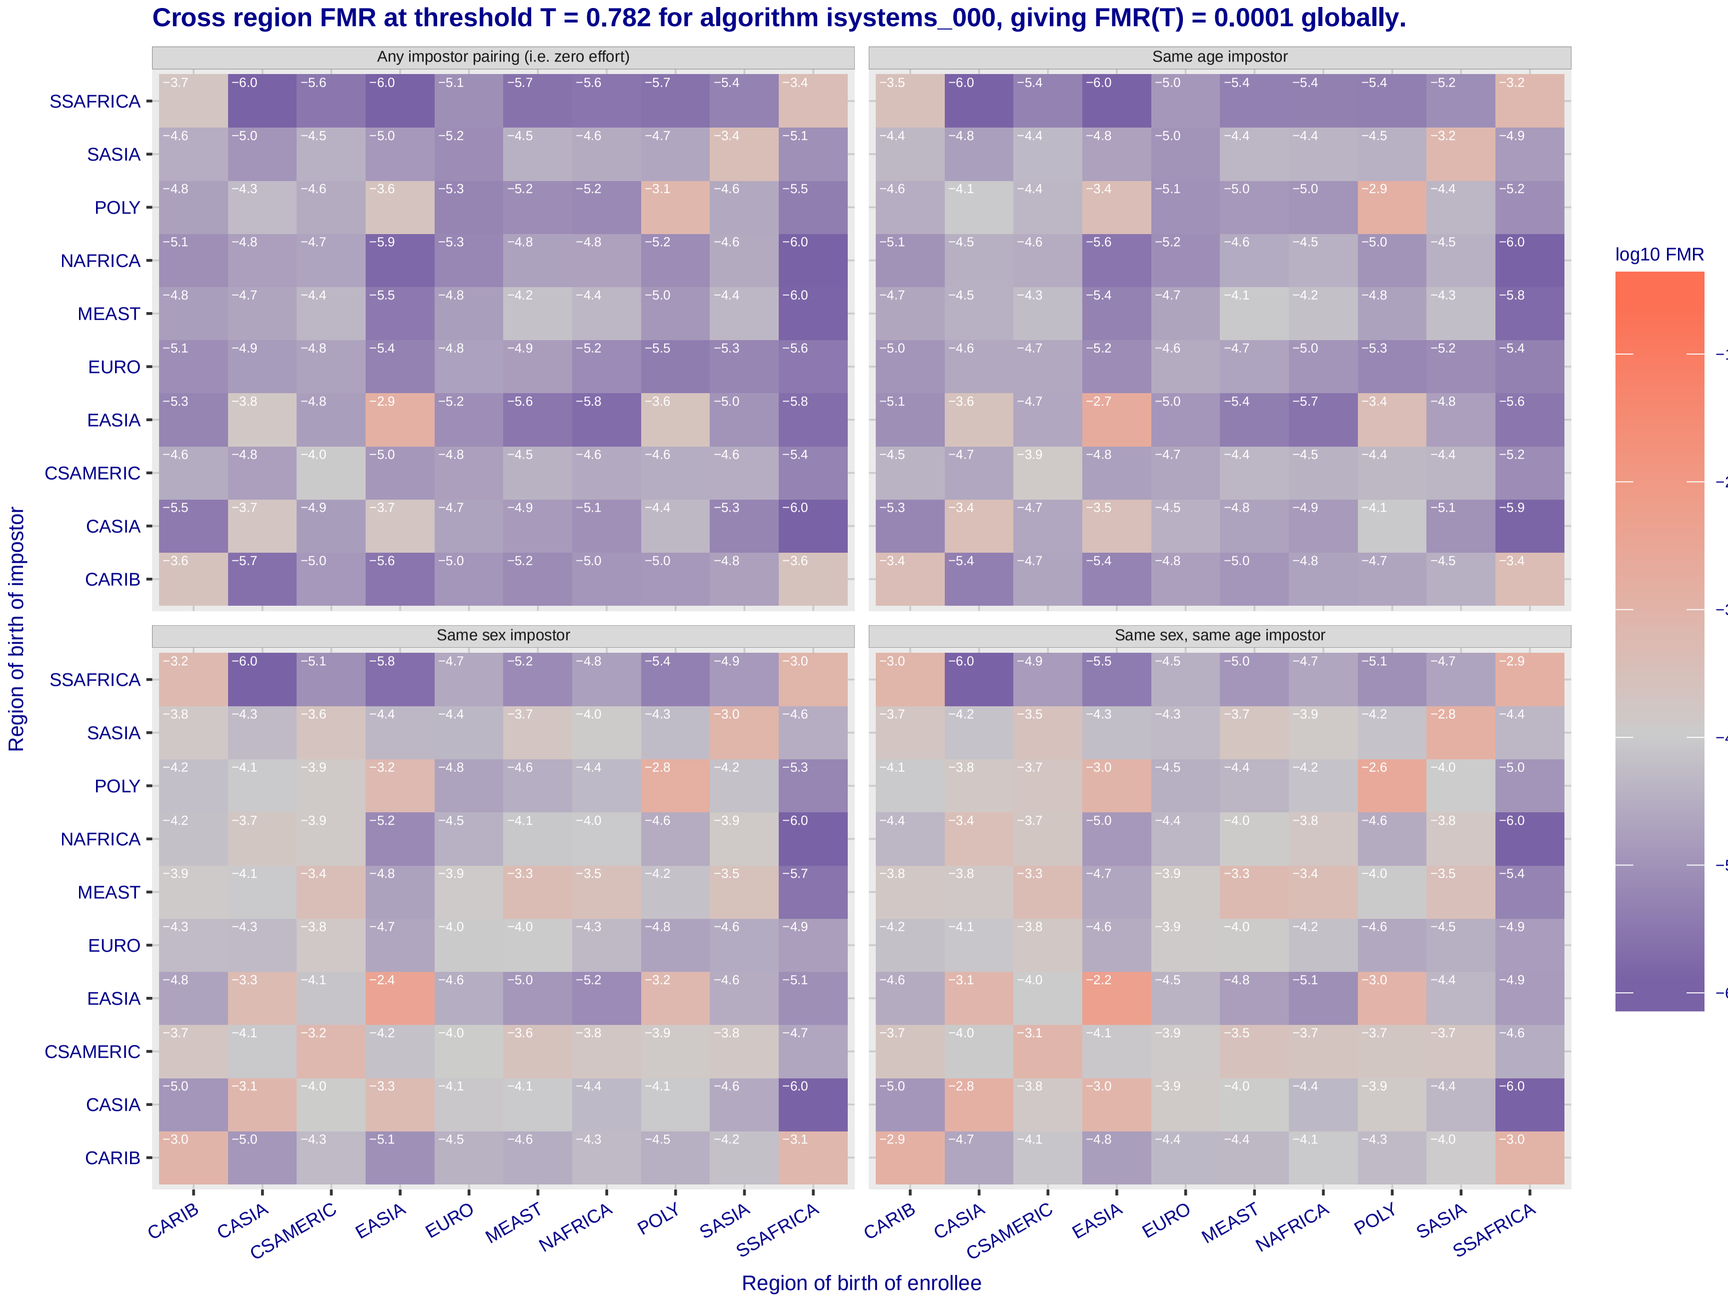

Figure 47: For algorithm isystems-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

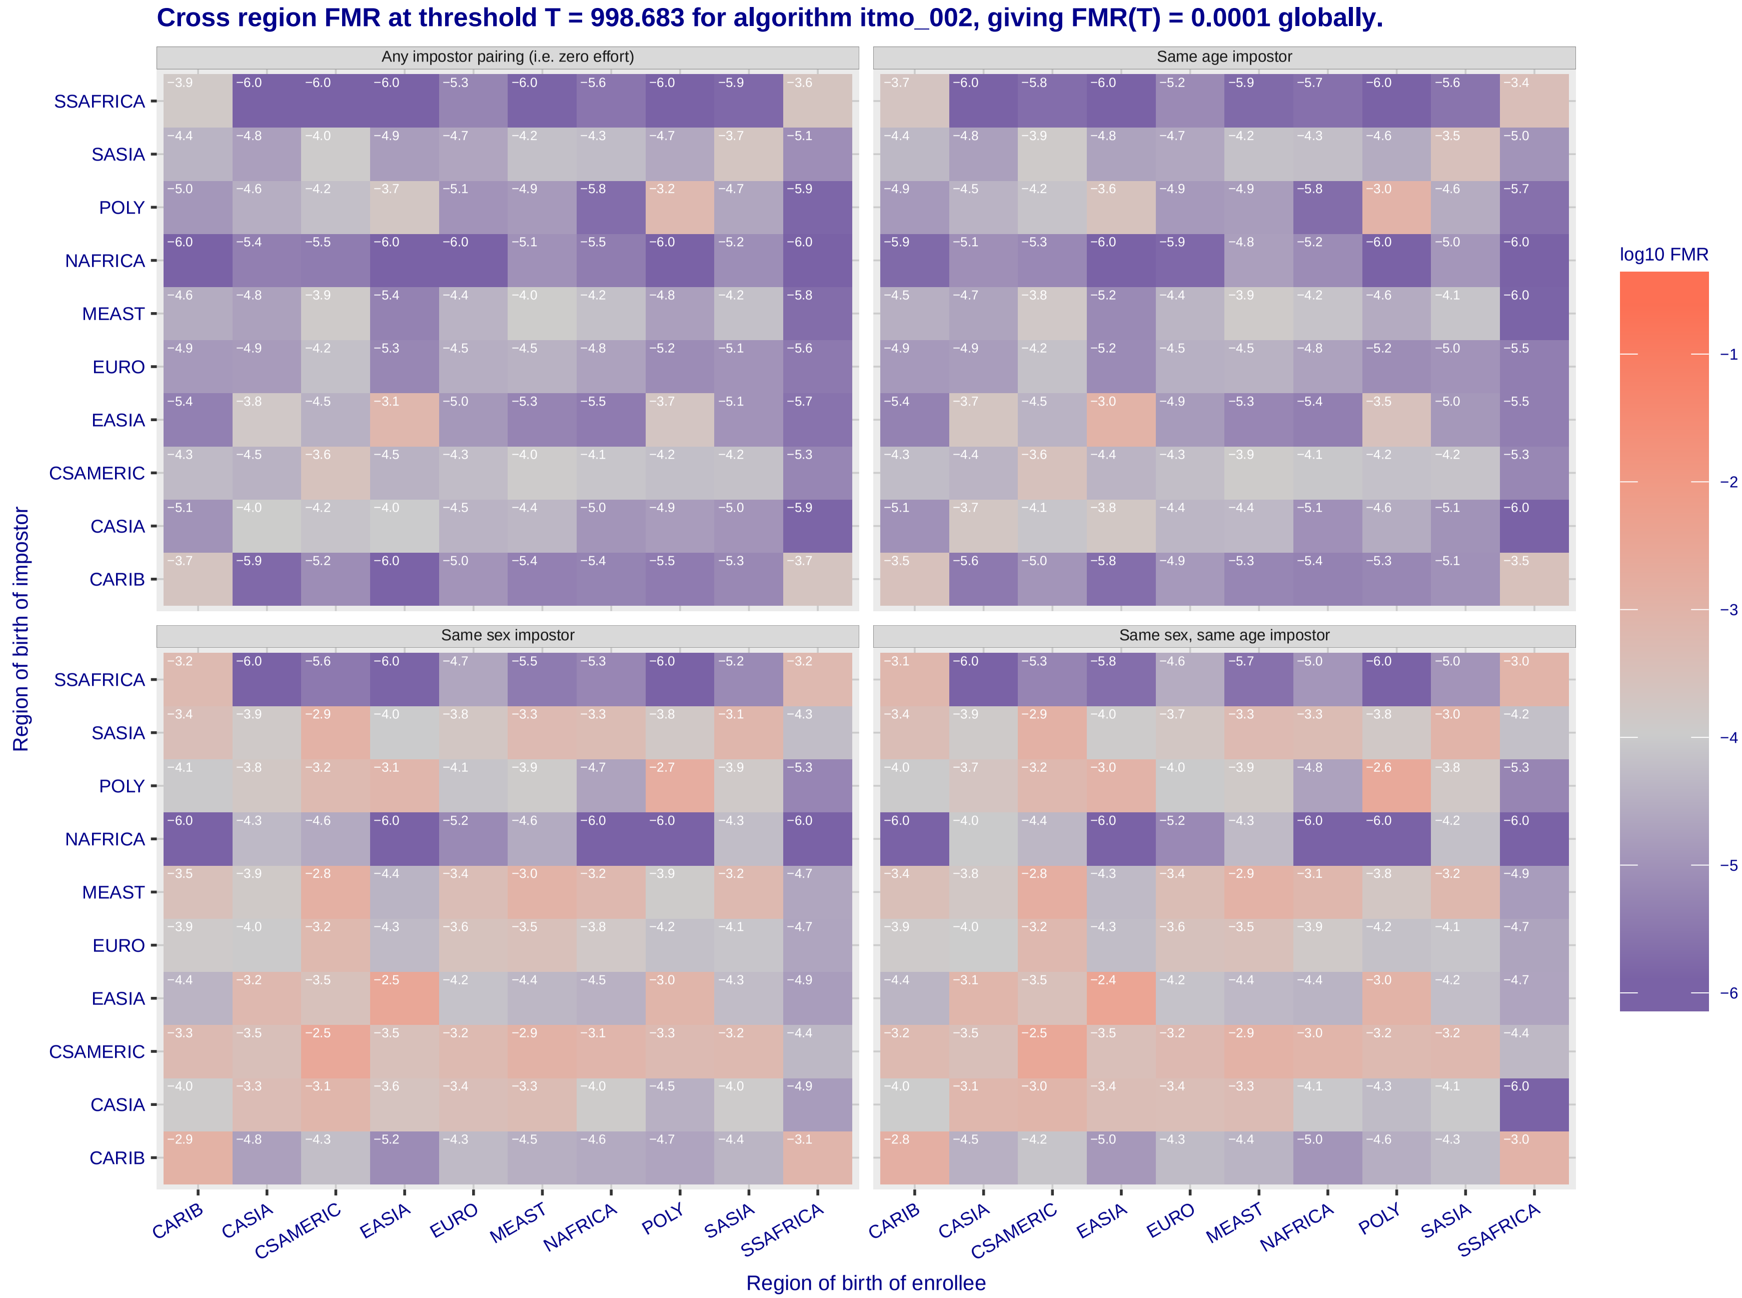

Figure 48: For algorithm itmo-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

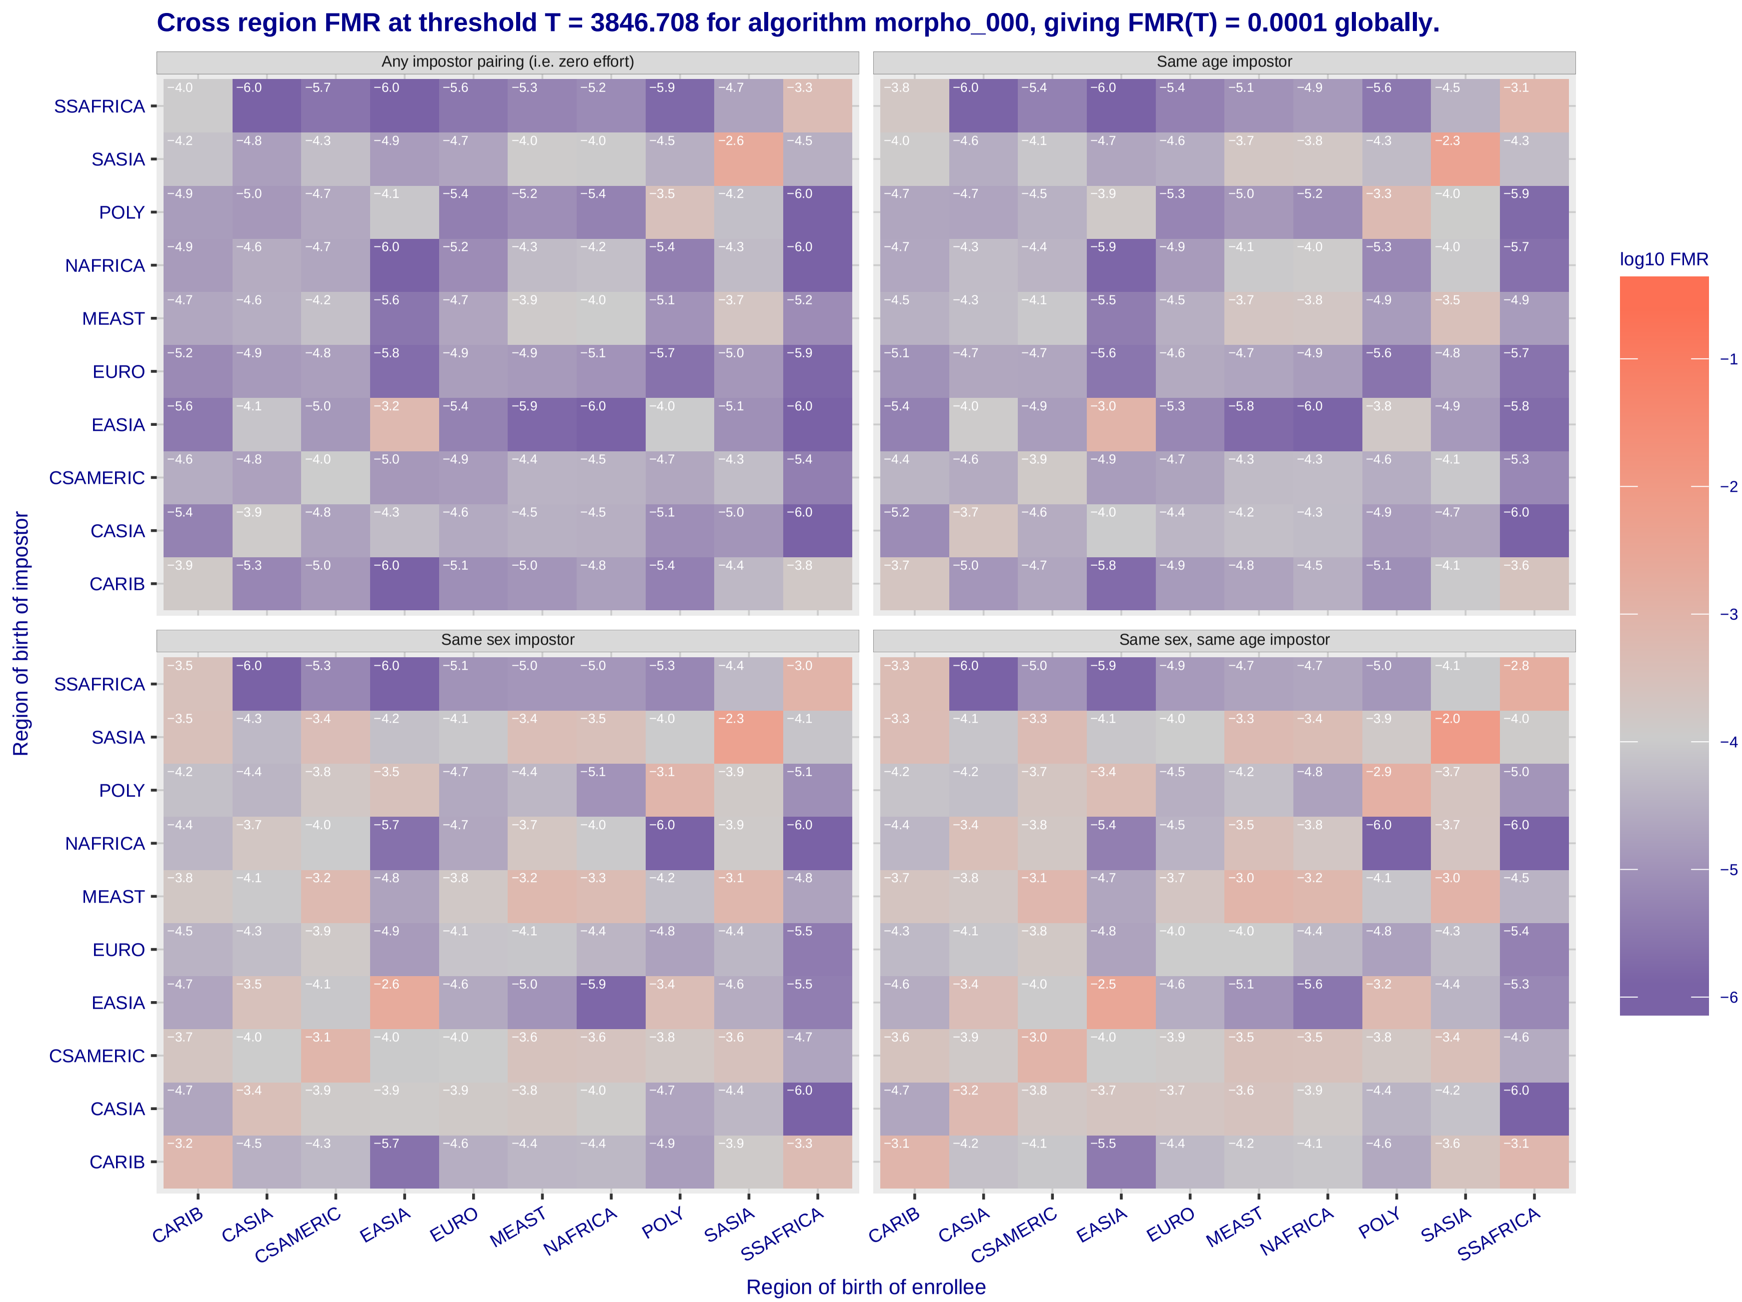

Figure 49: For algorithm morpho-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

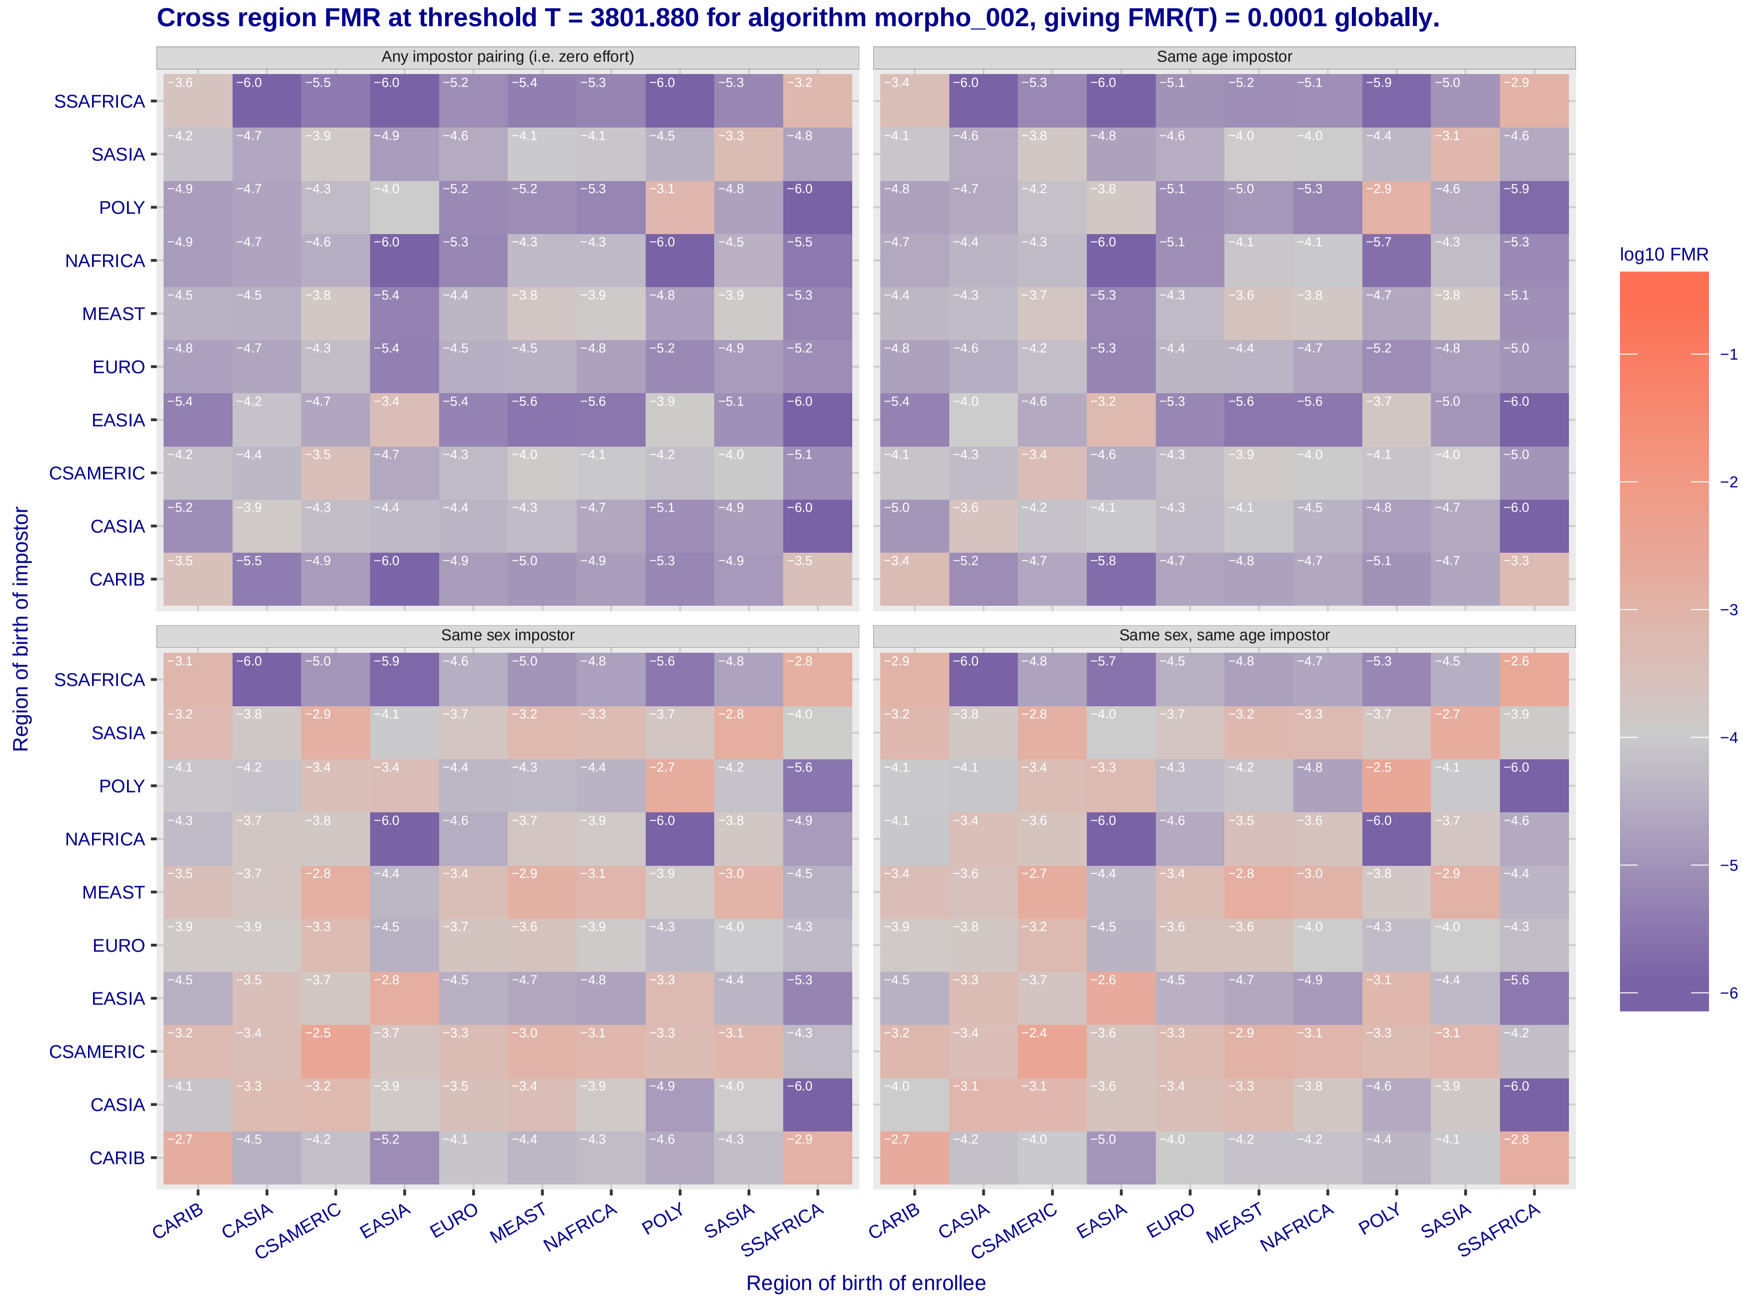

Figure 50: For algorithm morpho-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

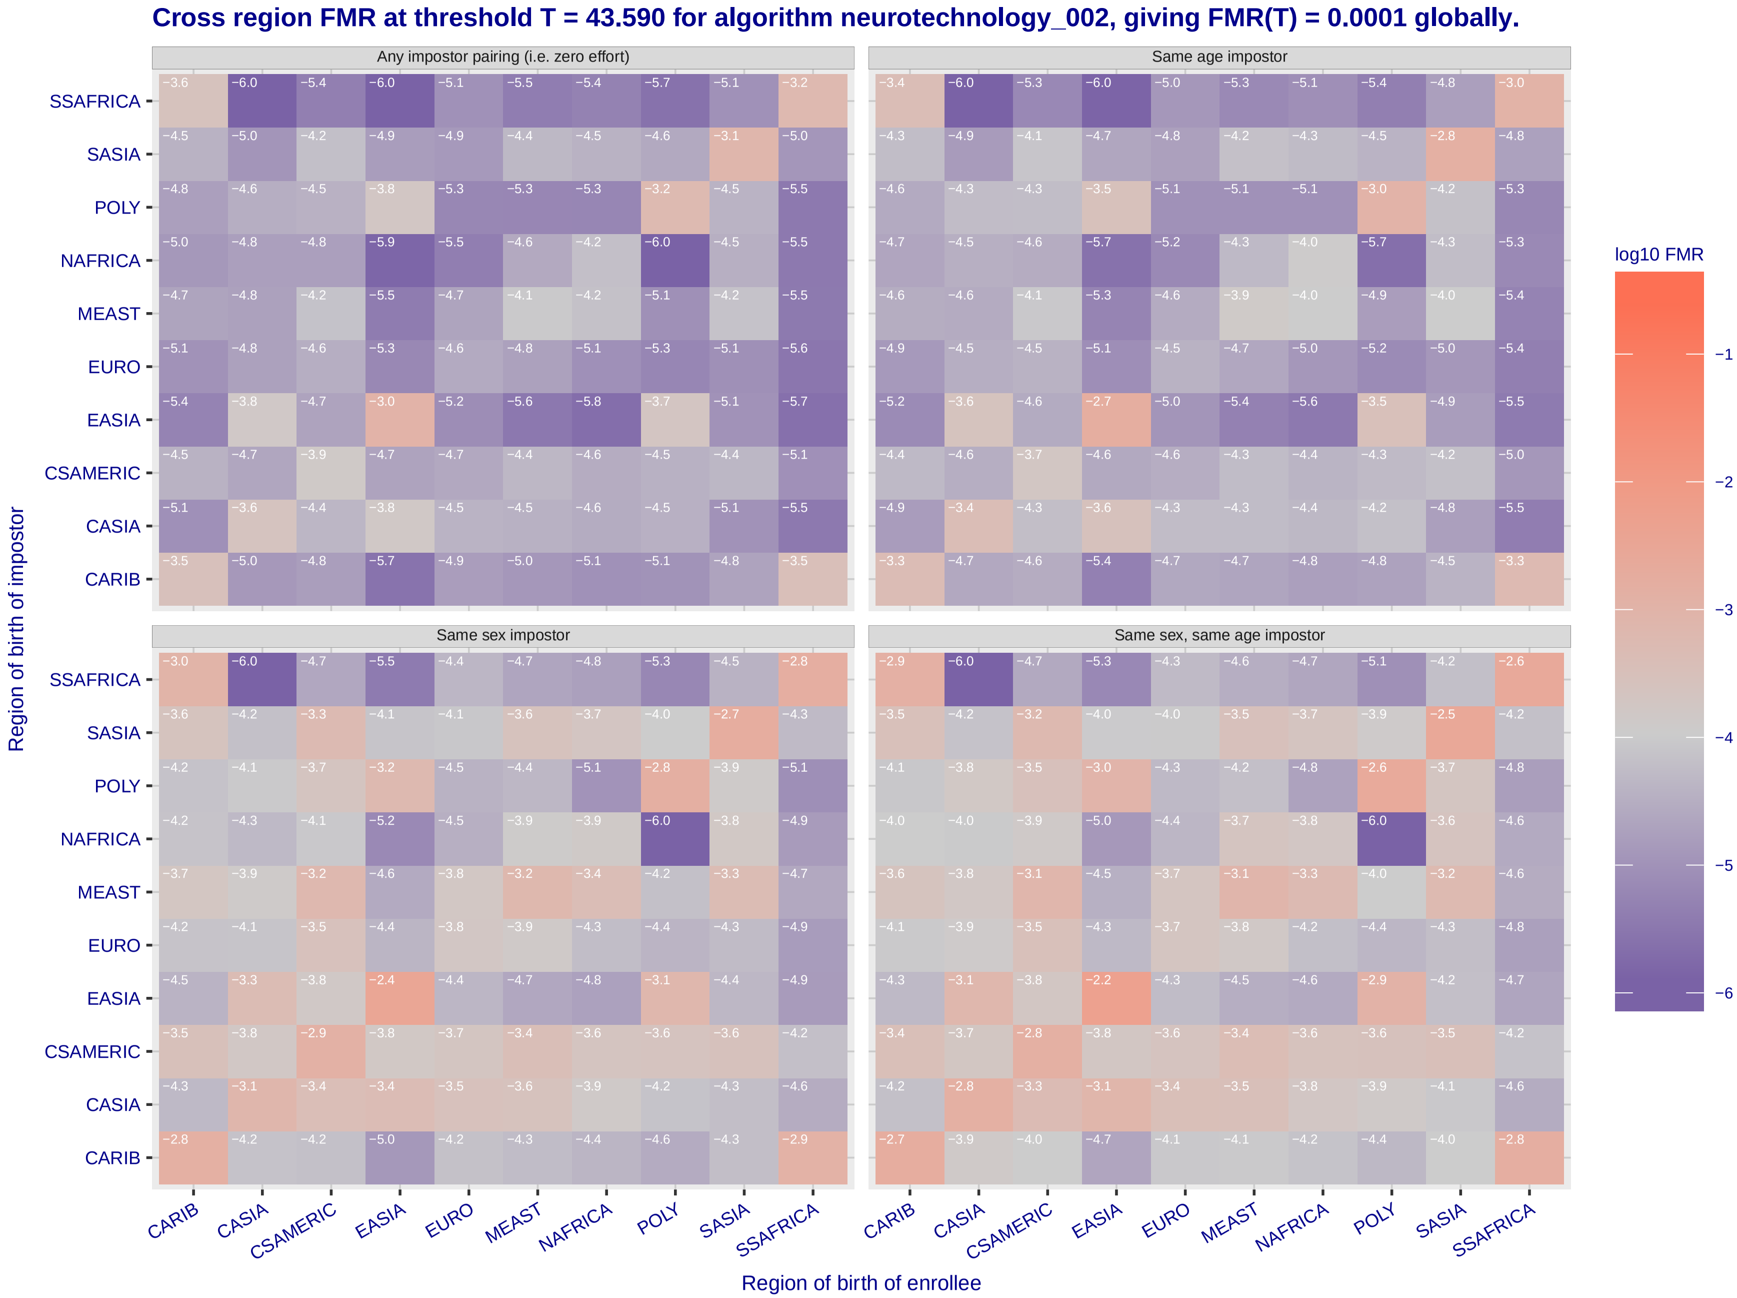

Figure 51: For algorithm neurotechnology-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

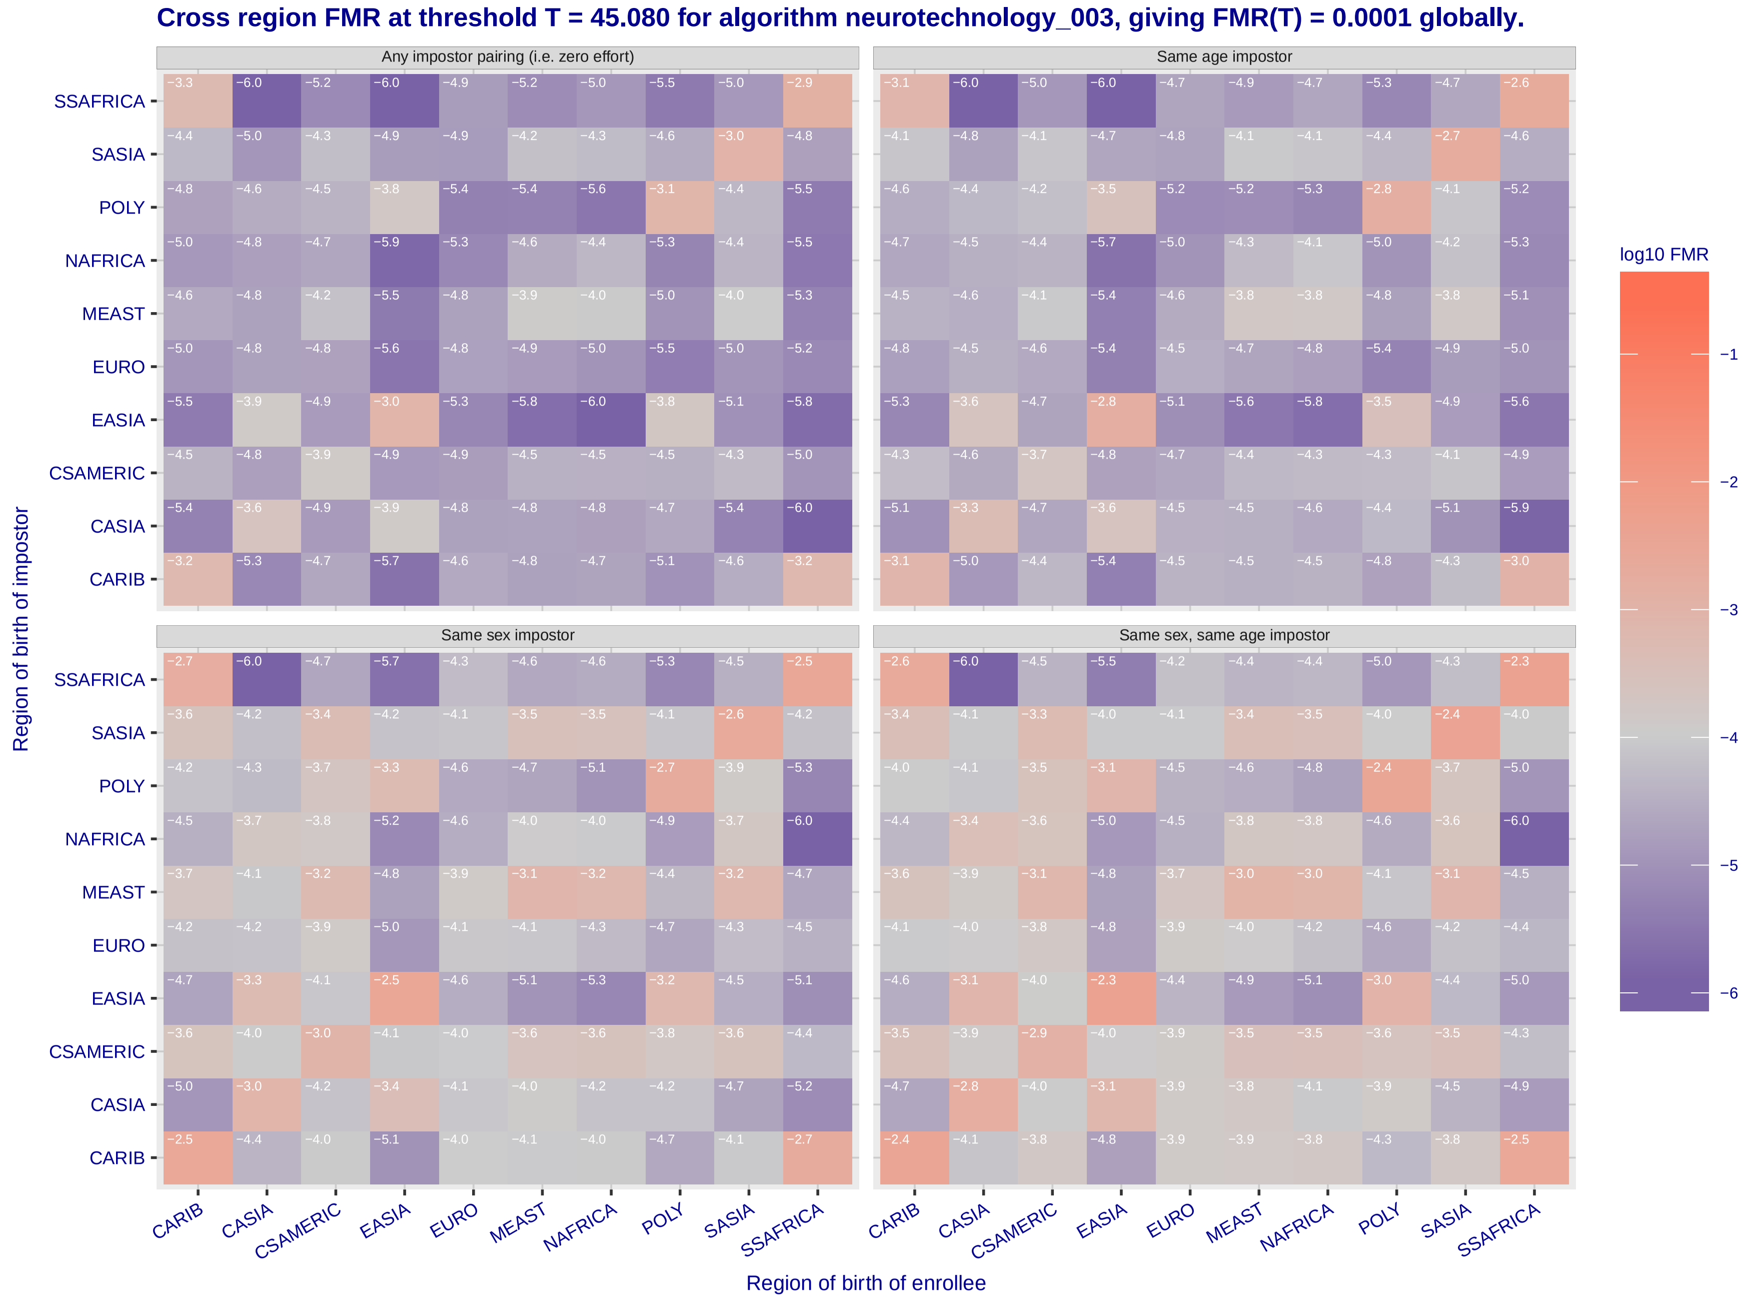

Figure 52: For algorithm neurotechnology-003 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

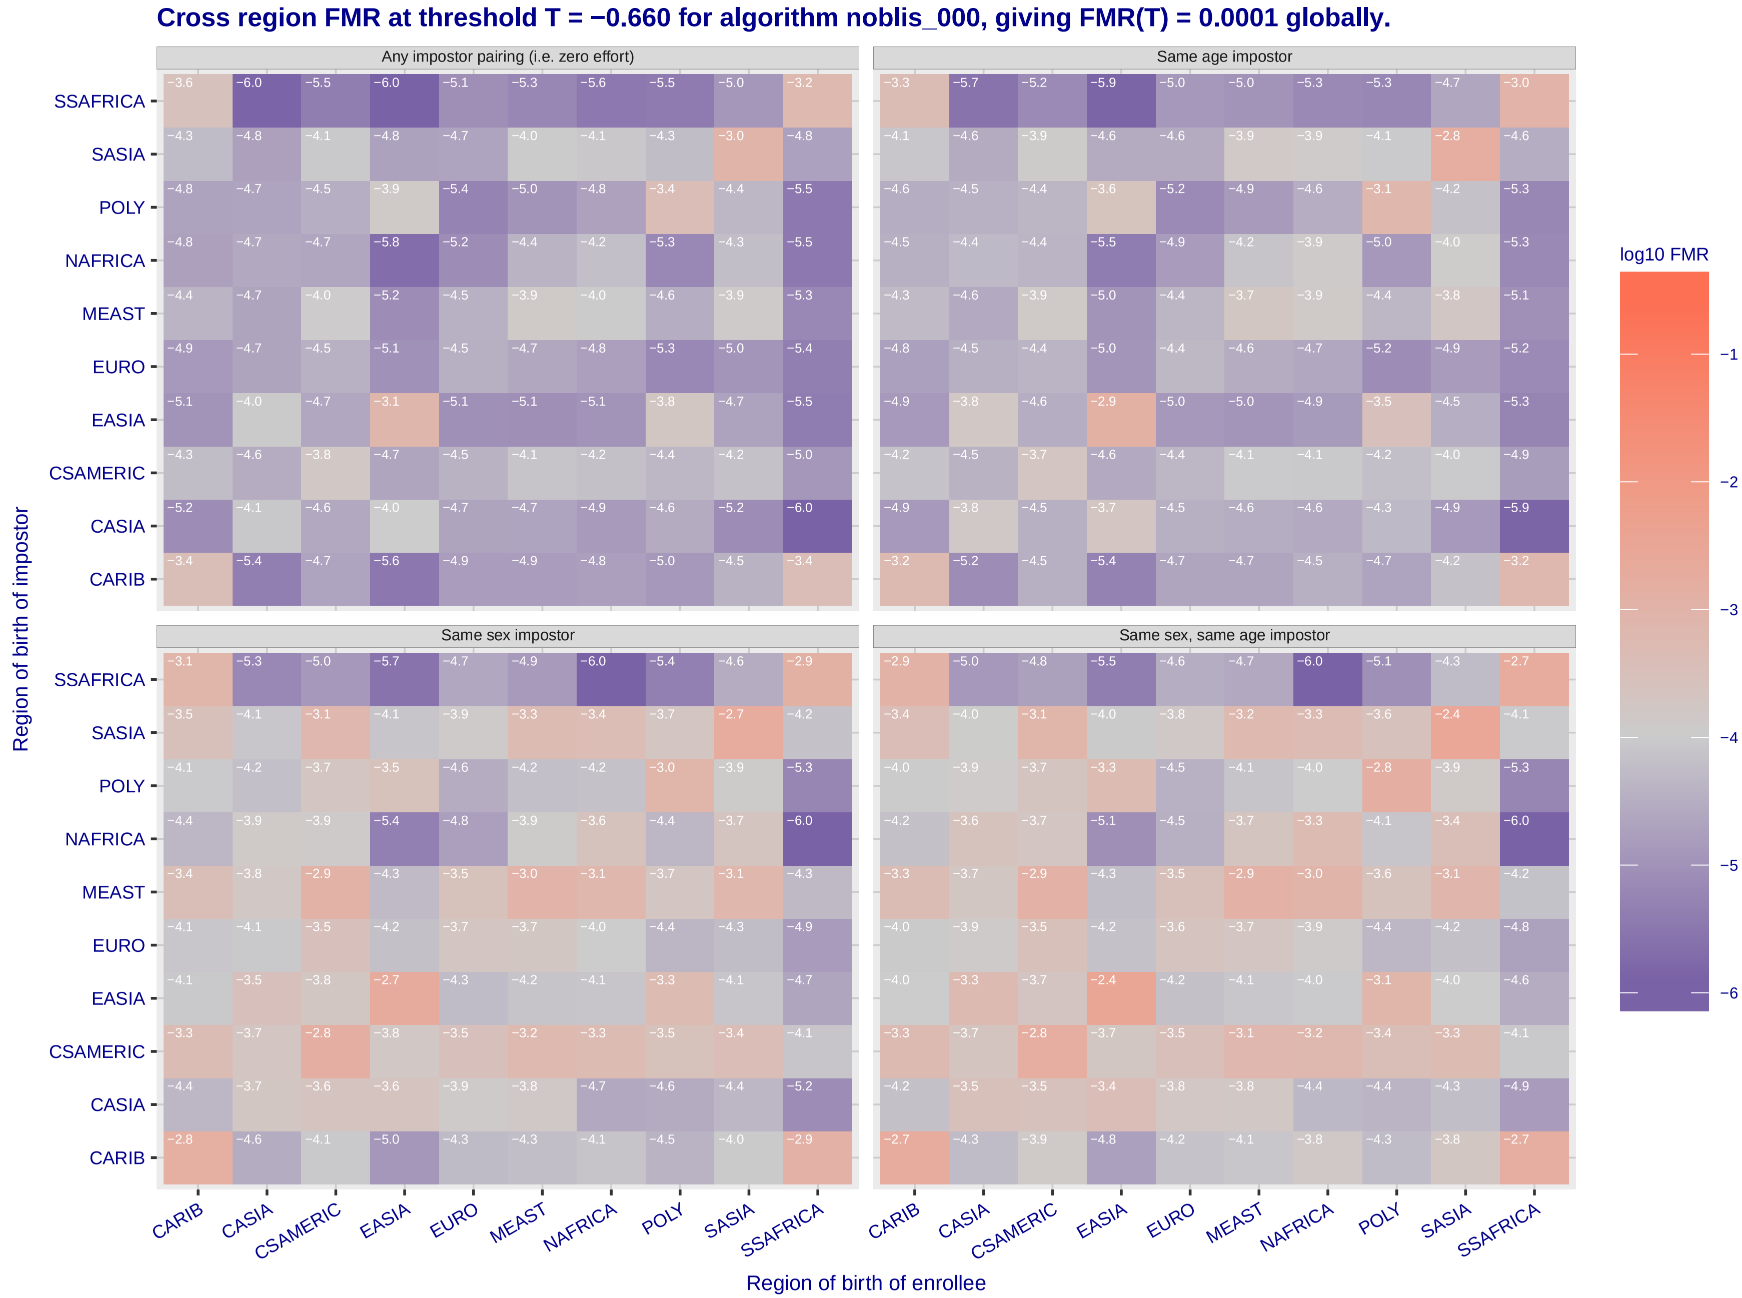

Figure 53: For algorithm noblis-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

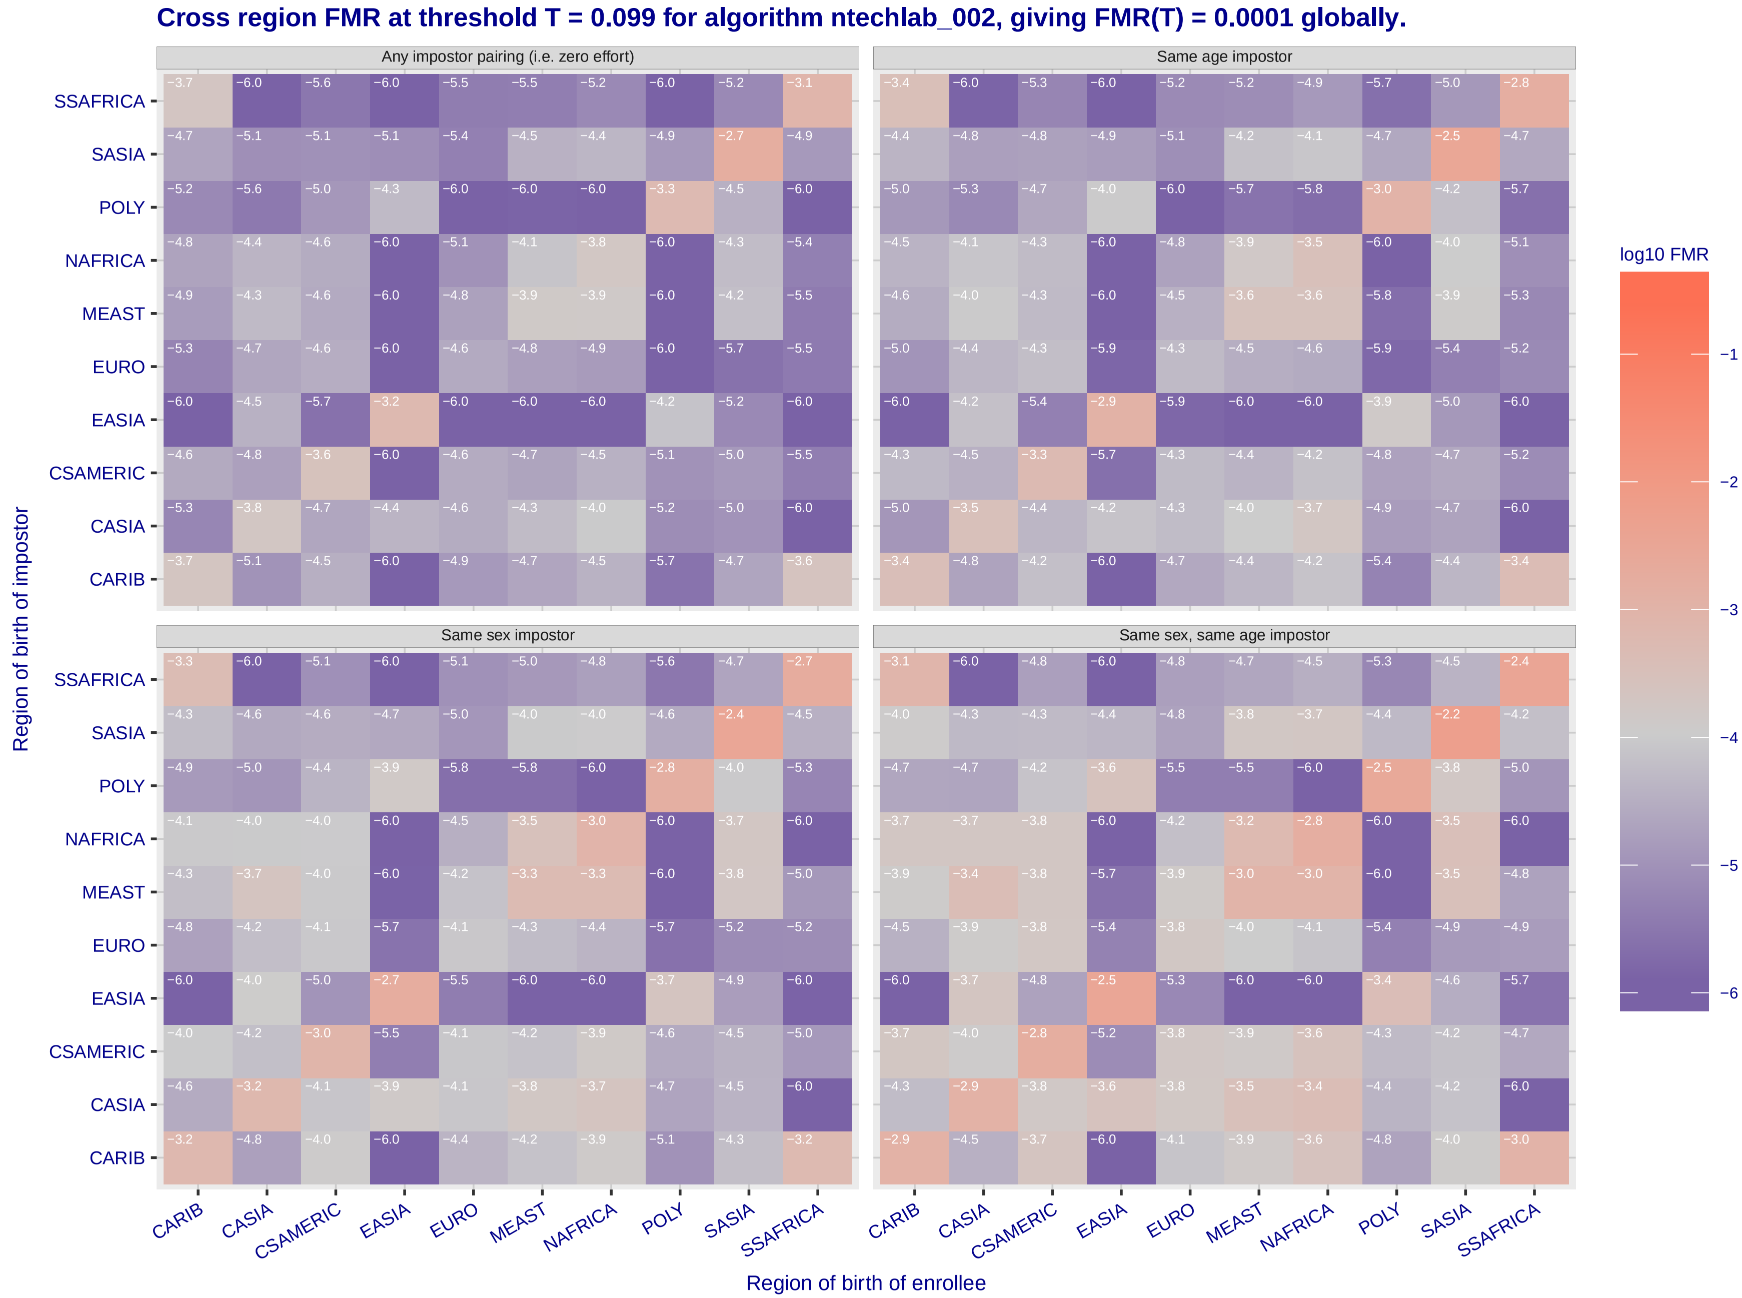

Figure 54: For algorithm ntechlab-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

Figure 55: For algorithm ntechlab-003 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

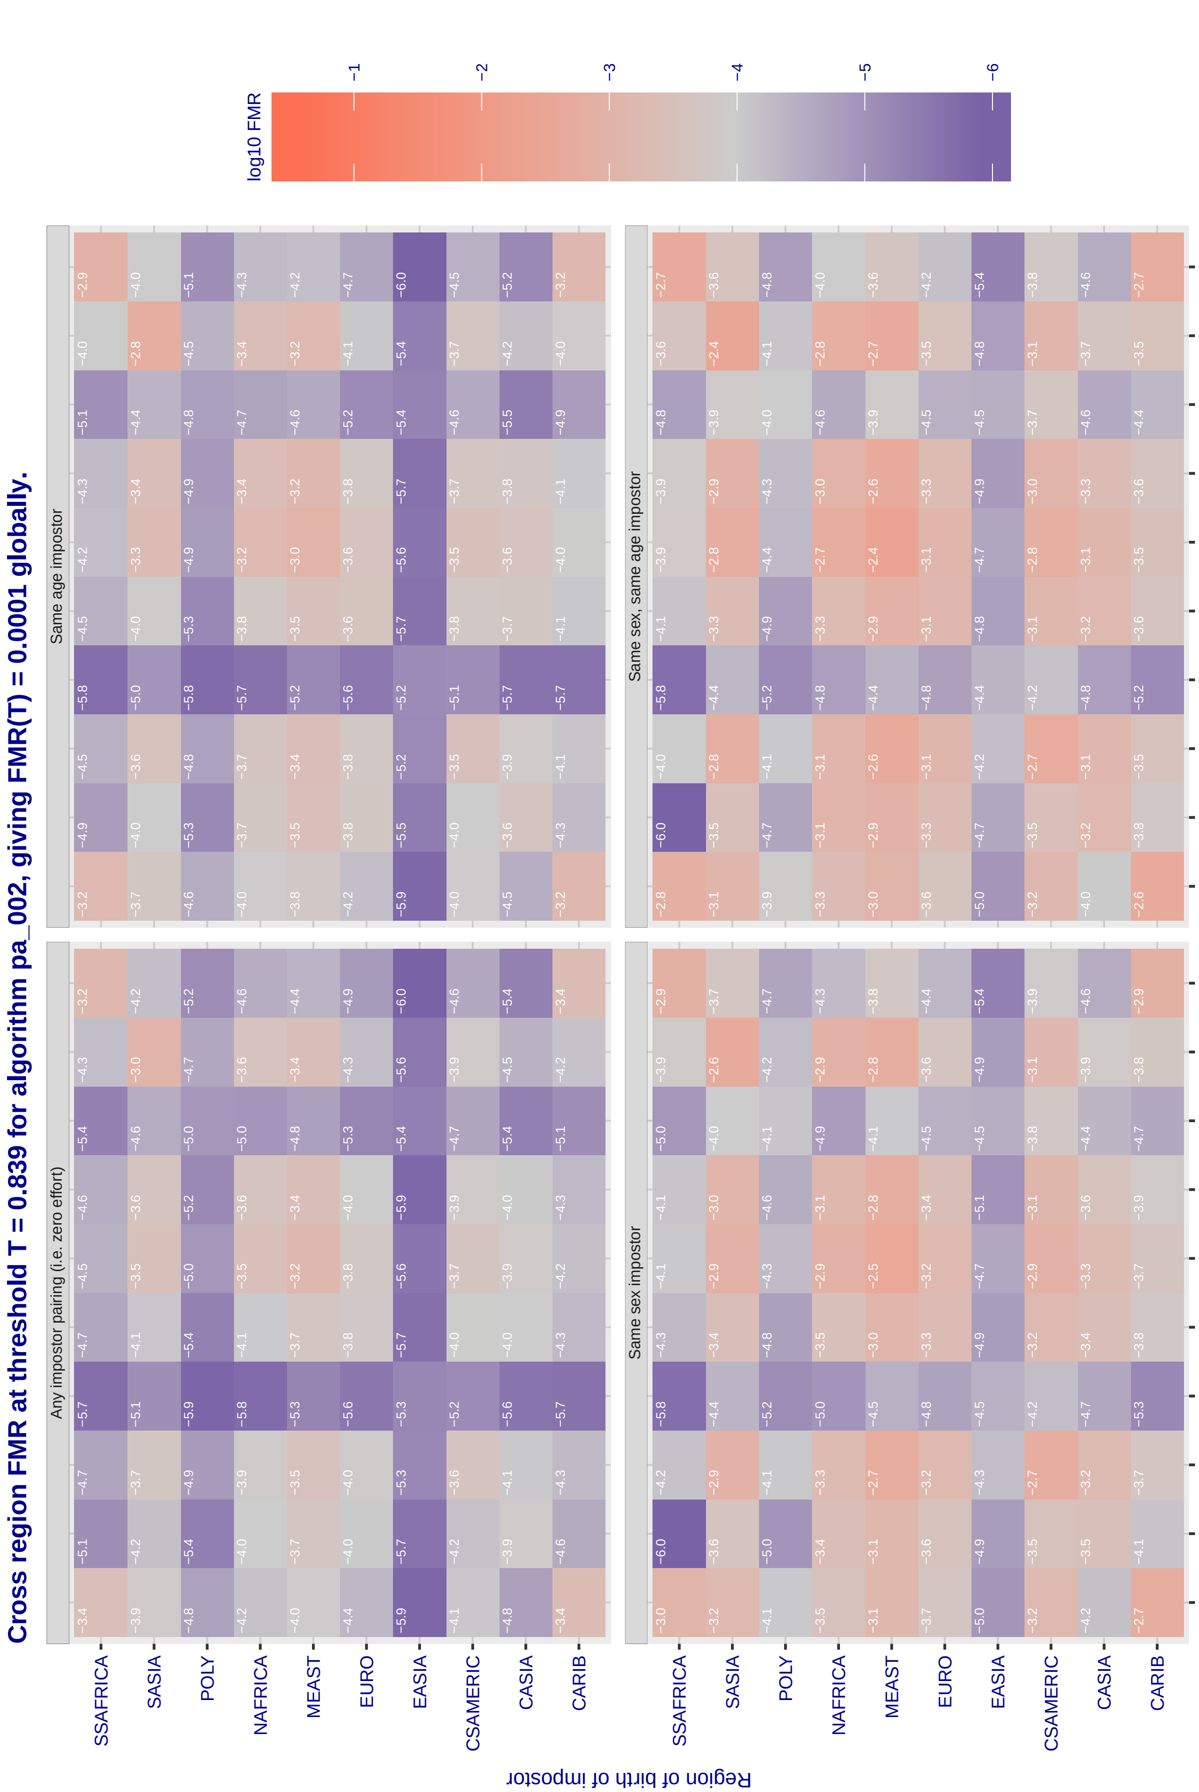

Figure 56: For algorithm pa-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

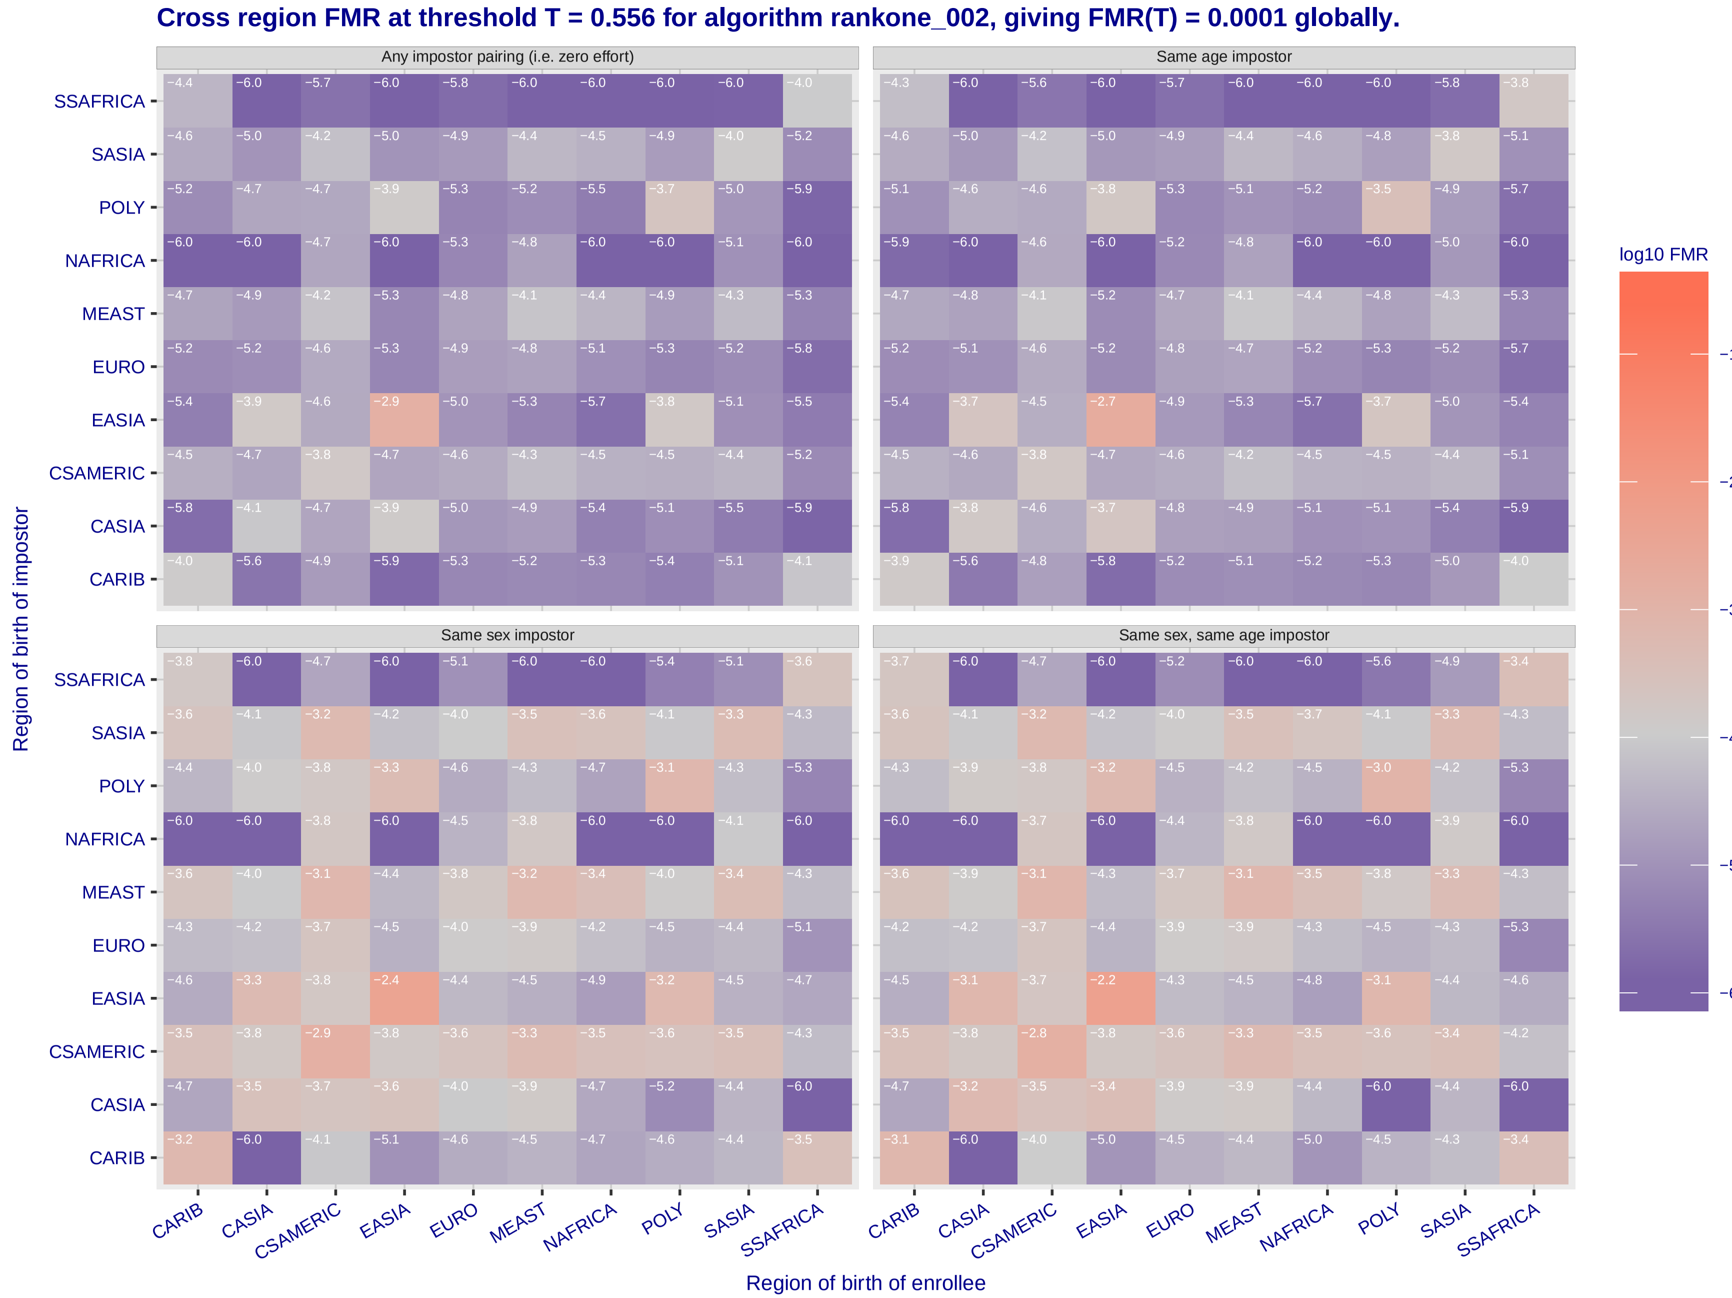

Figure 57: For algorithm rankone-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

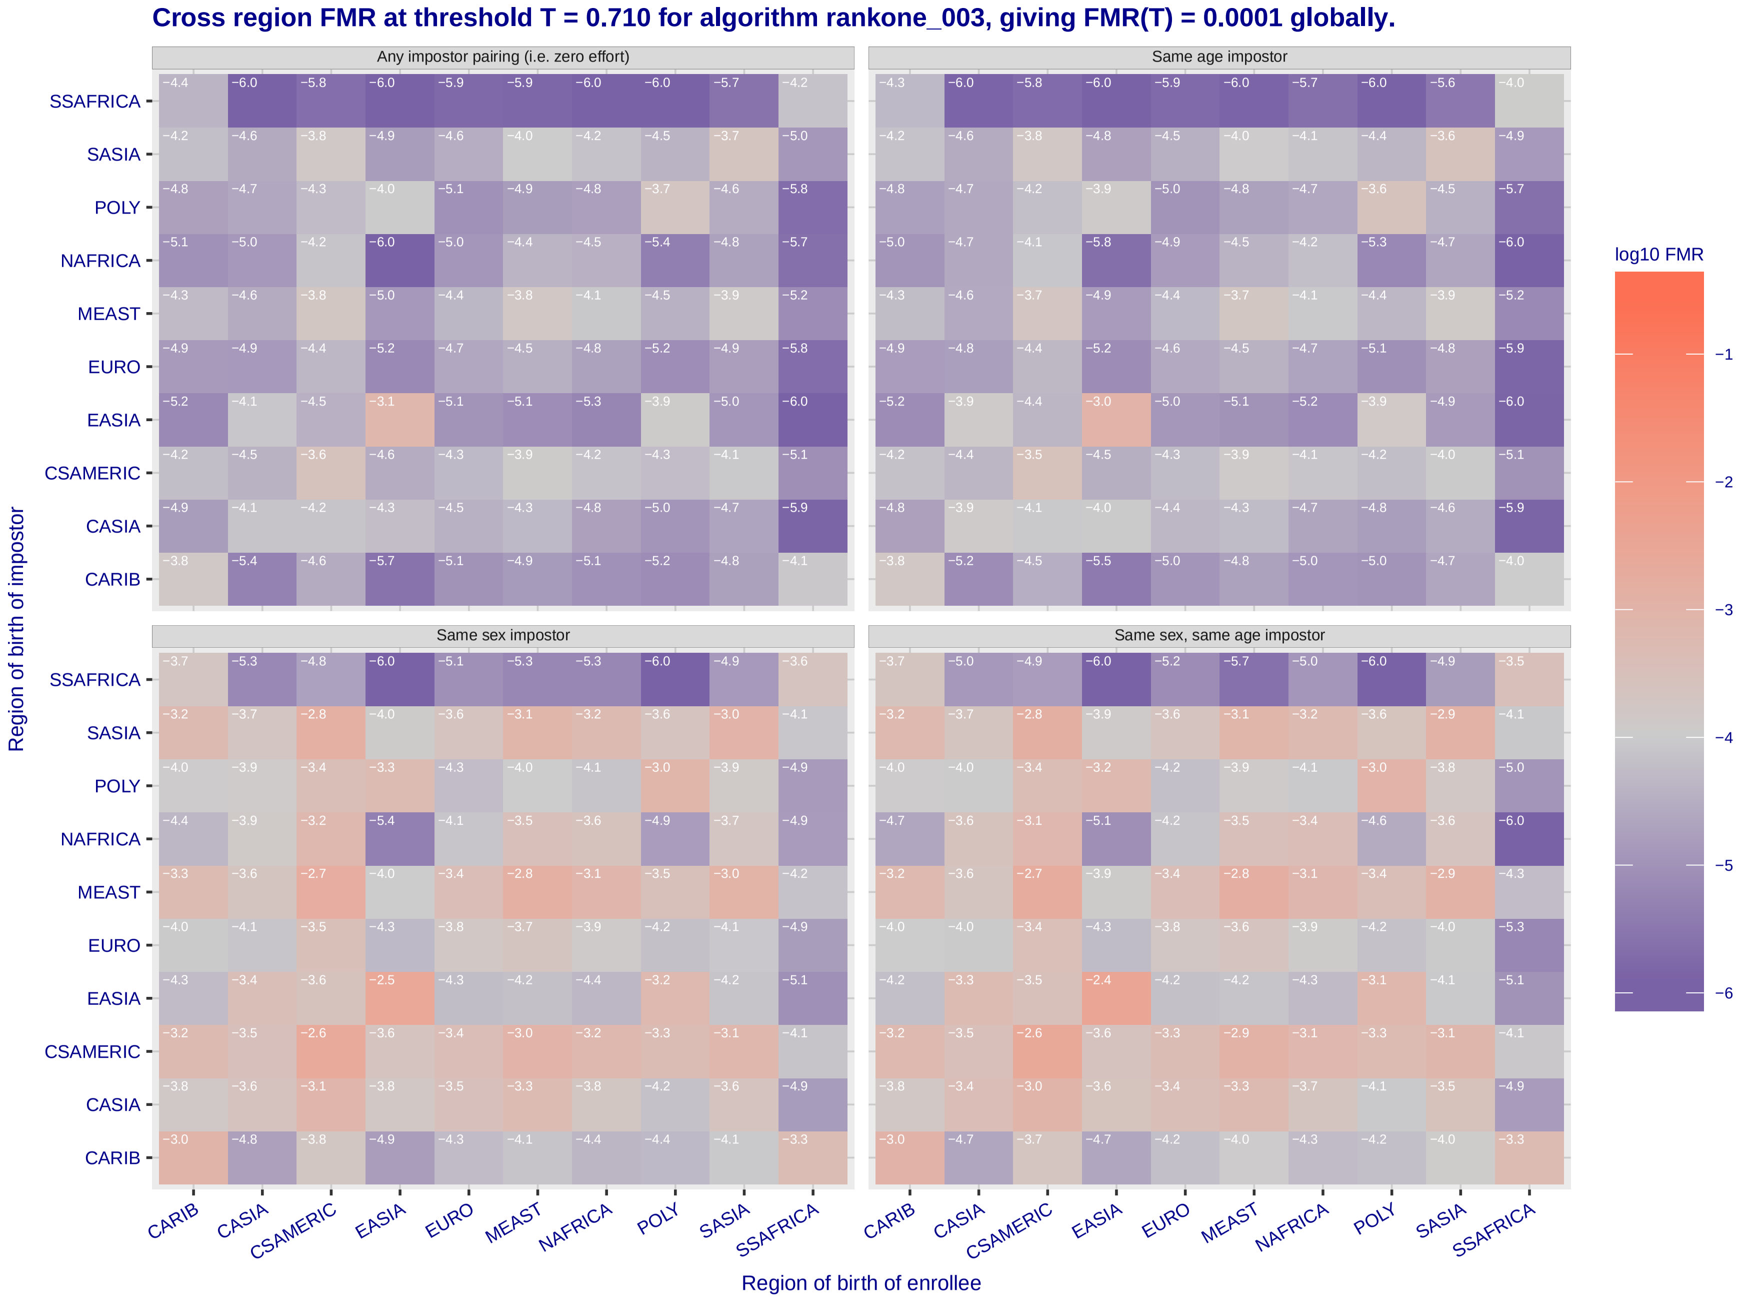

Figure 58: For algorithm rankone-003 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

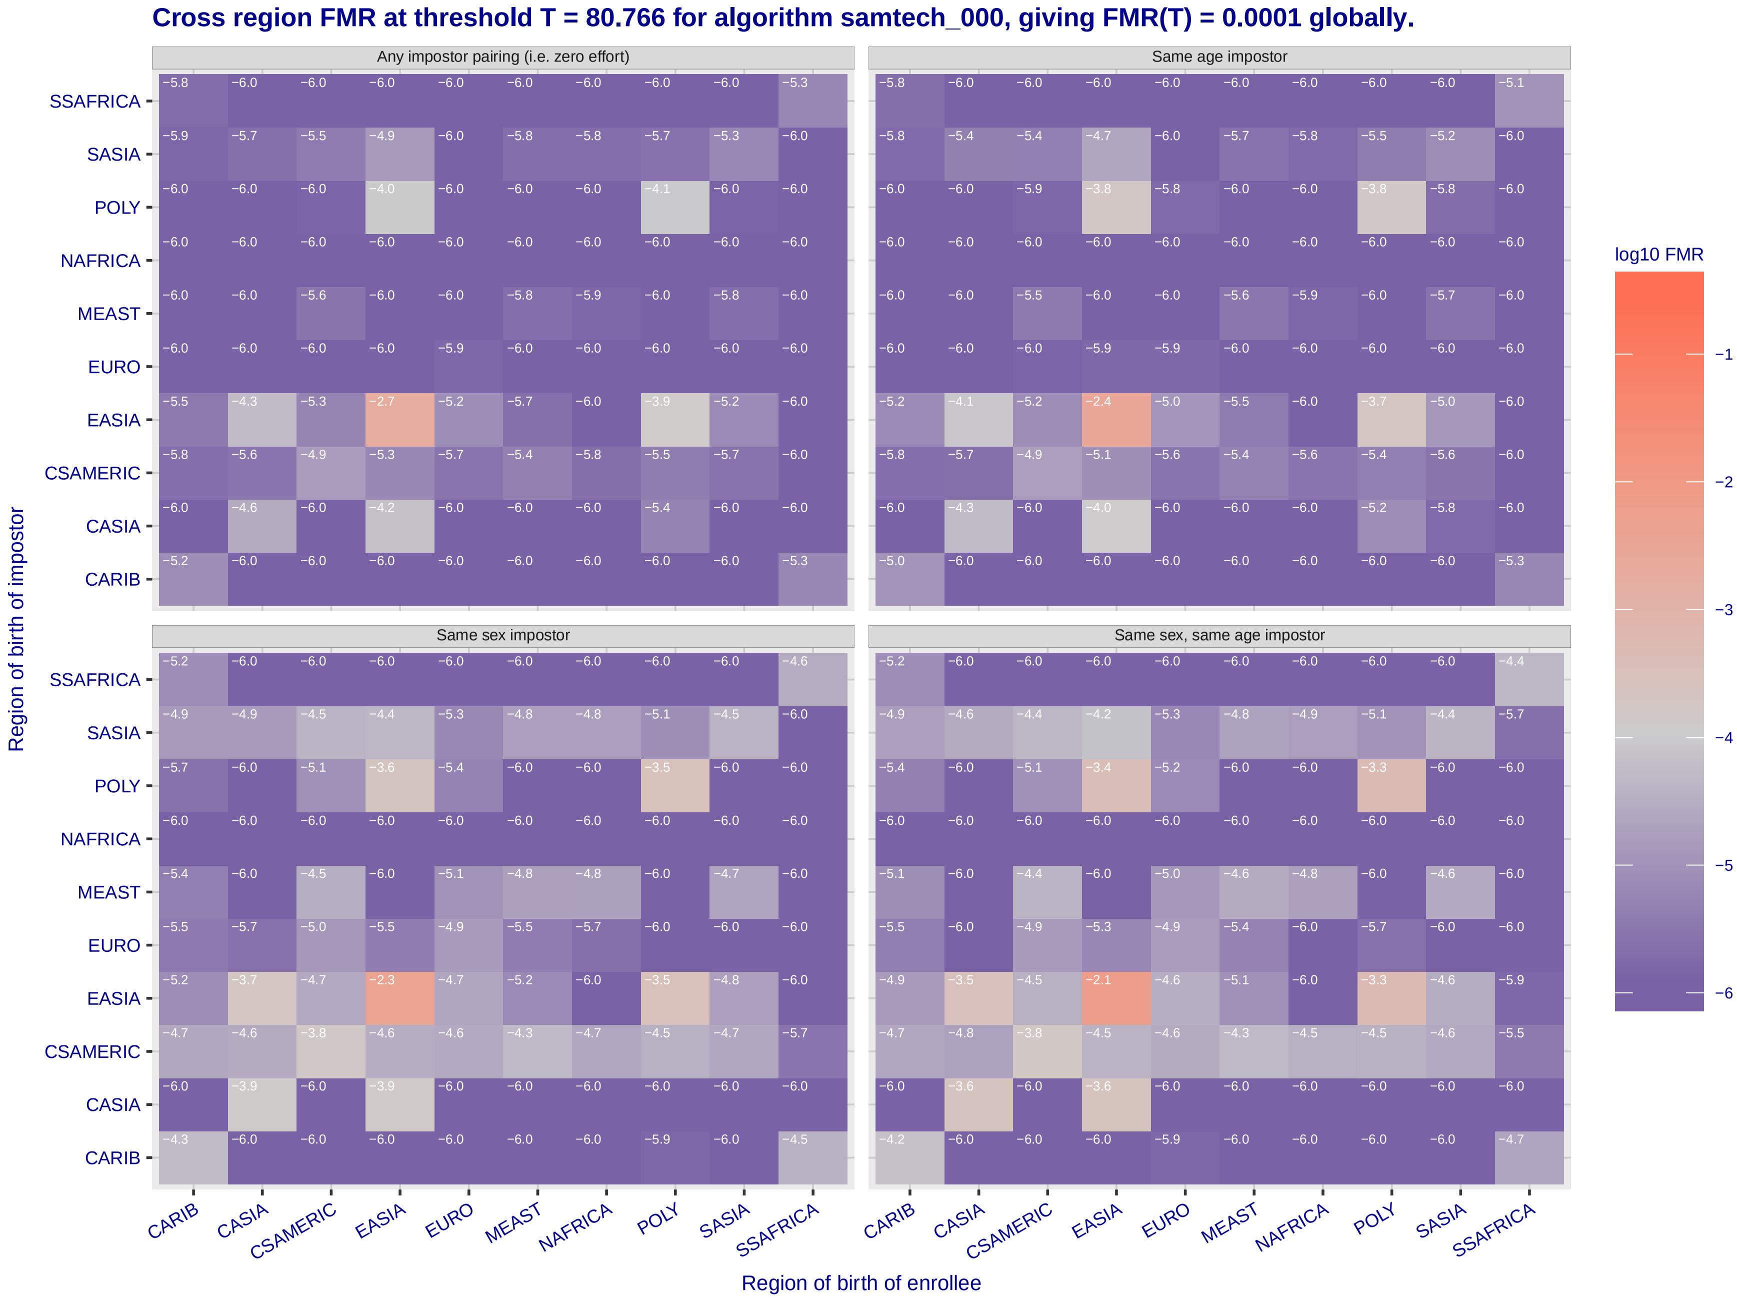

Figure 59: For algorithm samtech-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

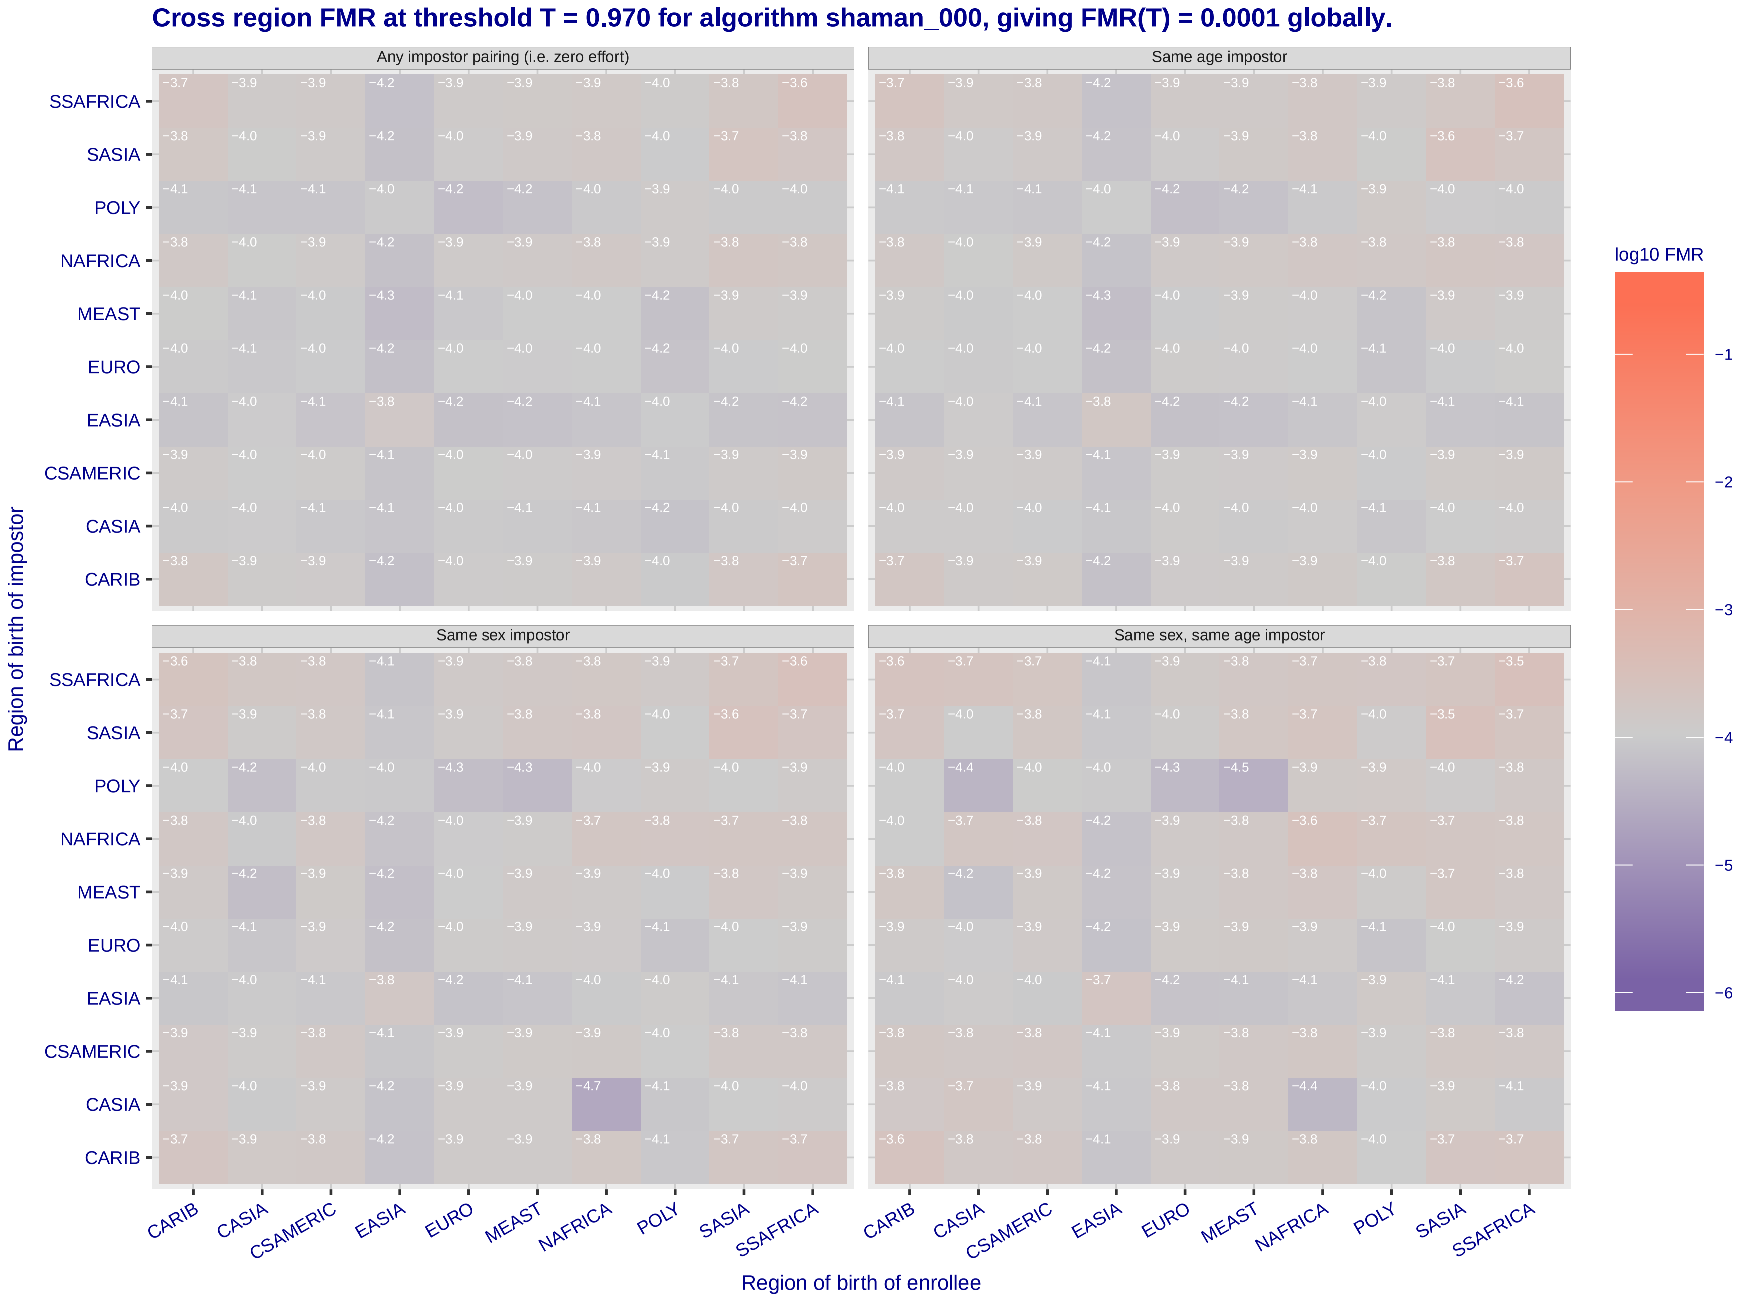

Figure 60: For algorithm shaman-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

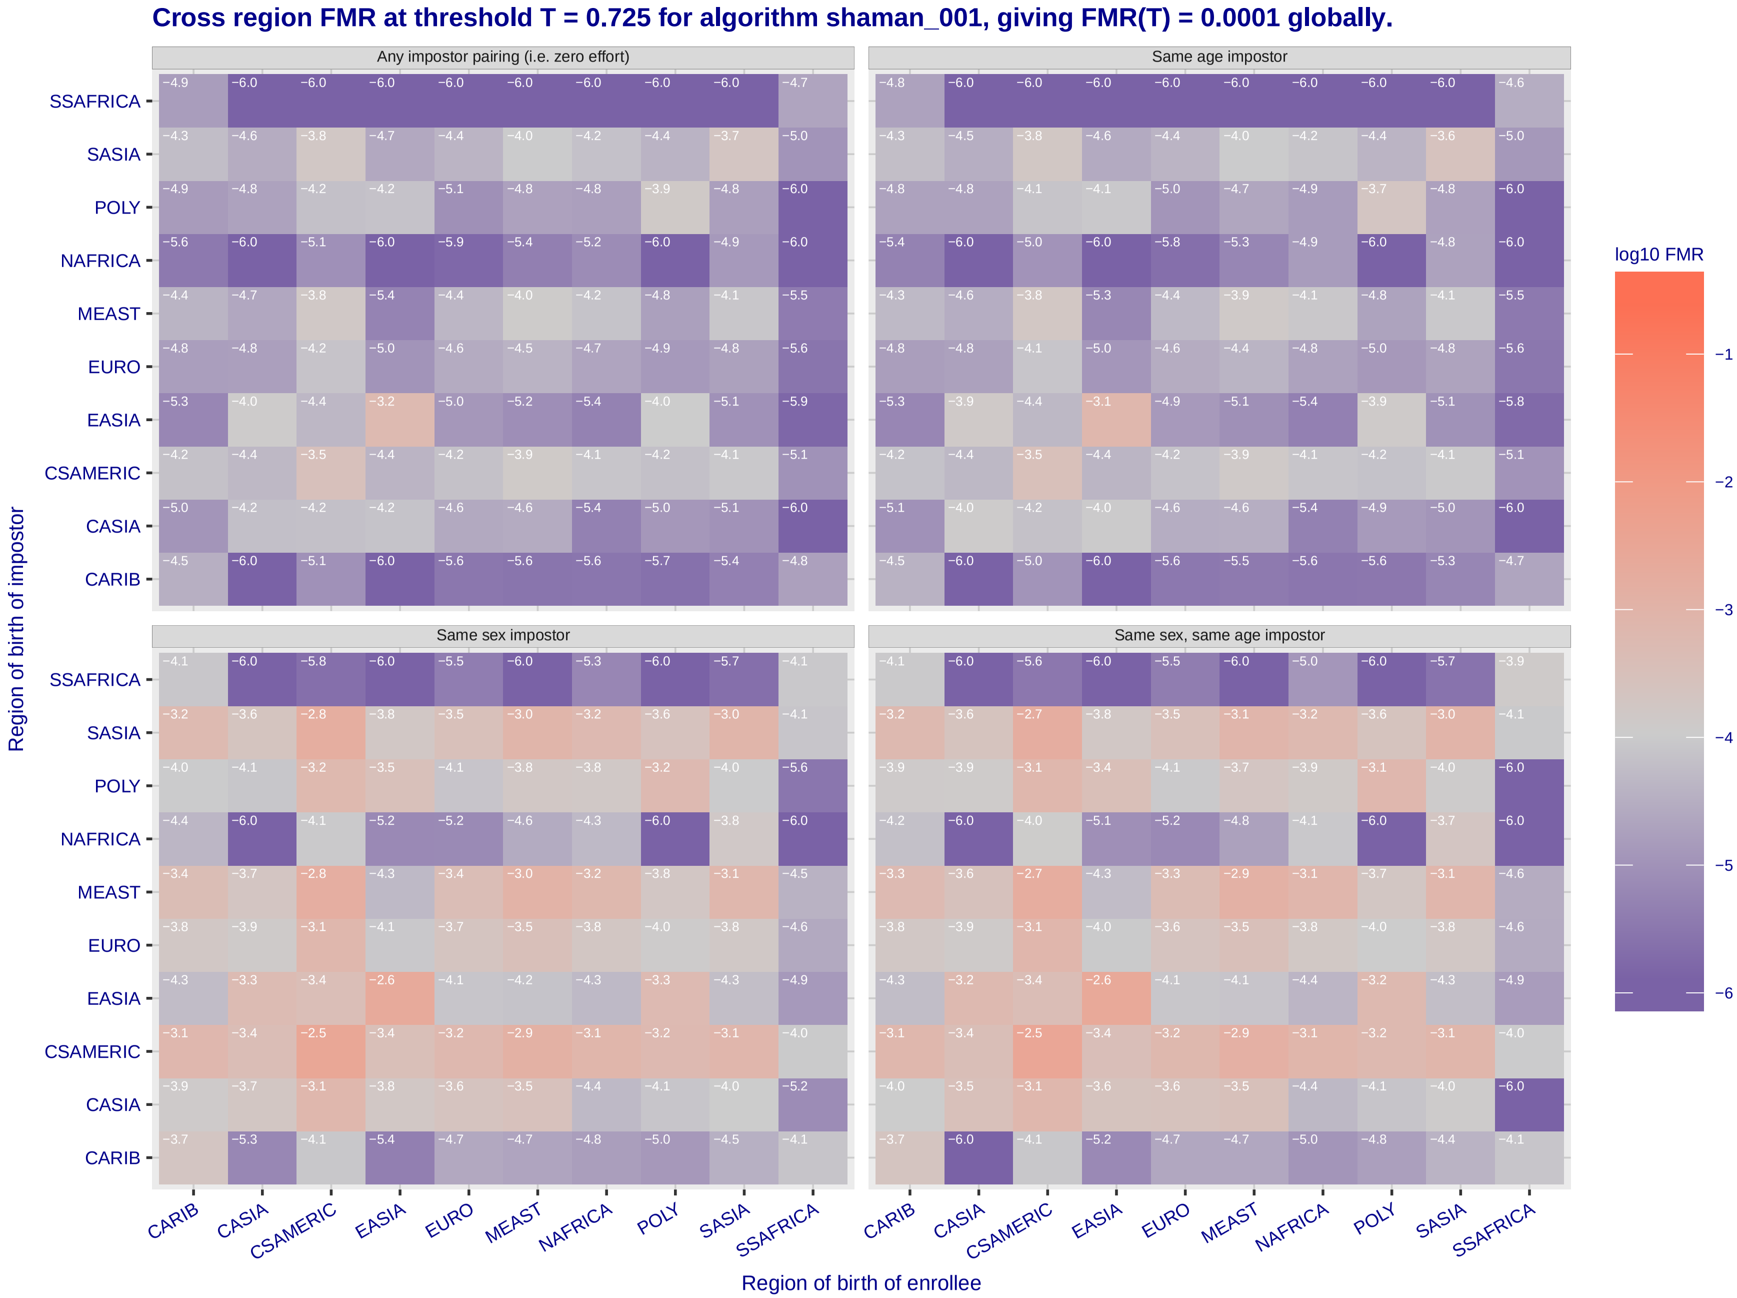

Figure 61: For algorithm shaman-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

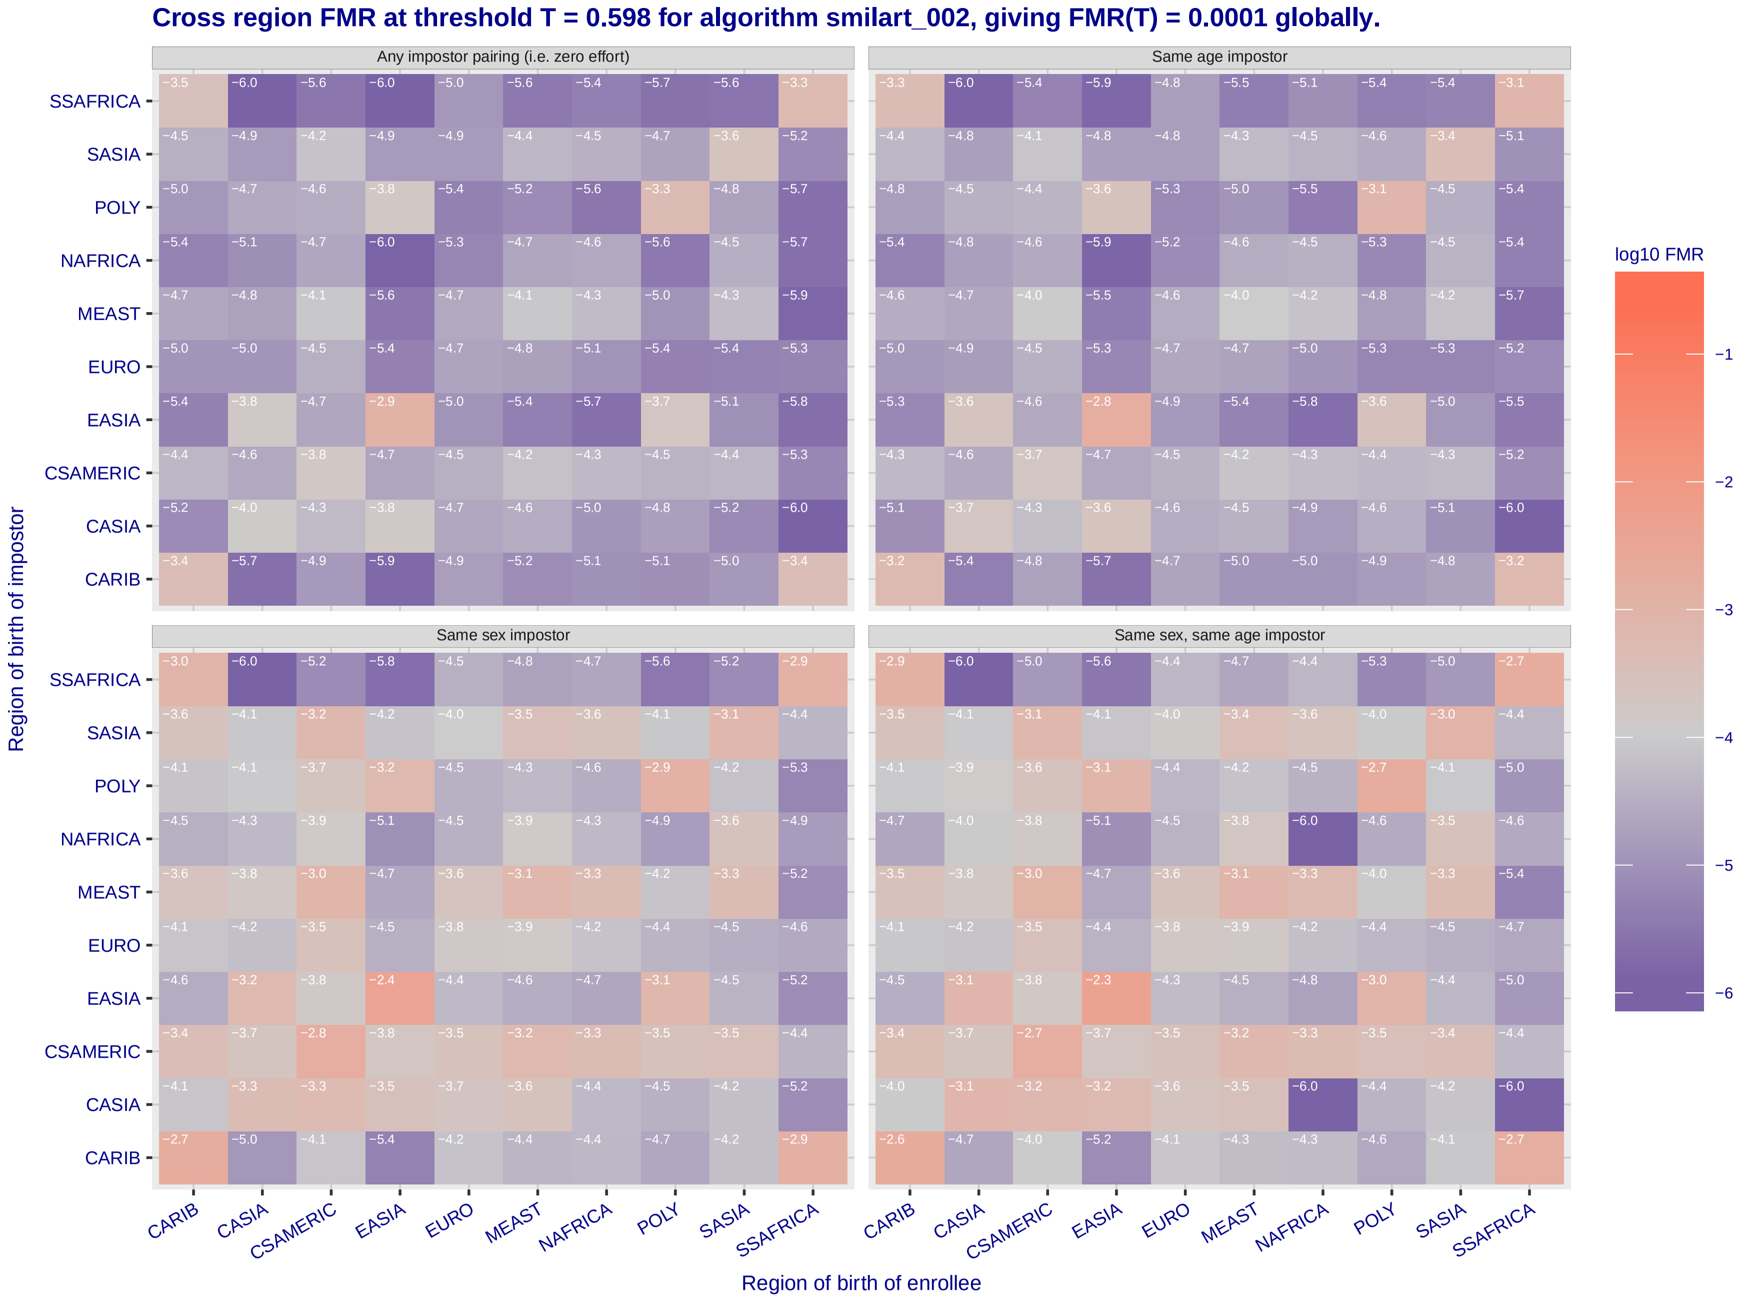

Figure 62: For algorithm smilart-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

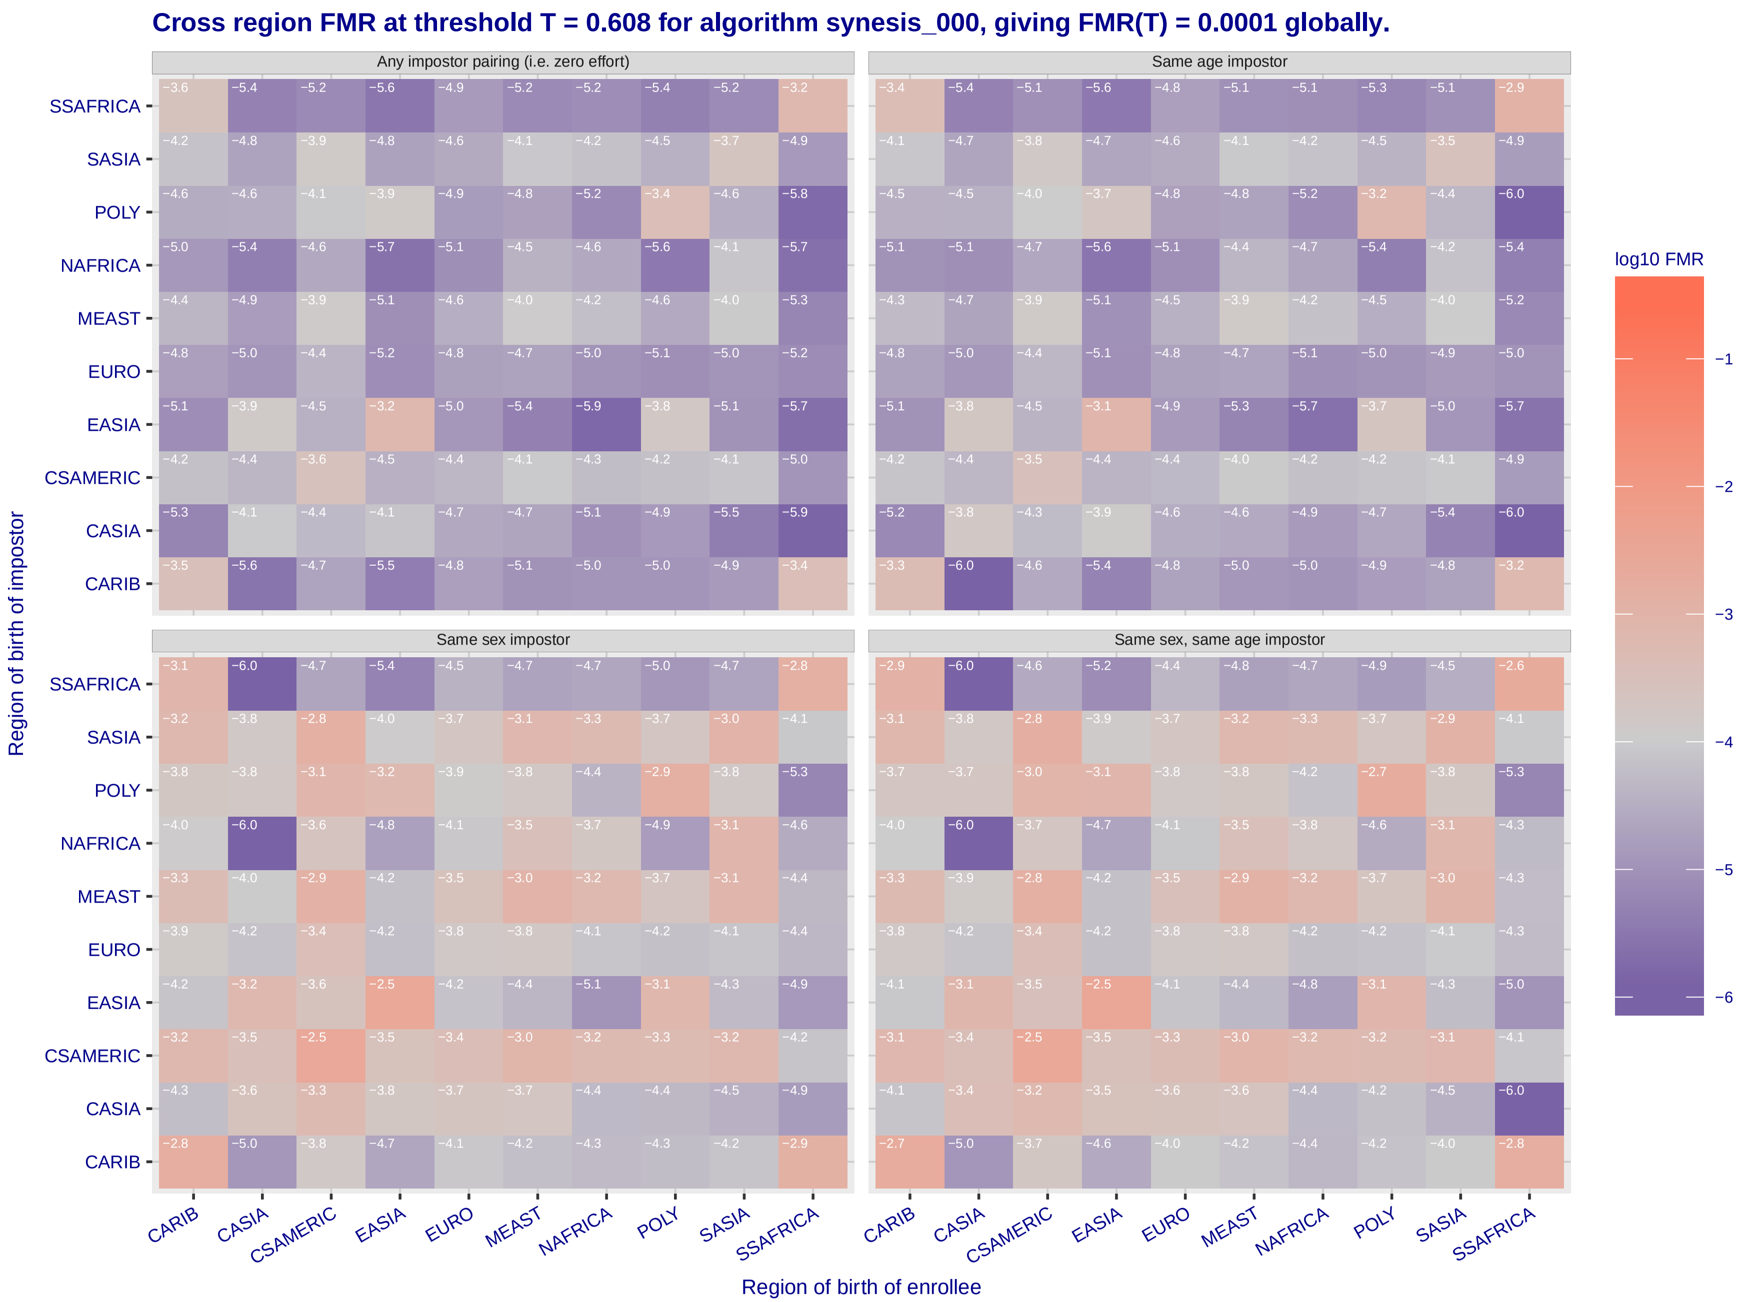

Figure 63: For algorithm synesis-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

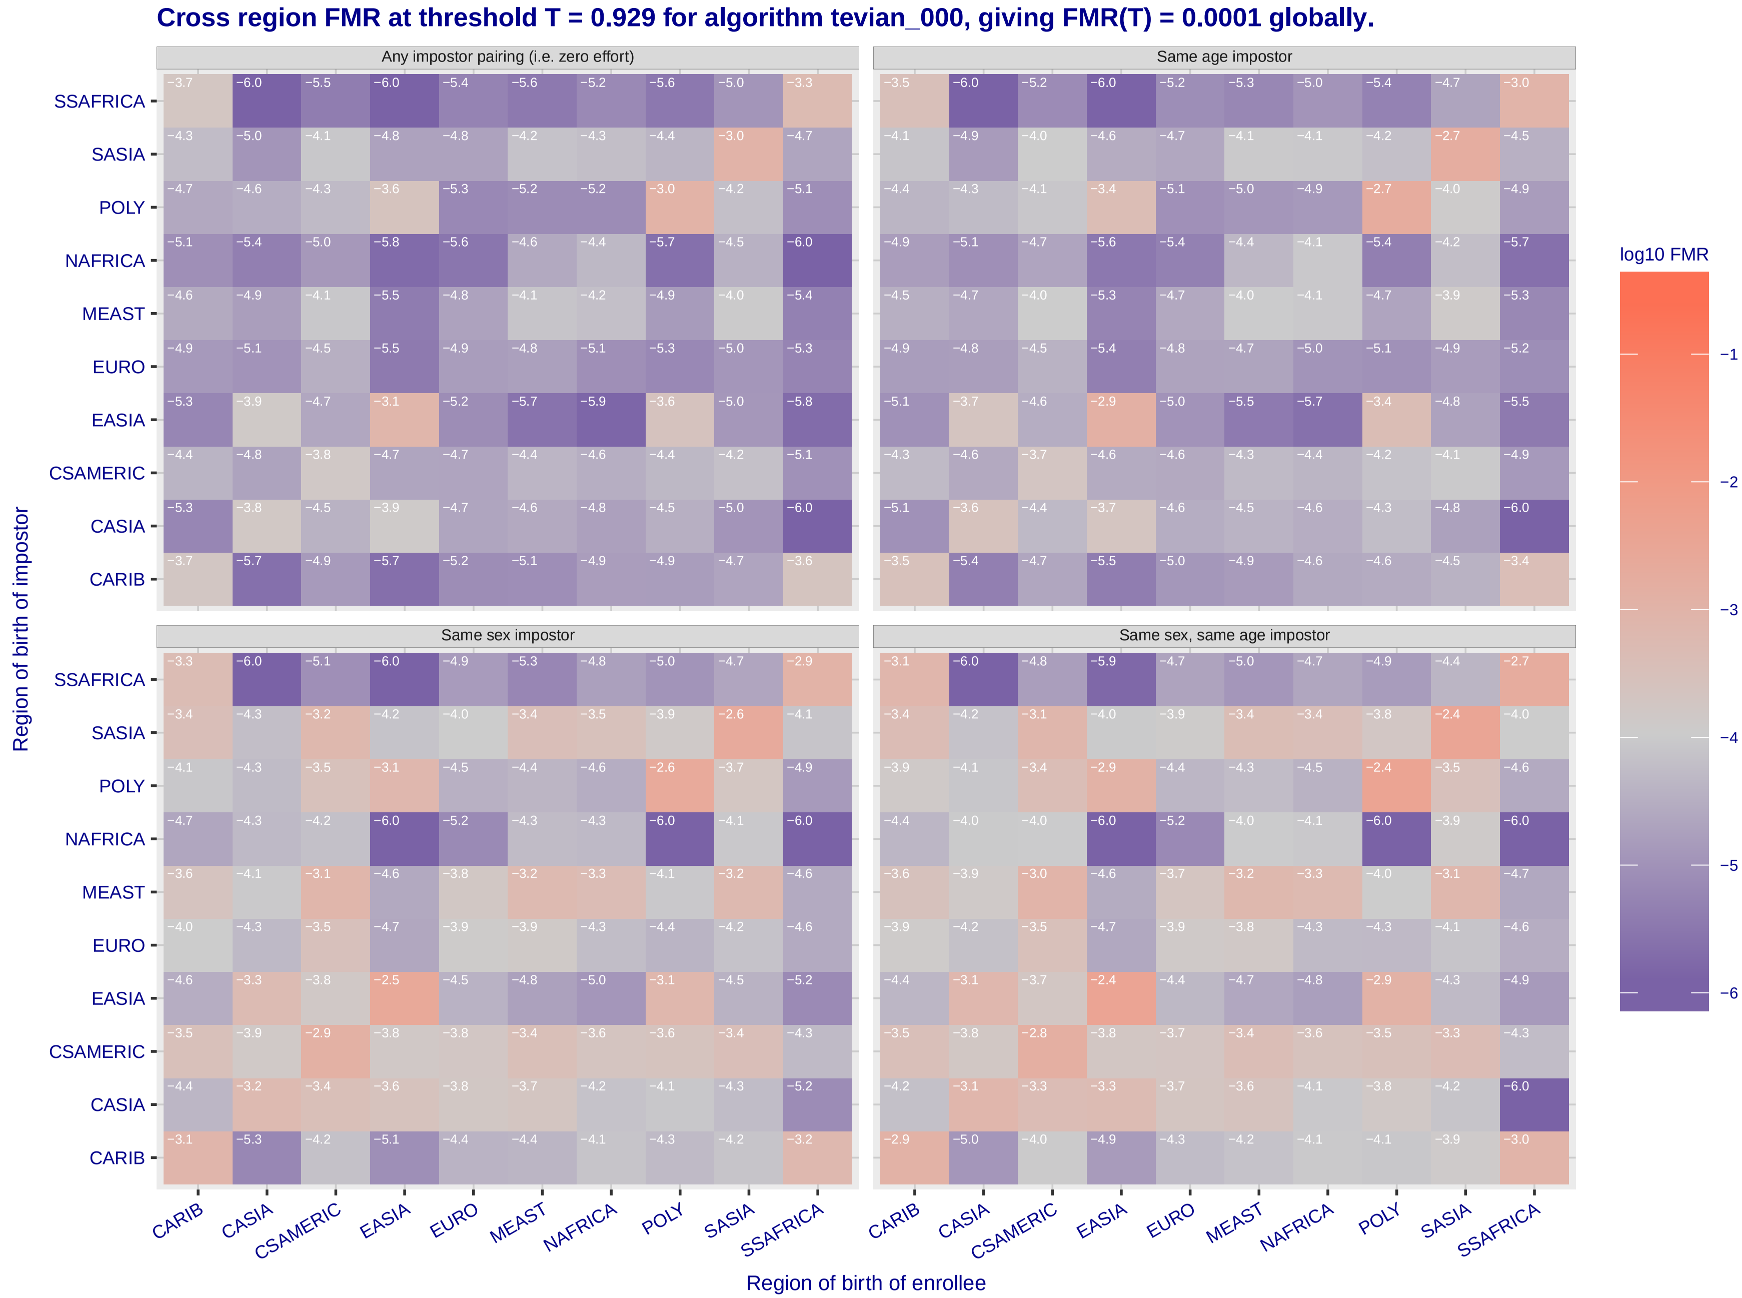

Figure 64: For algorithm tevian-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

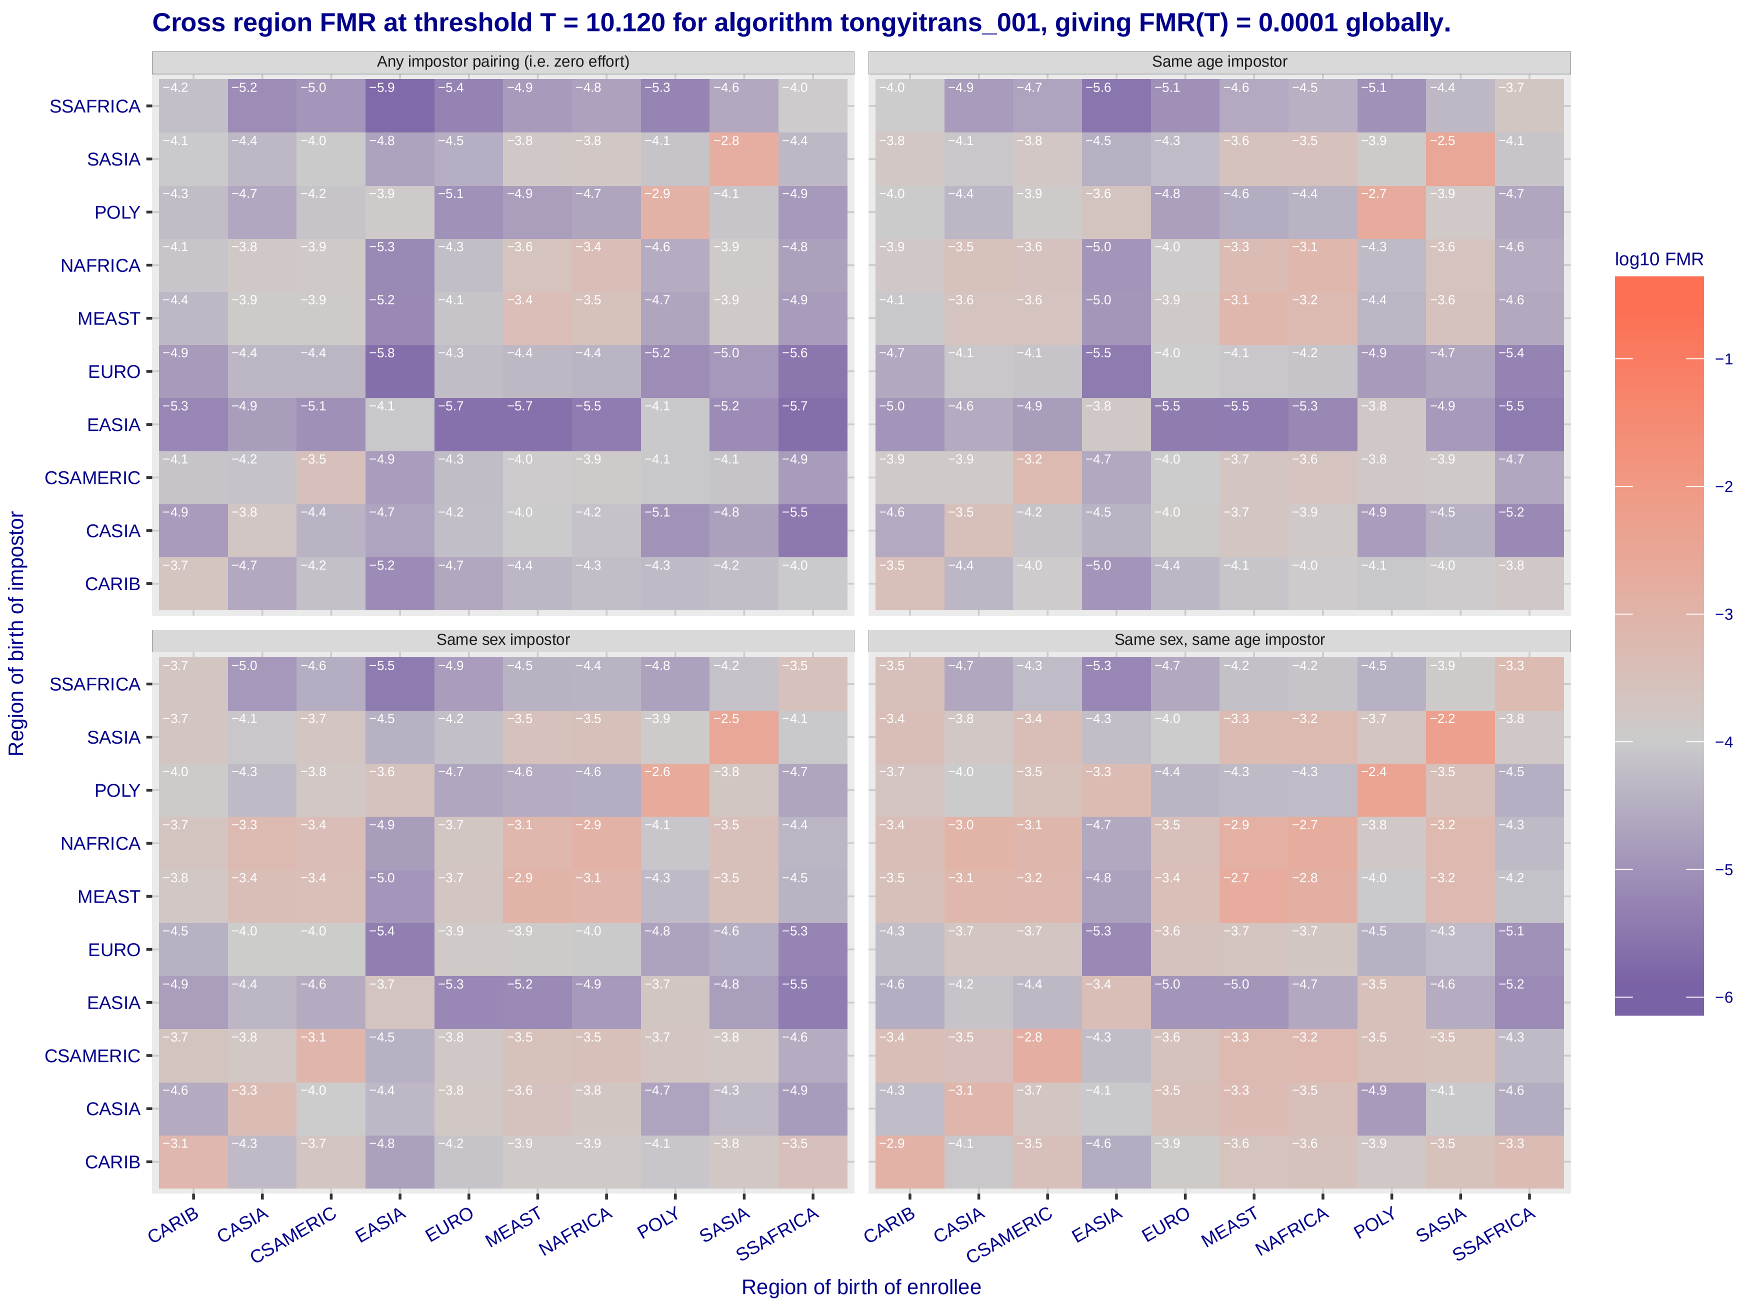

Figure 65: For algorithm tongyitrans-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

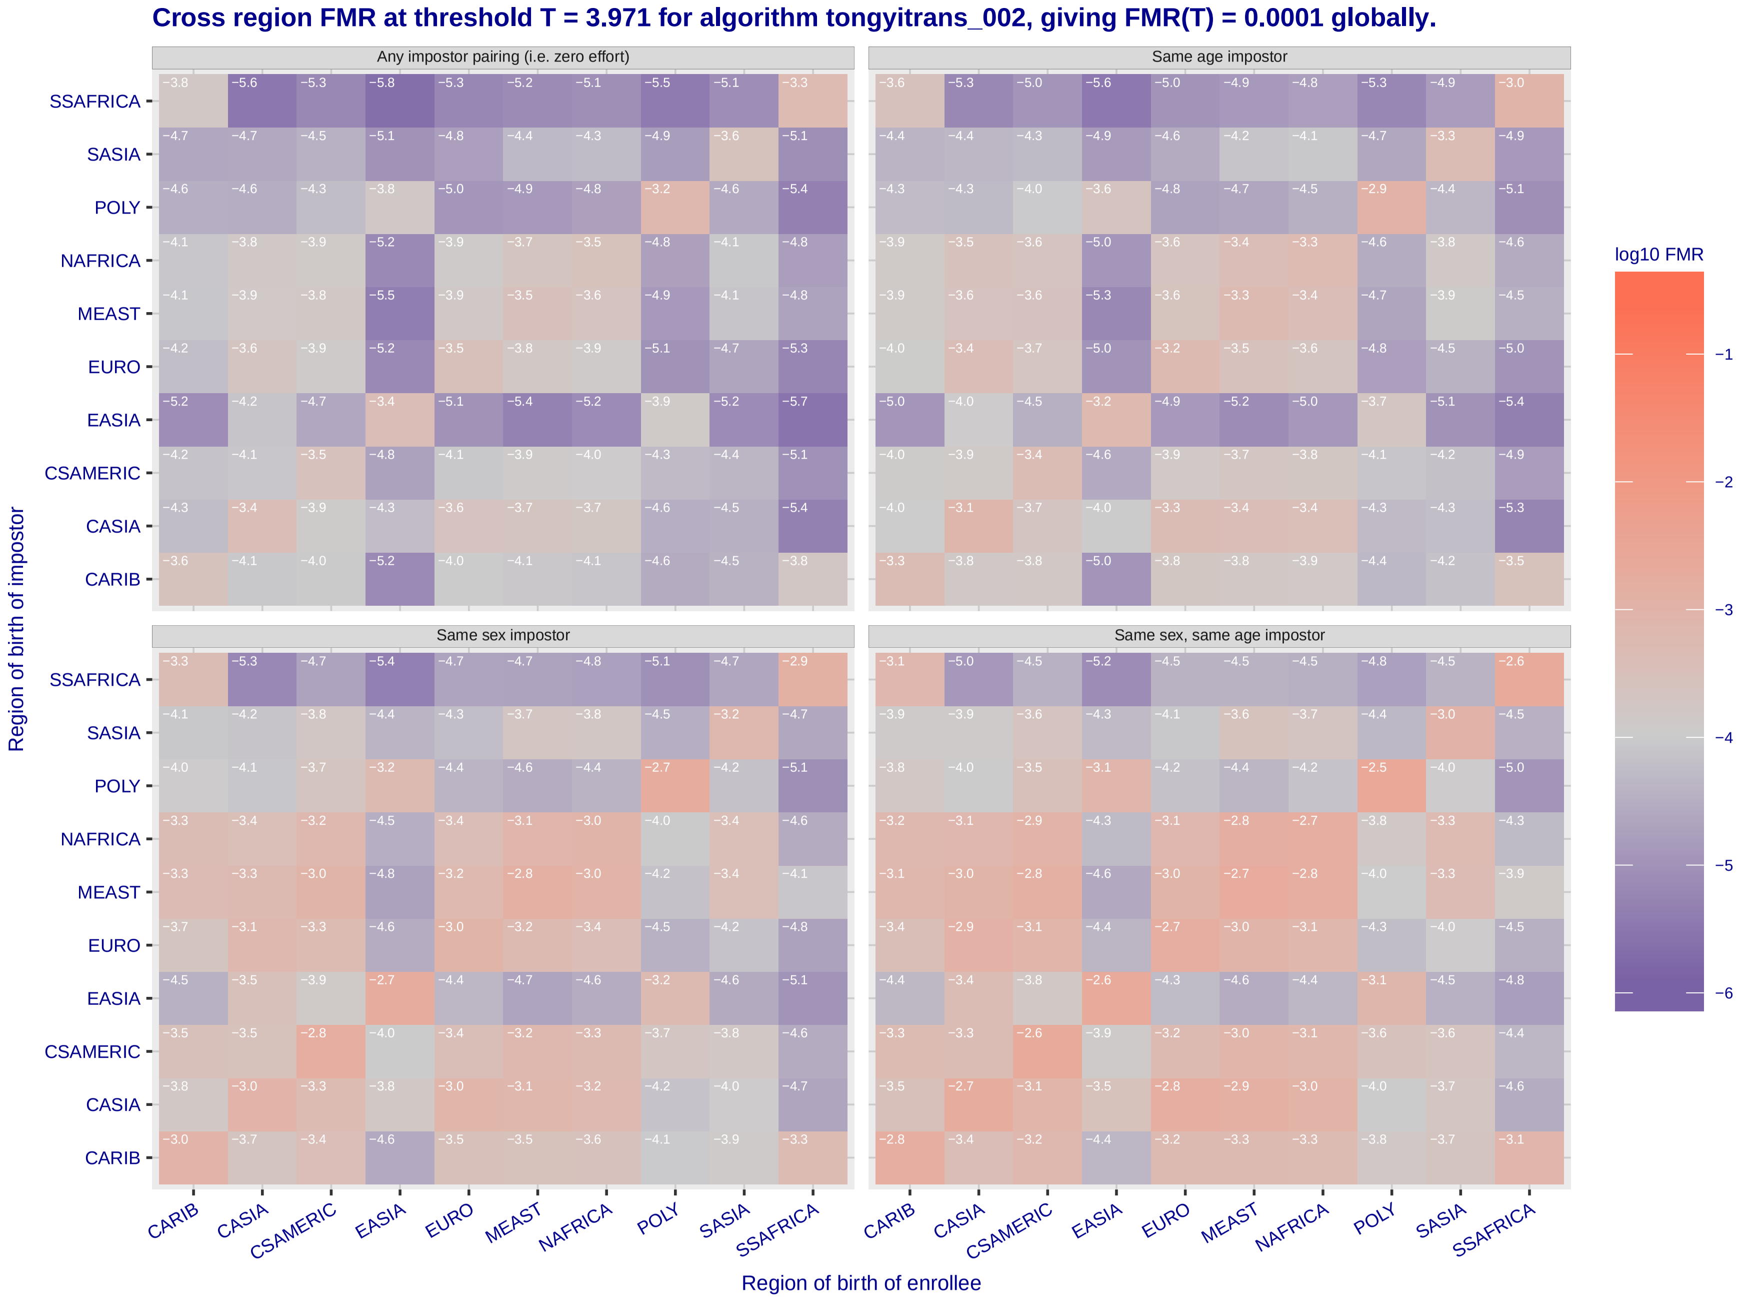

Figure 66: For algorithm tongyitrans-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

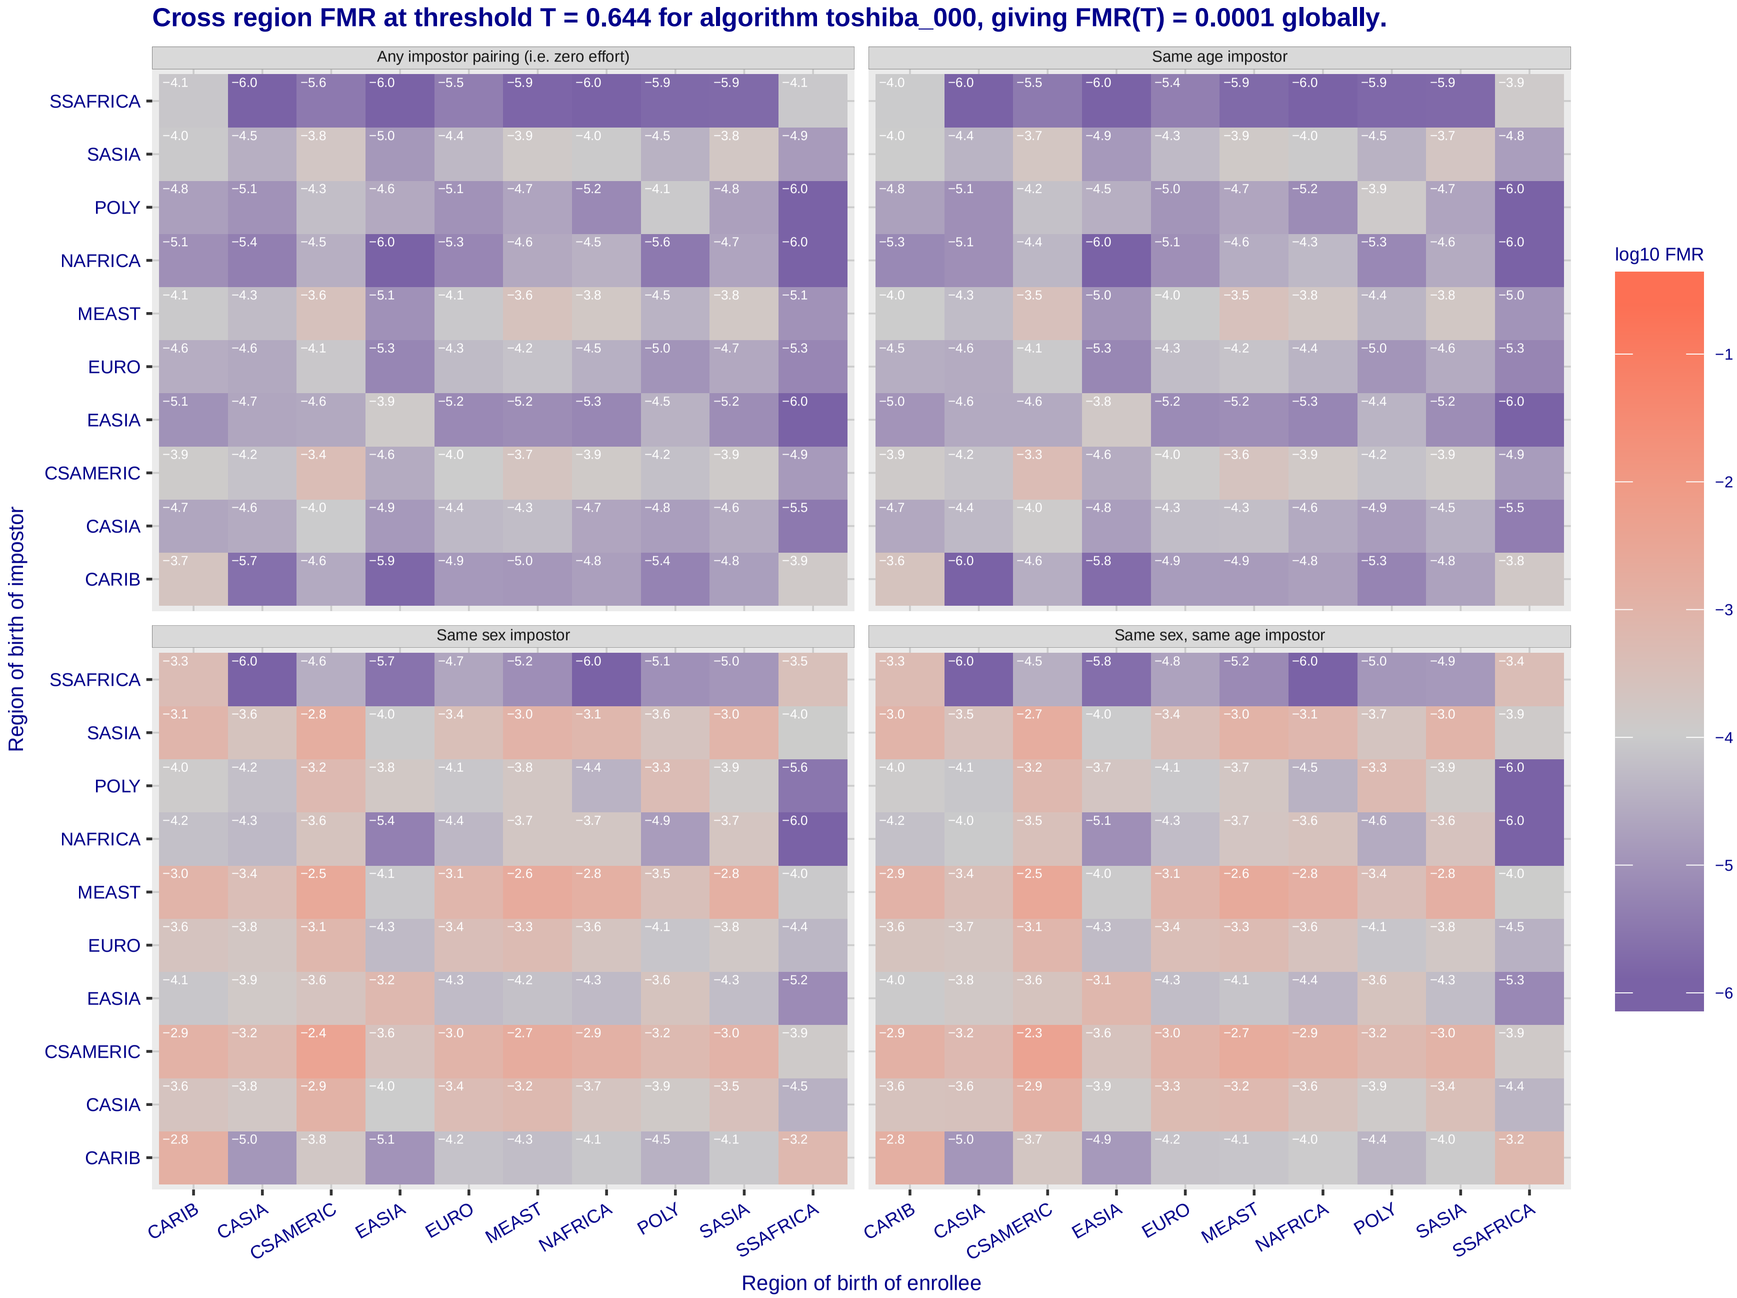

Figure 67: For algorithm toshiba-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

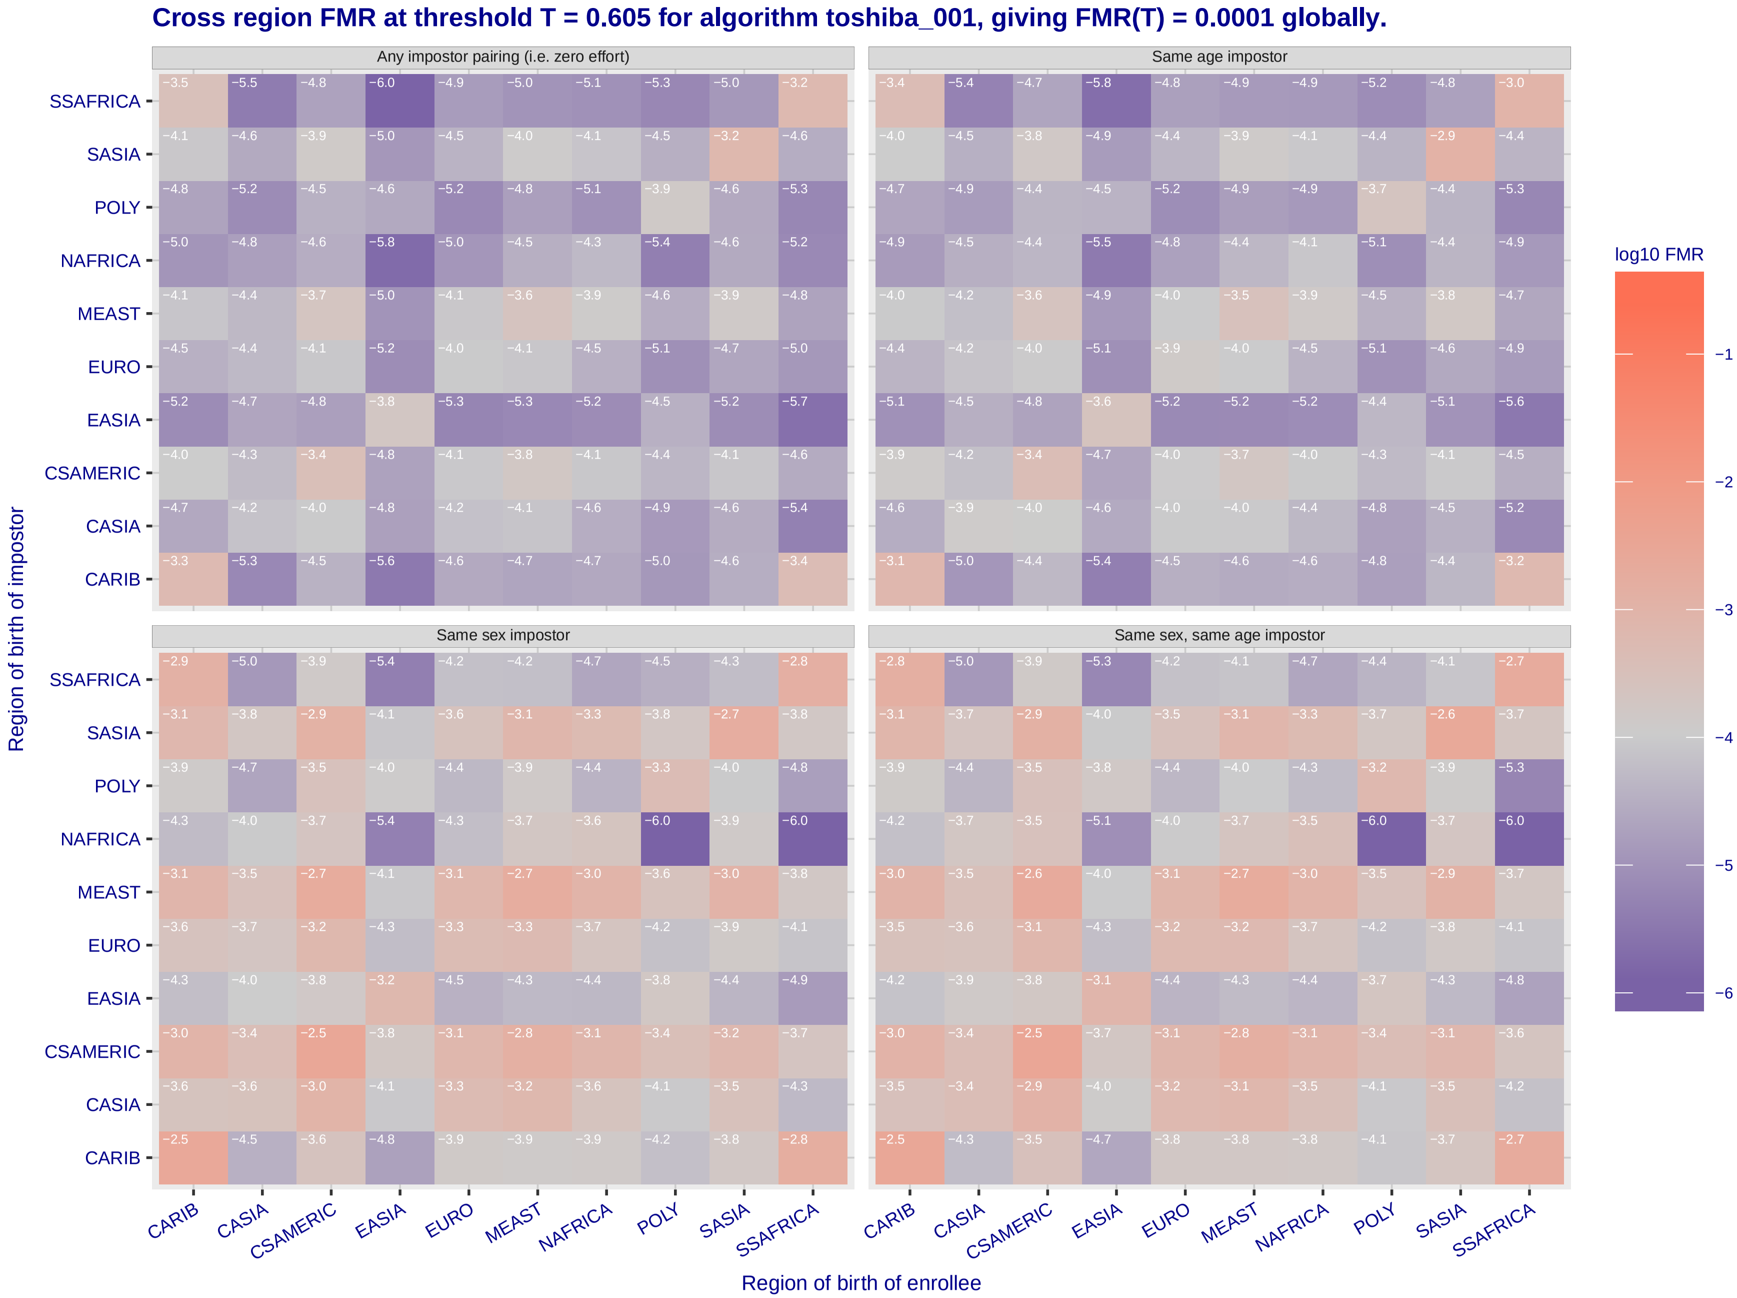

Figure 68: For algorithm toshiba-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

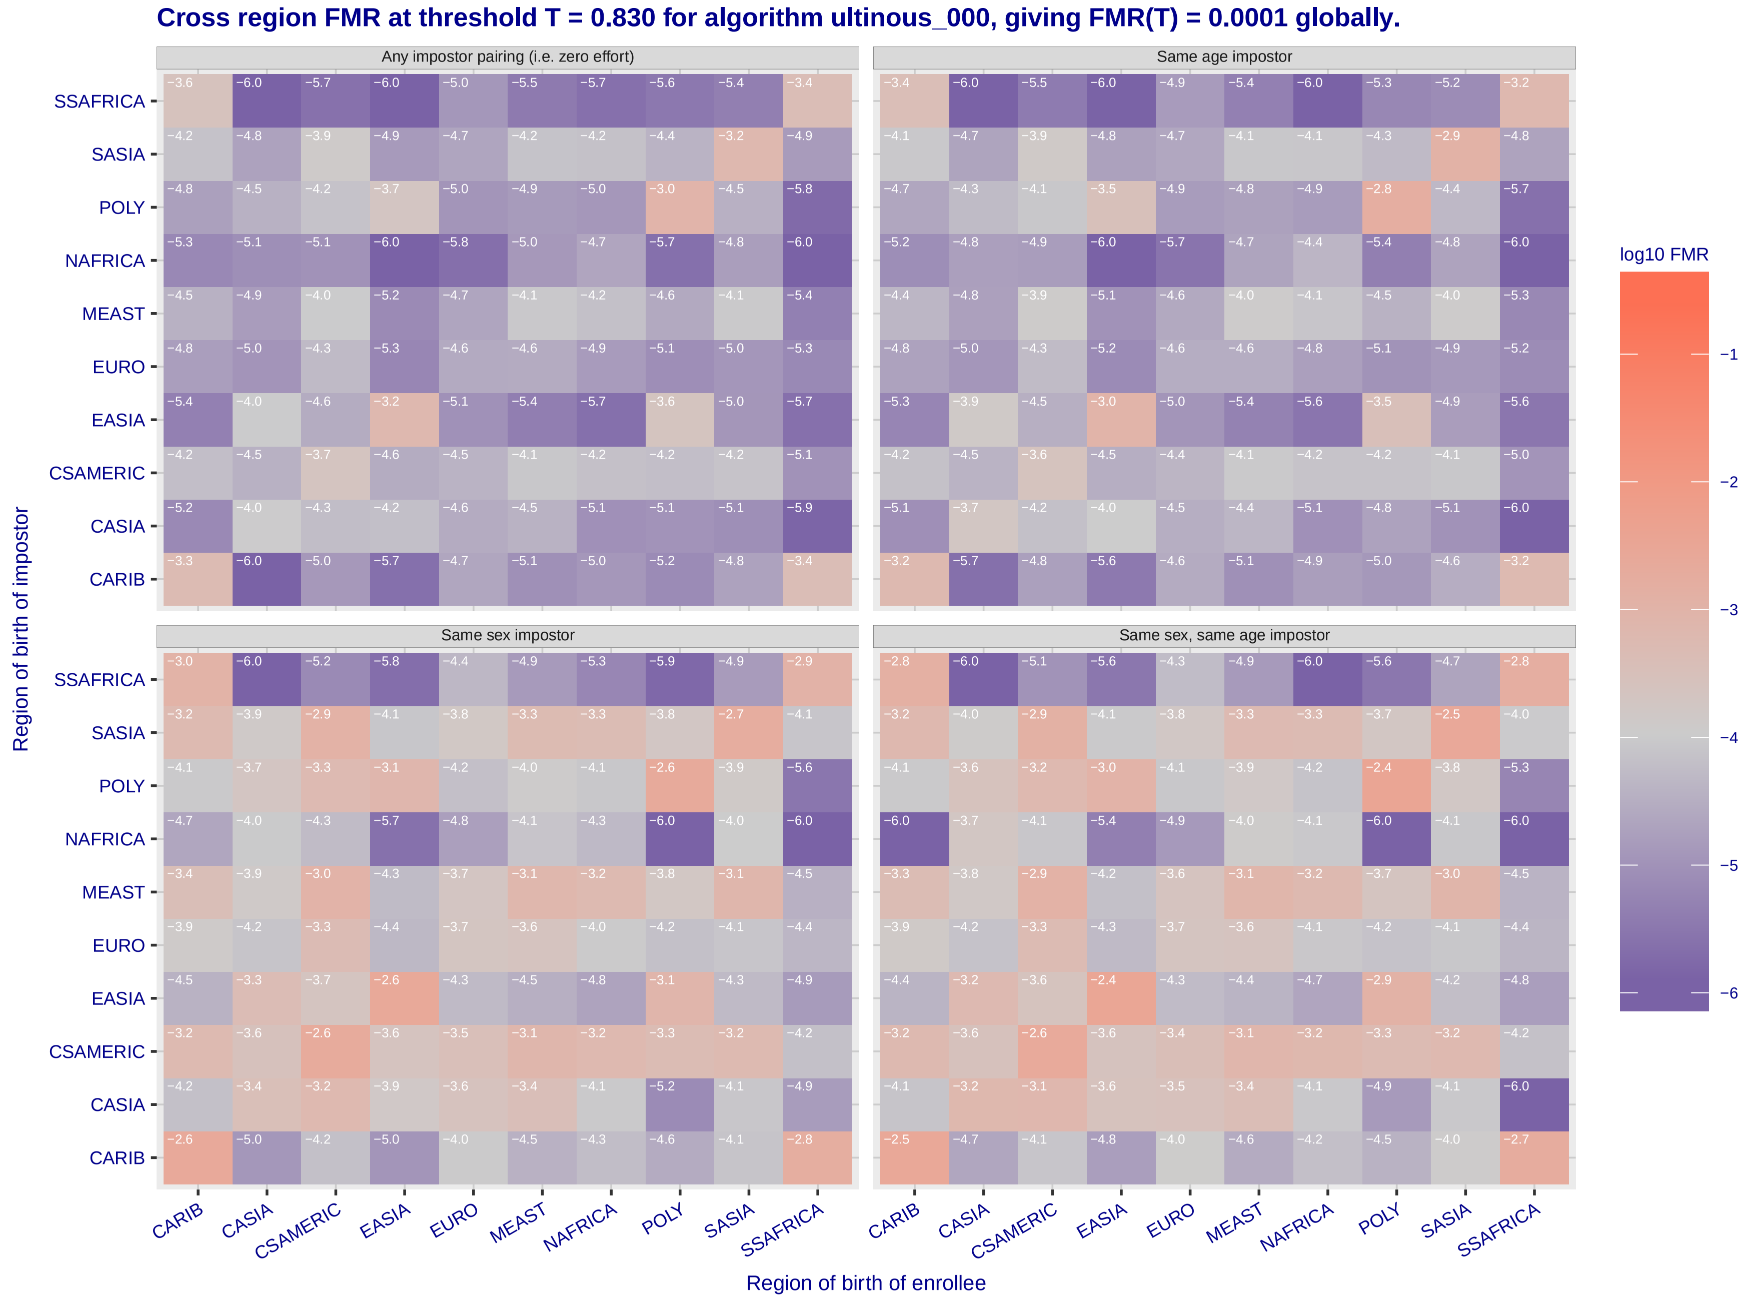

Figure 69: For algorithm ultinous-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

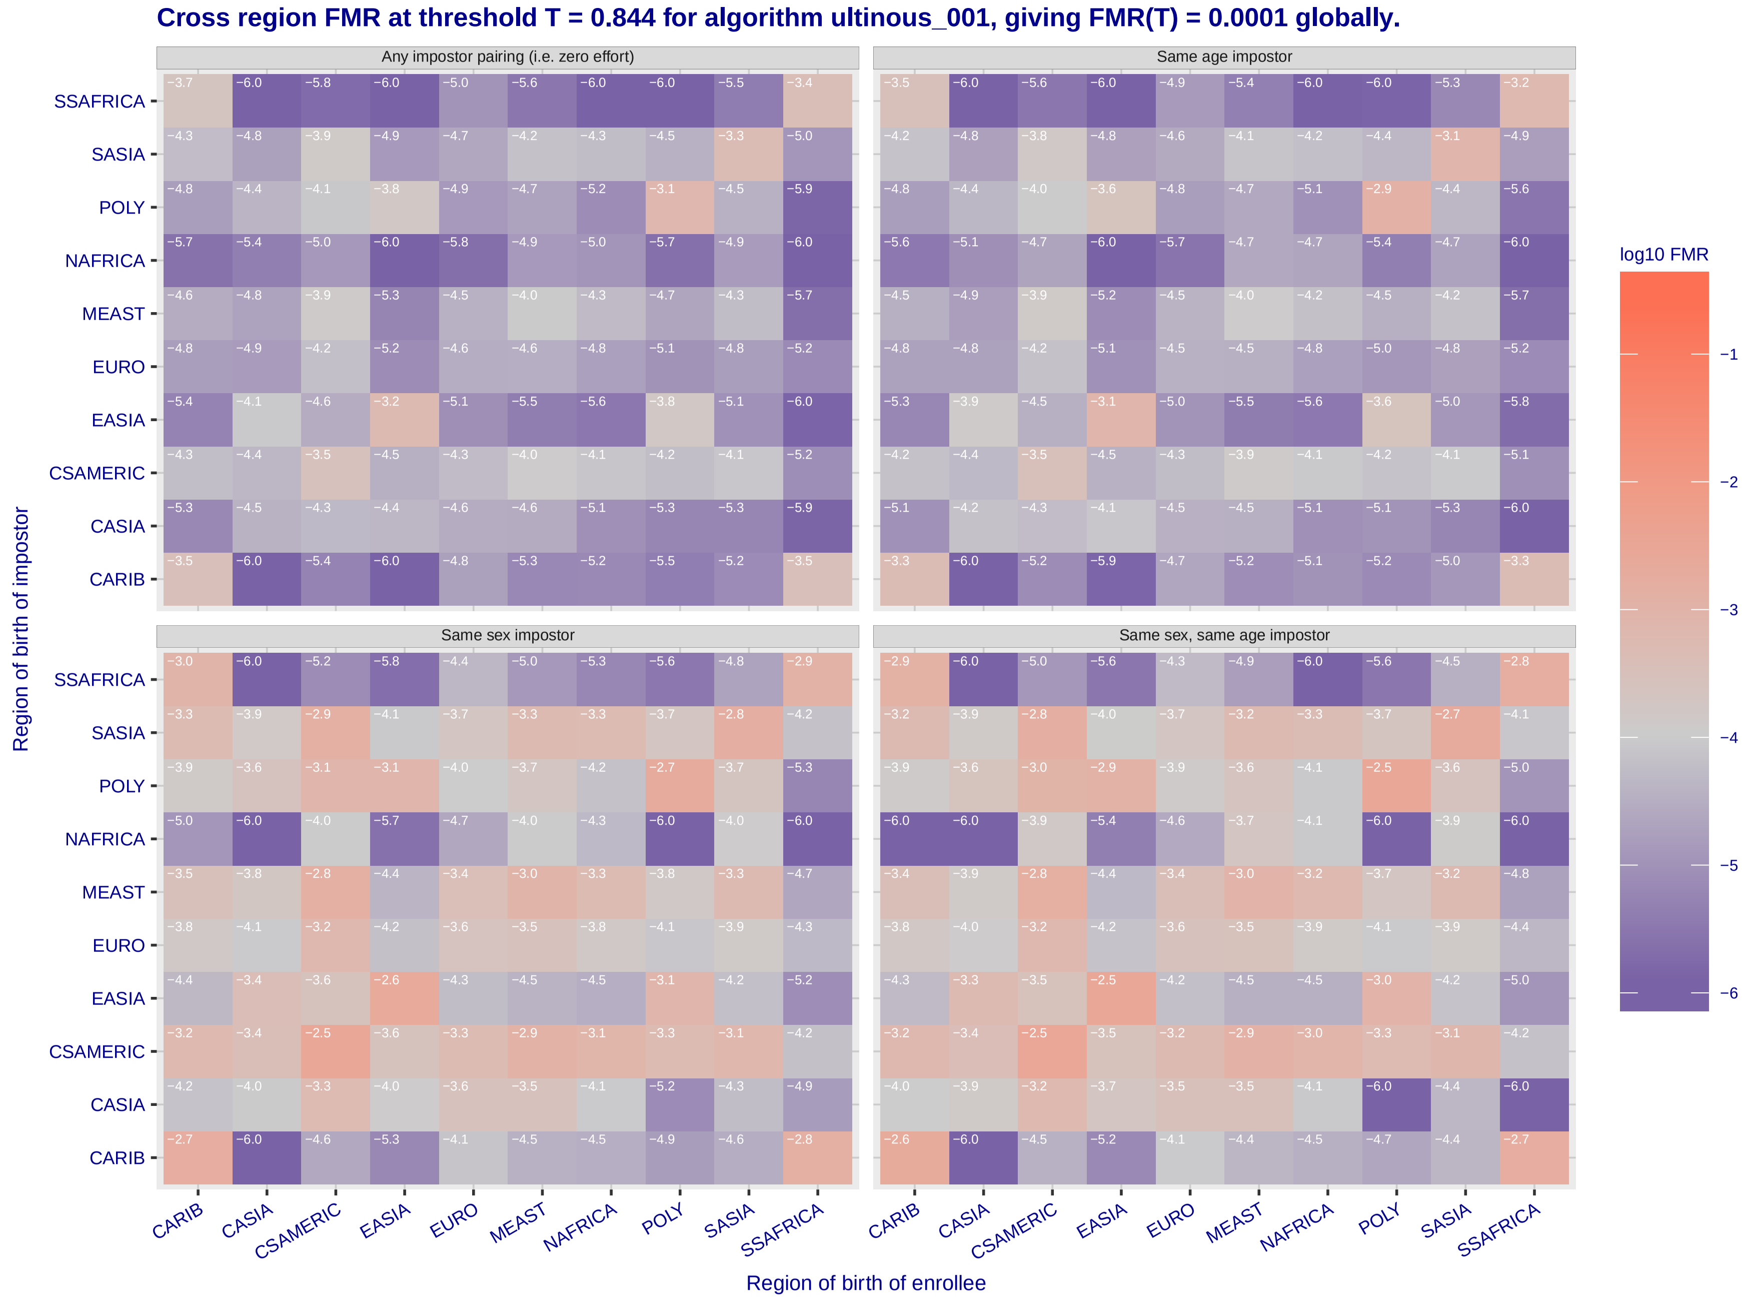

Figure 70: For algorithm ultinous-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

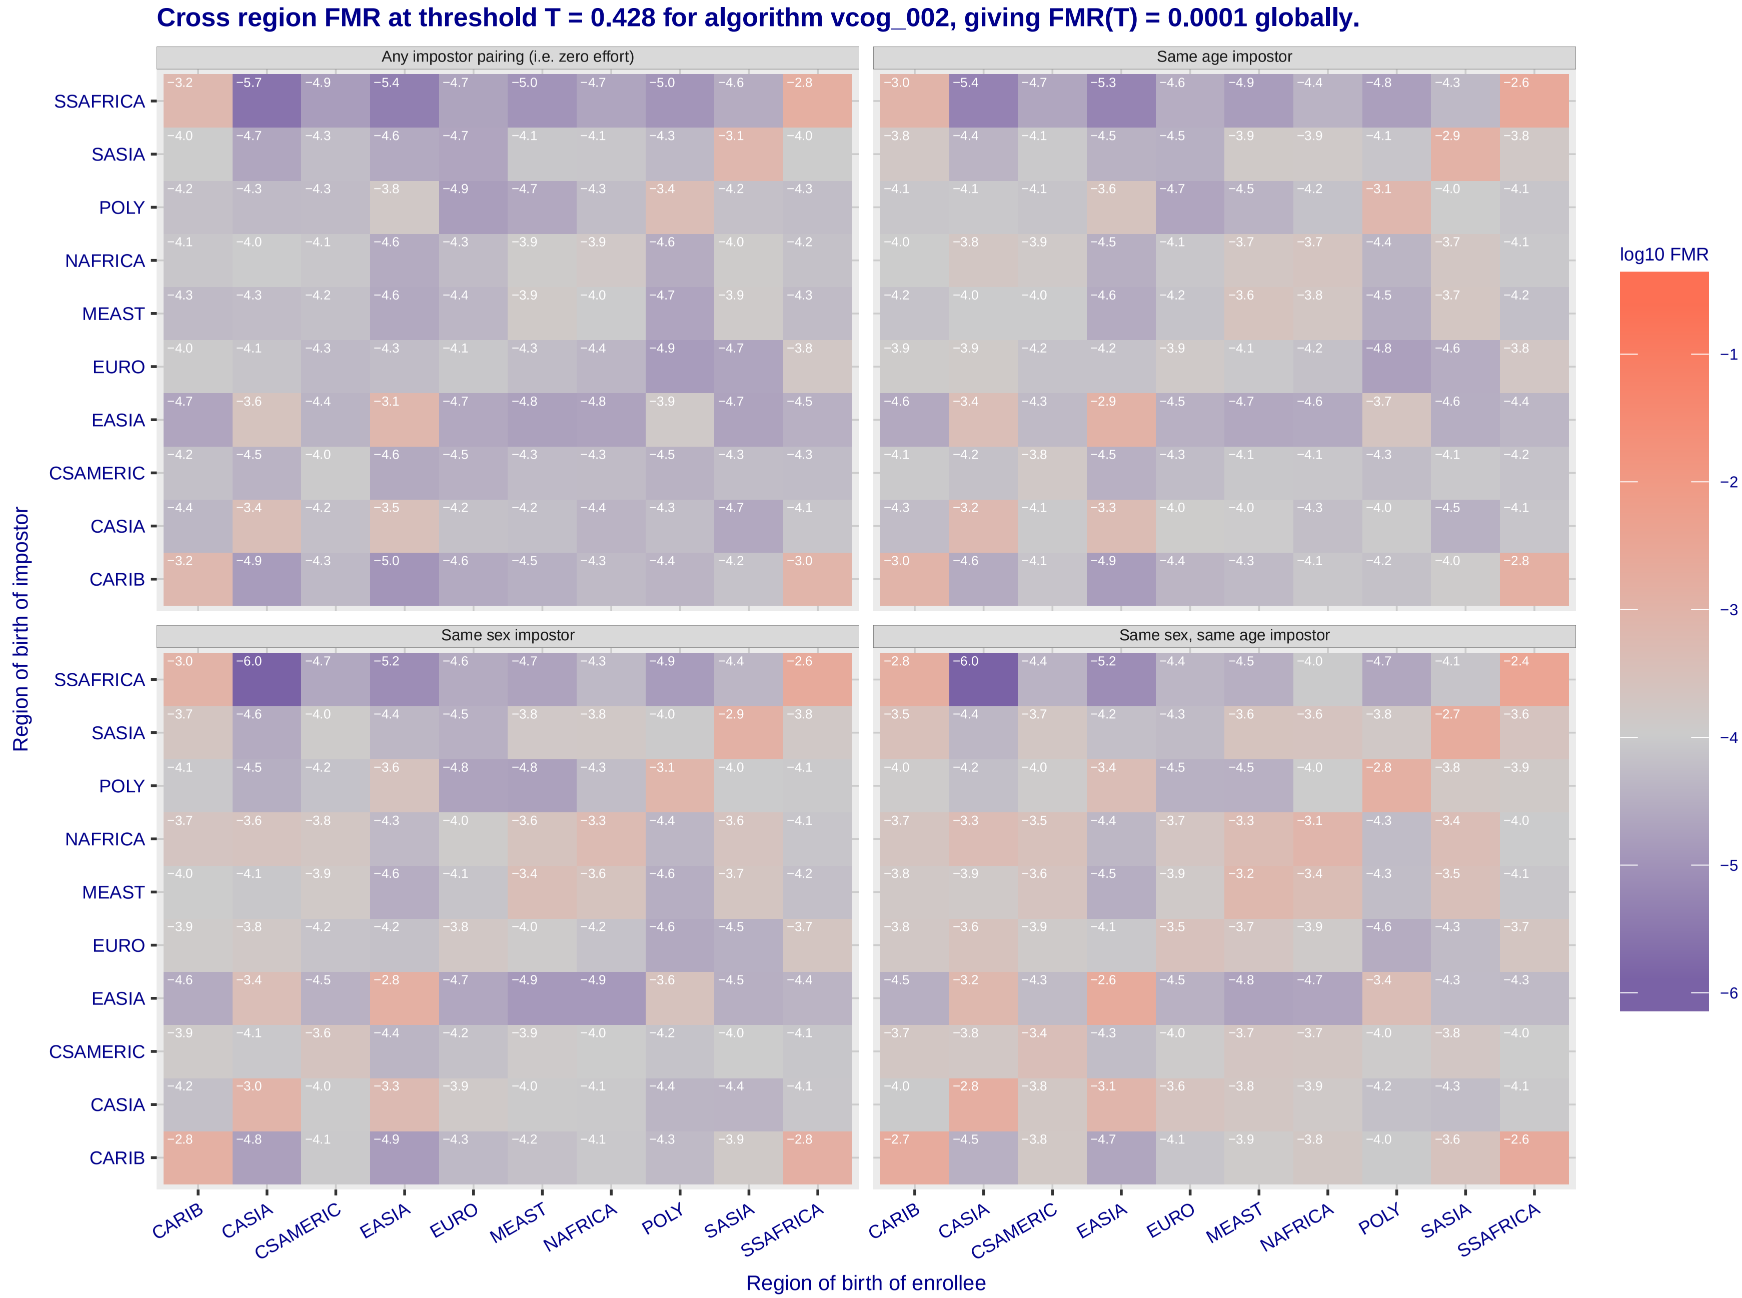

Figure 71: For algorithm vcog-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

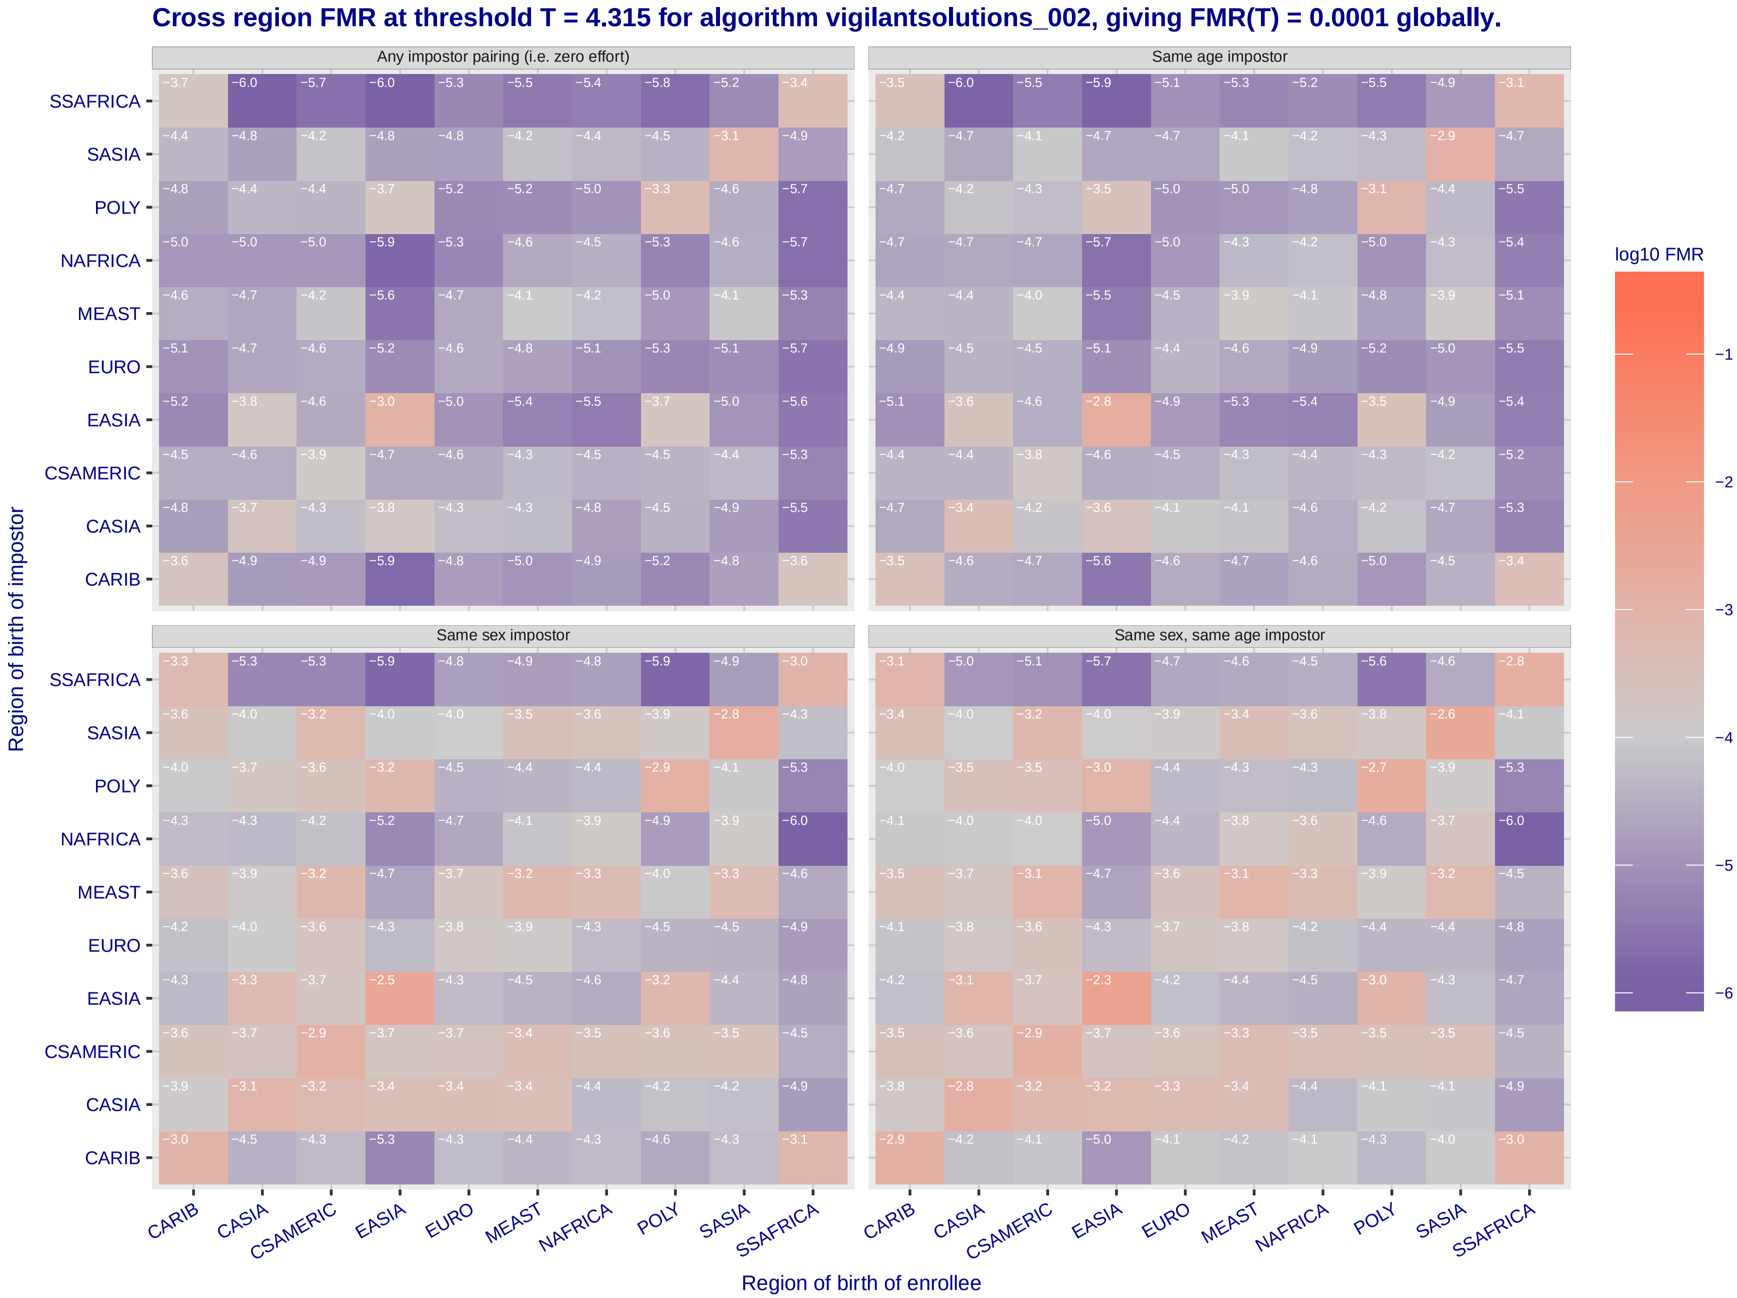

Figure 72: For algorithm vigilantsolutions-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

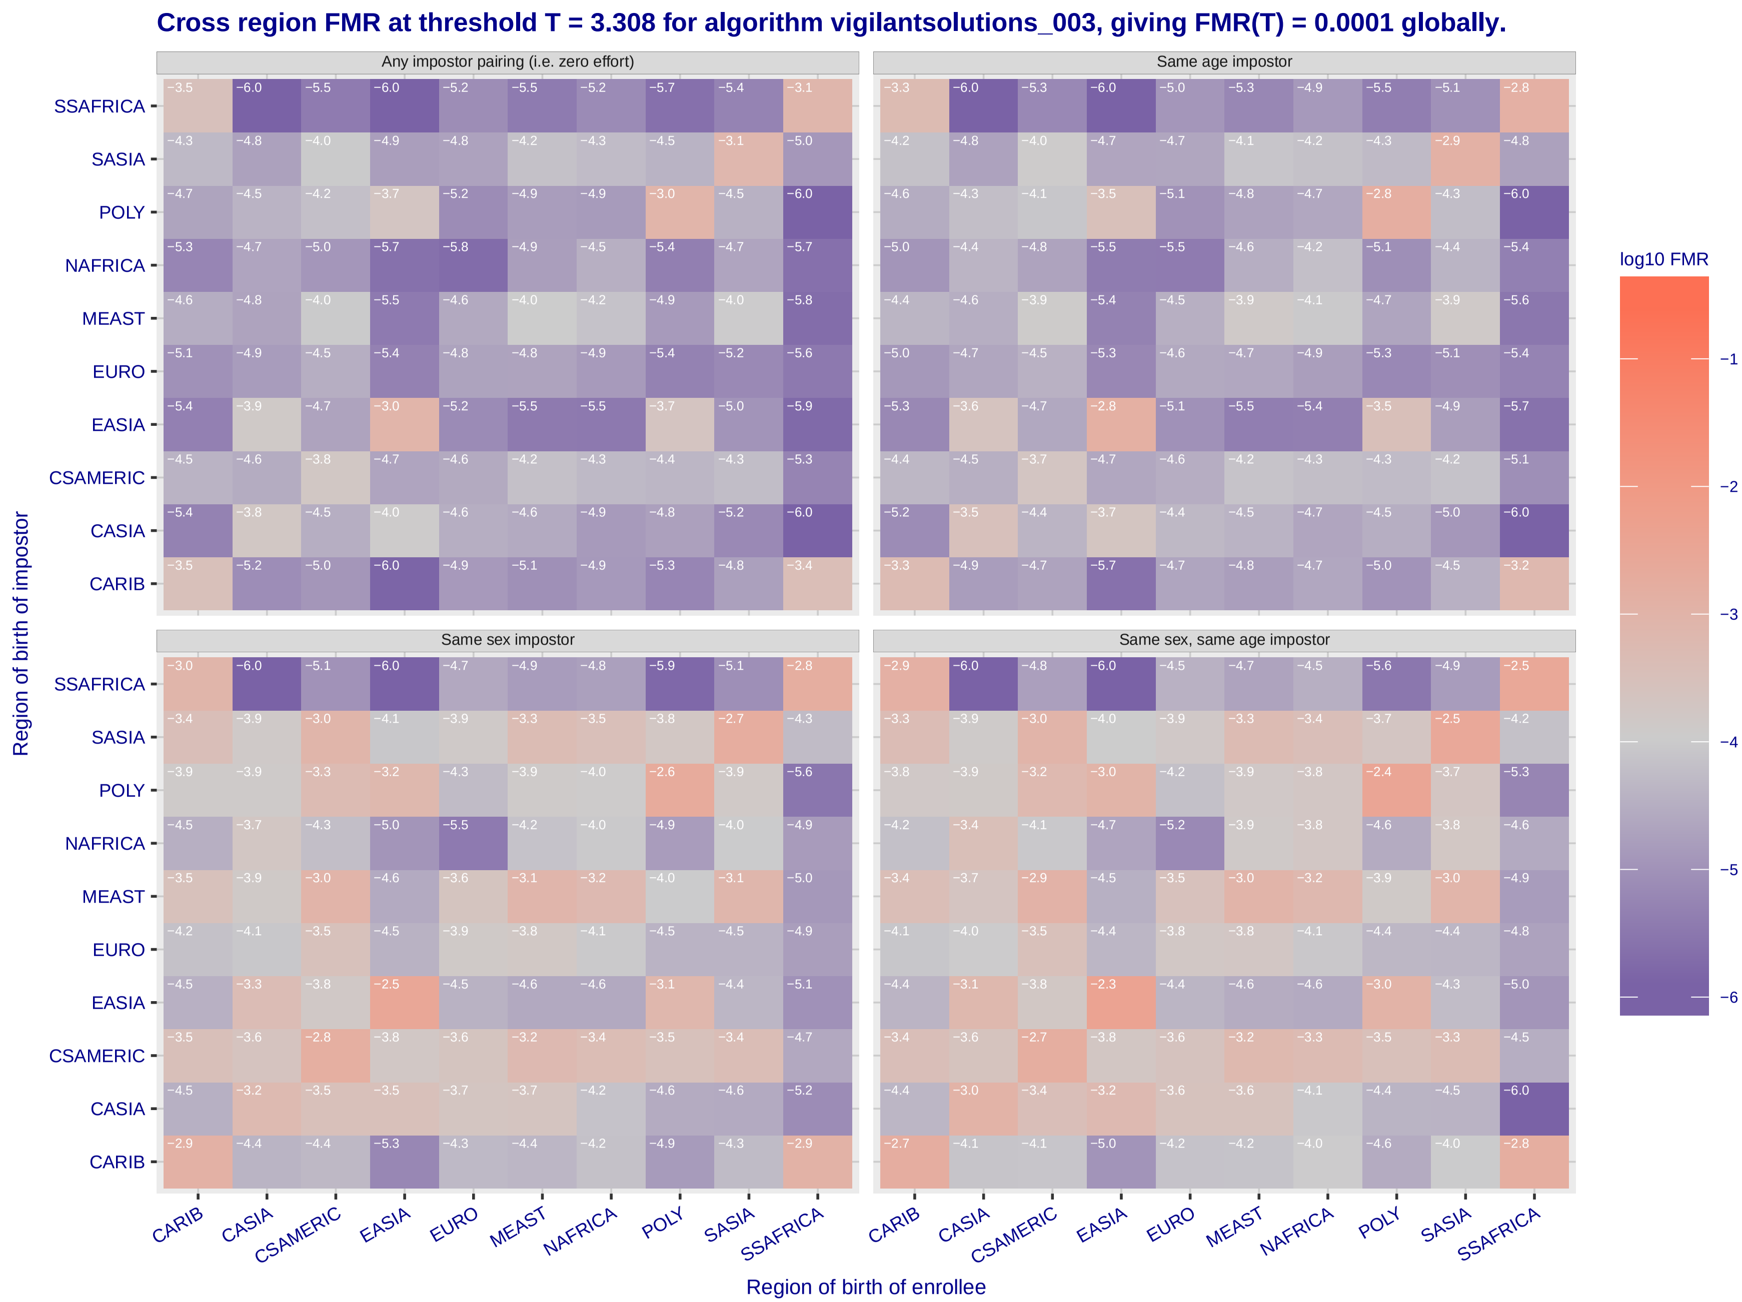

Figure 73: For algorithm vigilantsolutions-003 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

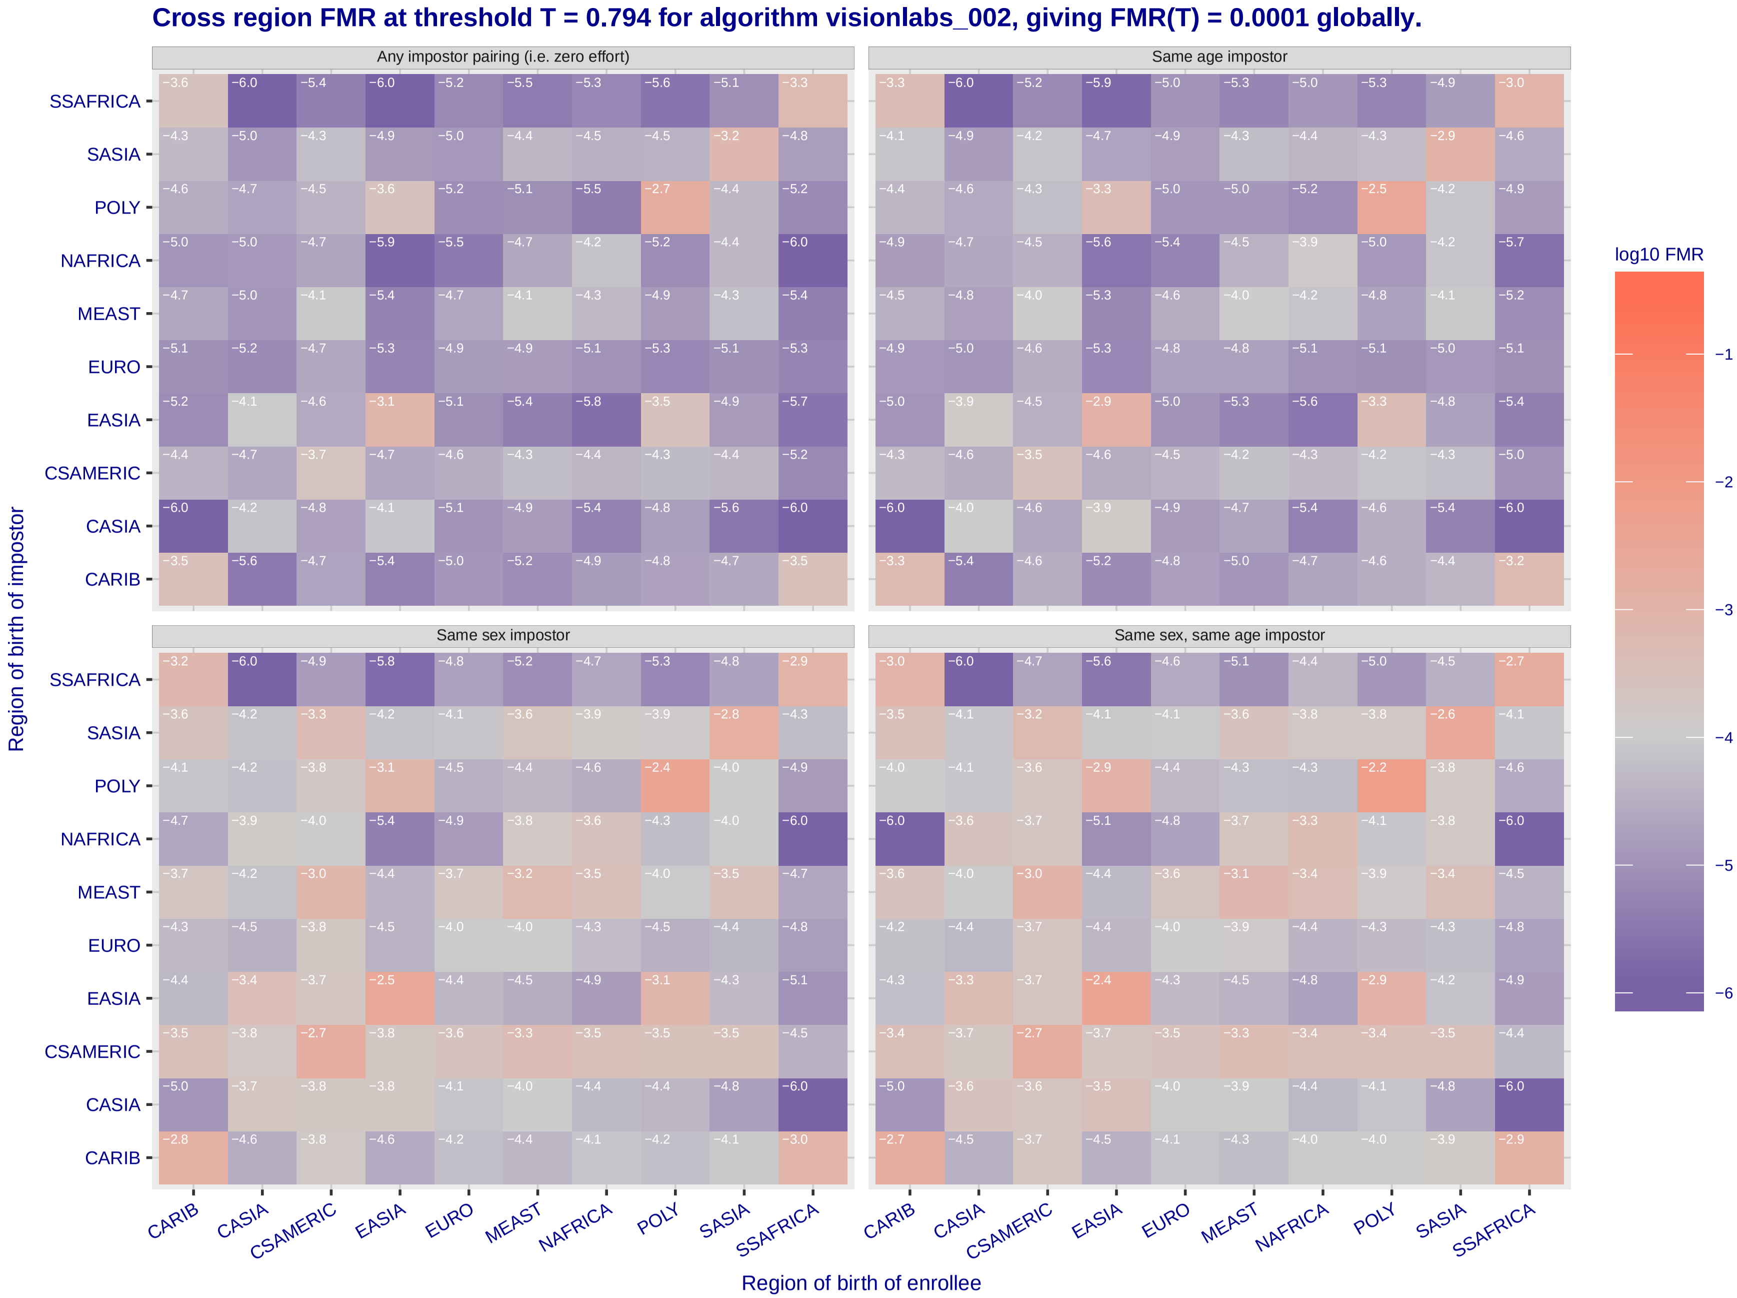

Figure 74: For algorithm visionlabs-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

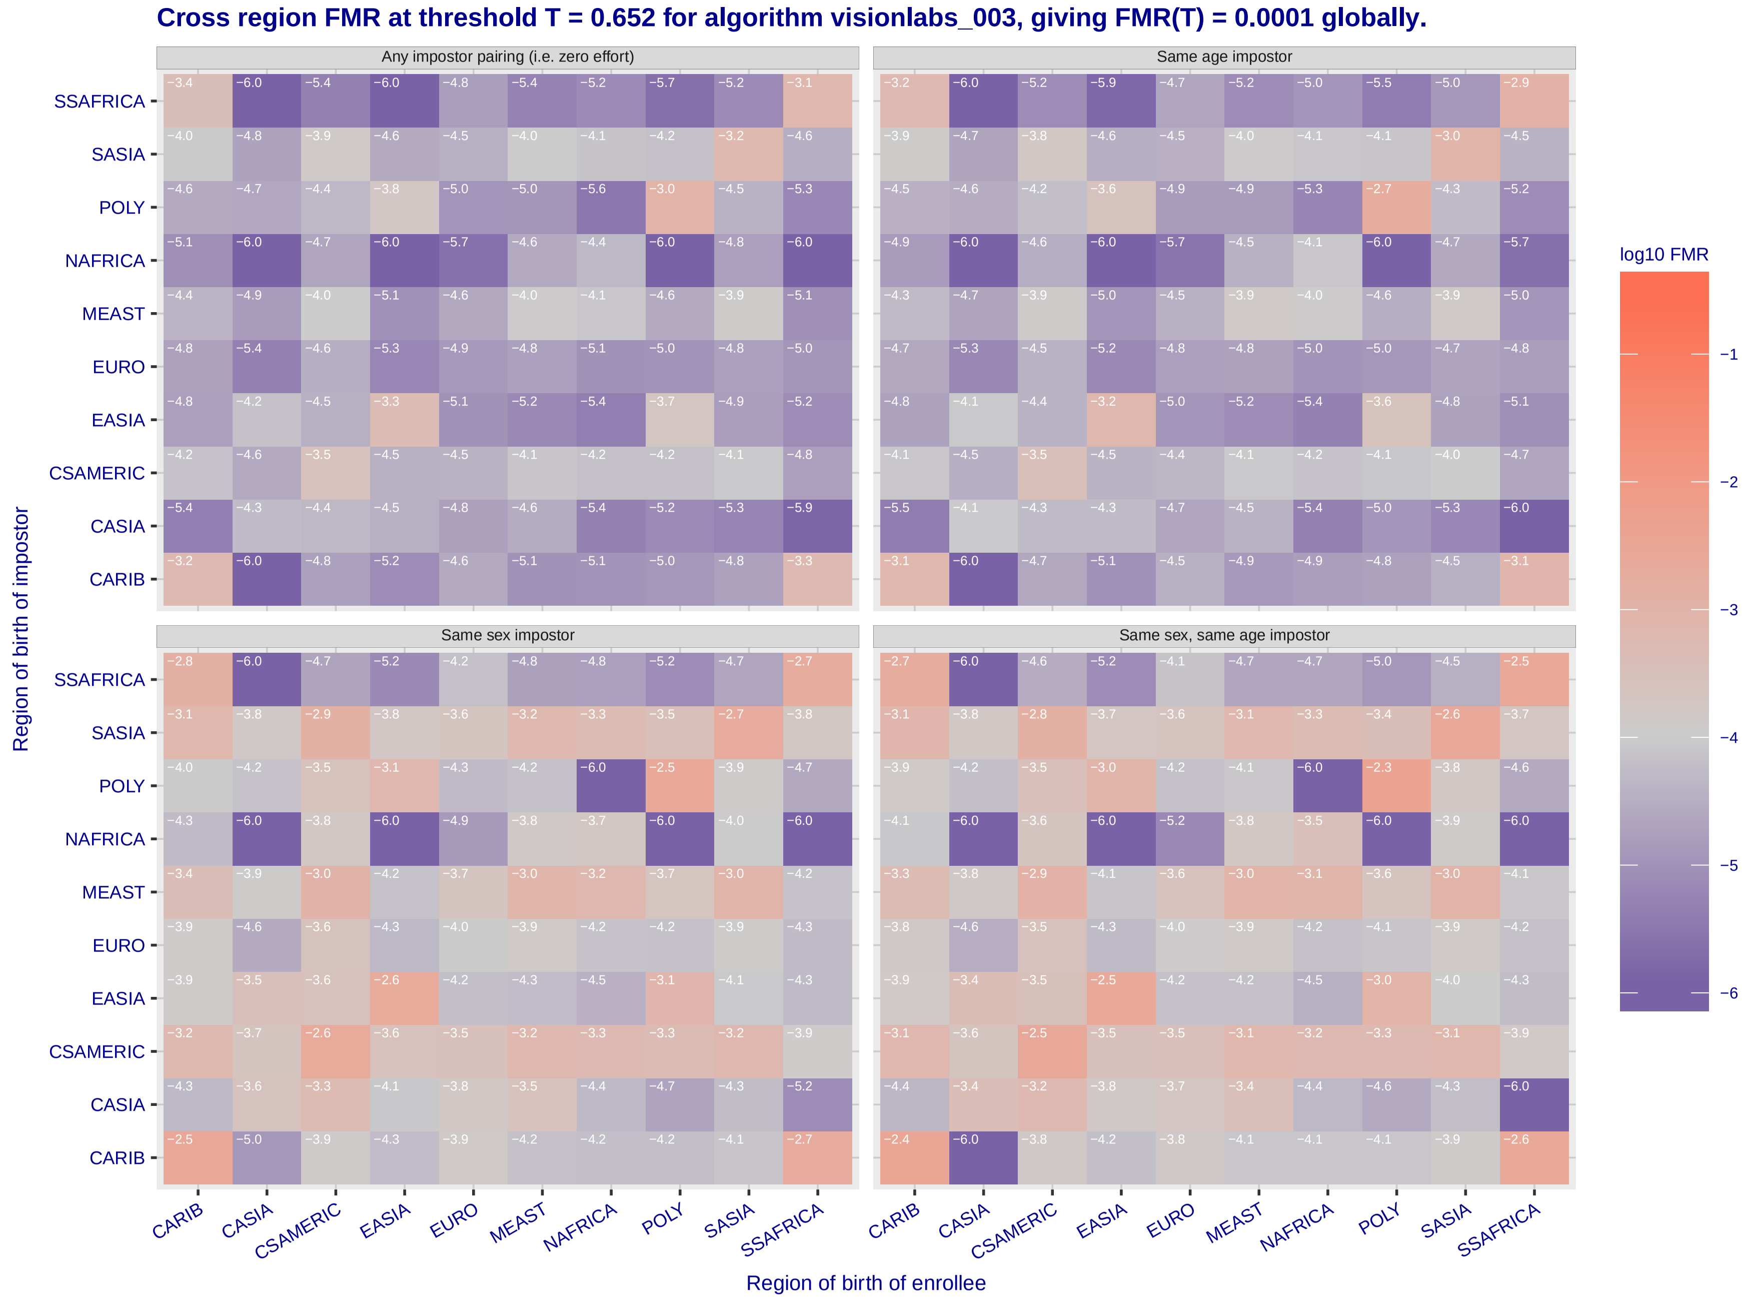

Figure 75: For algorithm visionlabs-003 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

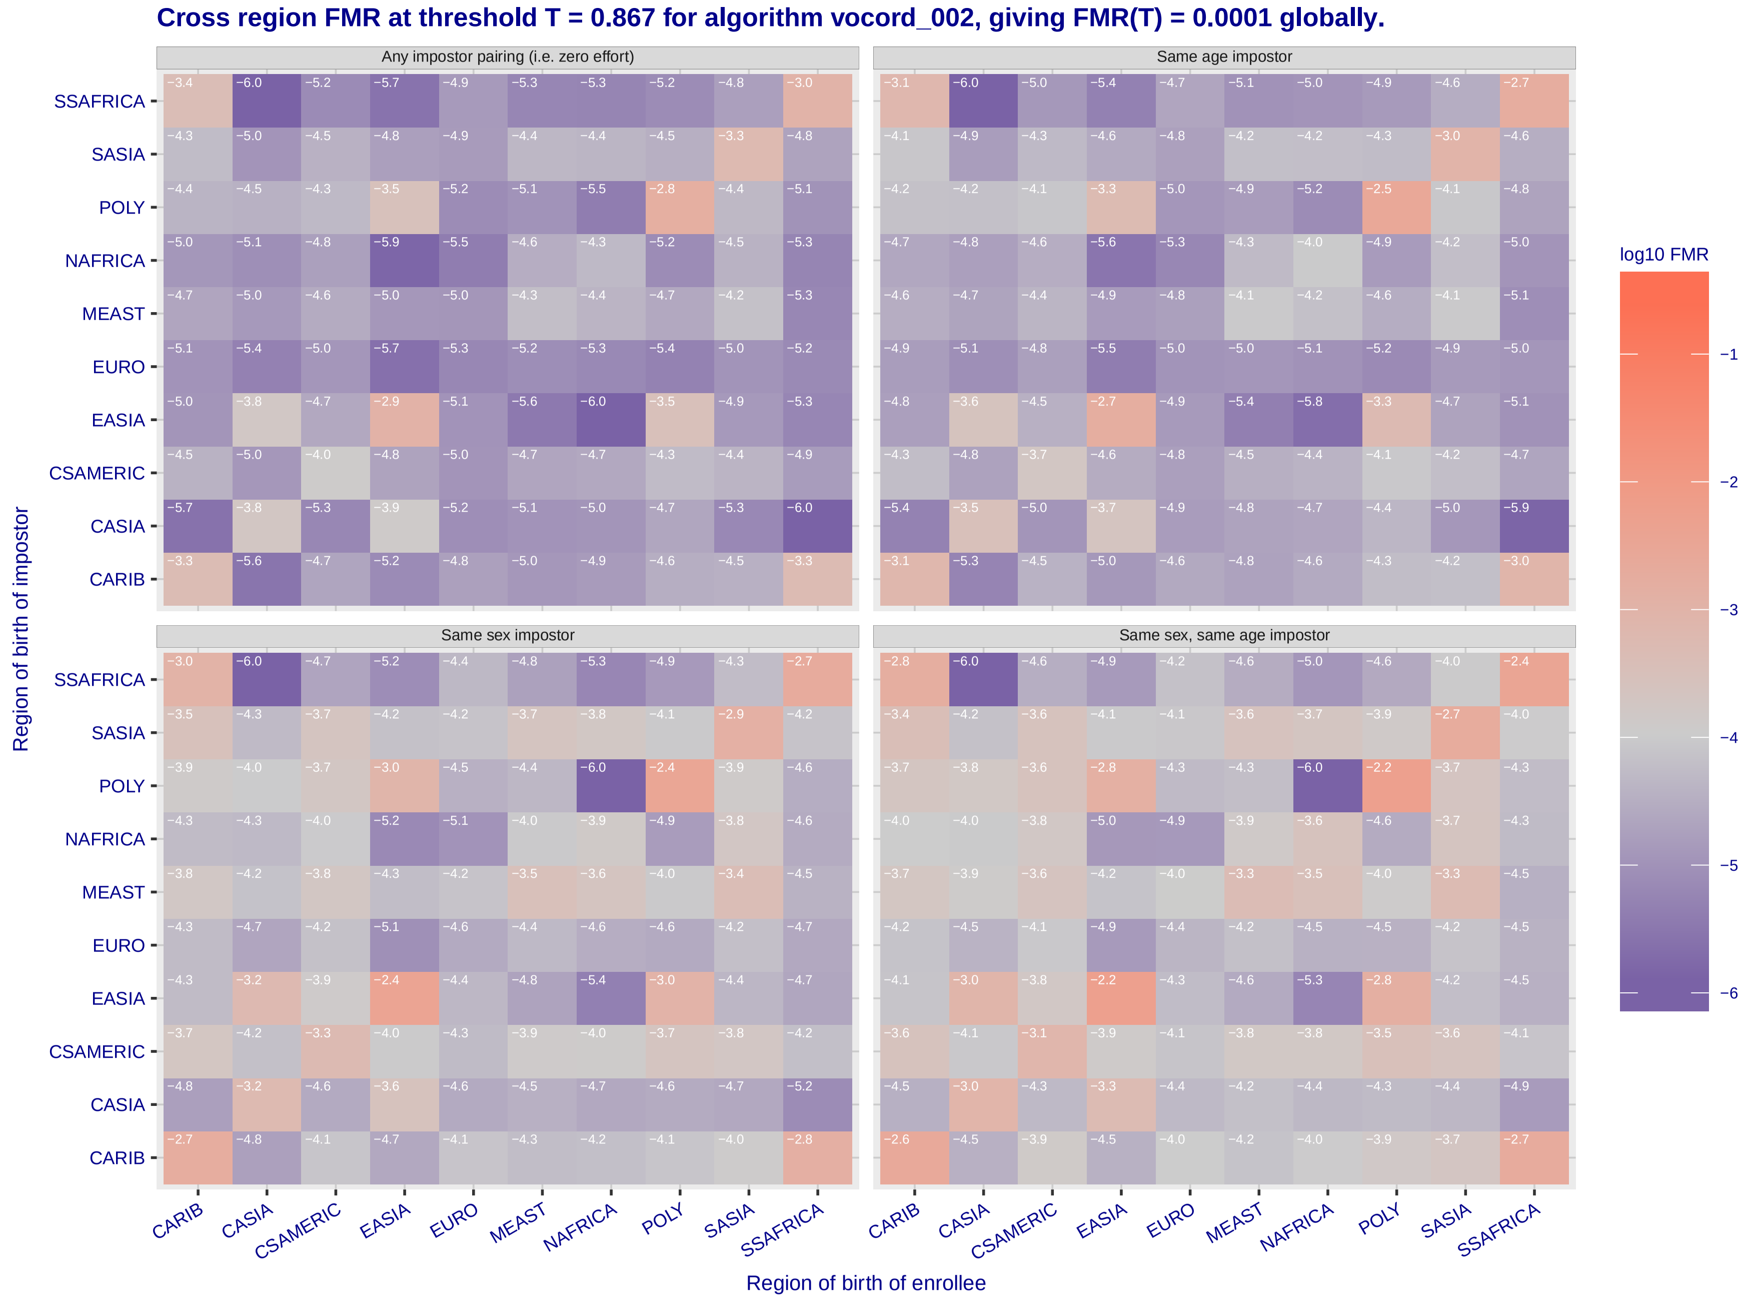

Figure 76: For algorithm vocord-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

Figure 77: For algorithm vocord-003 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

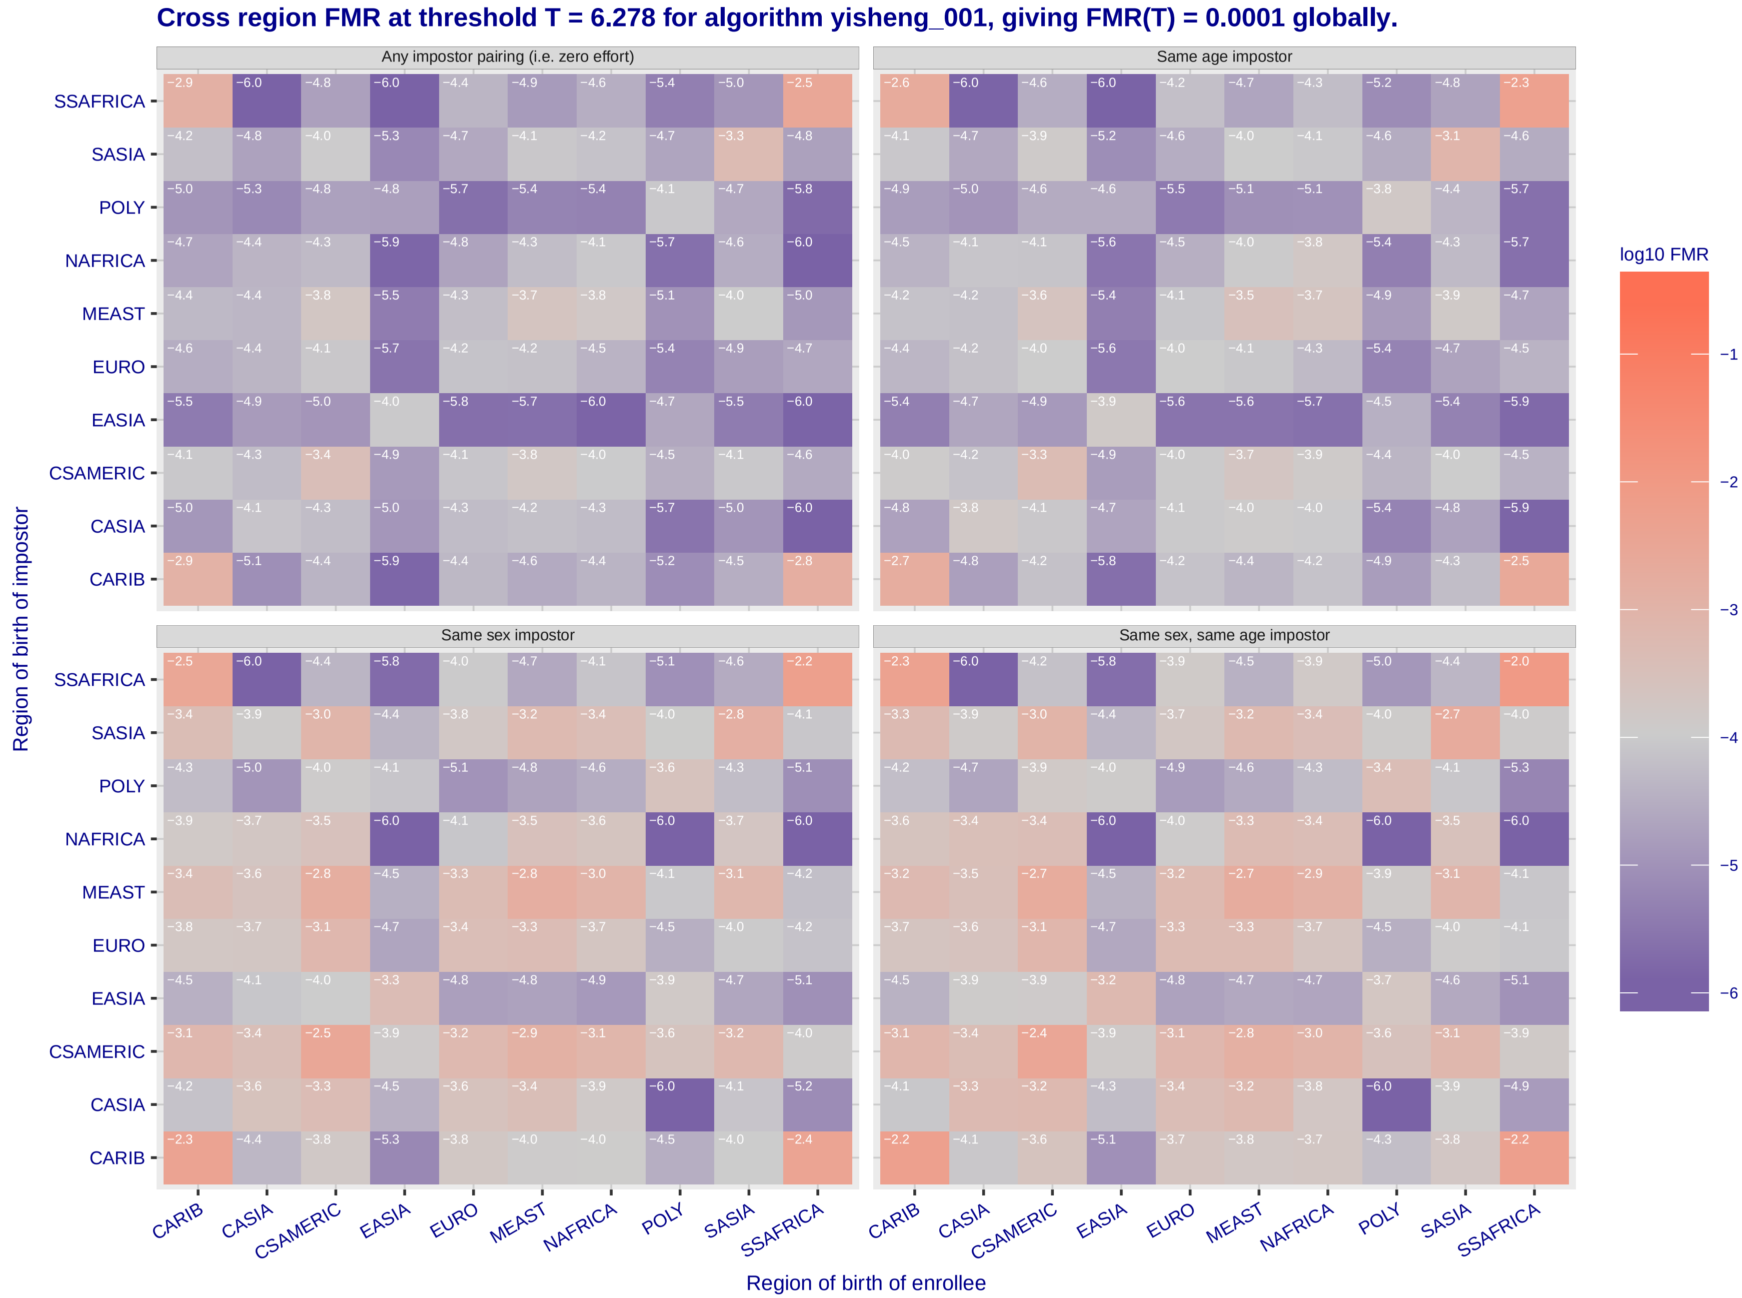

Figure 78: For algorithm yisheng-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

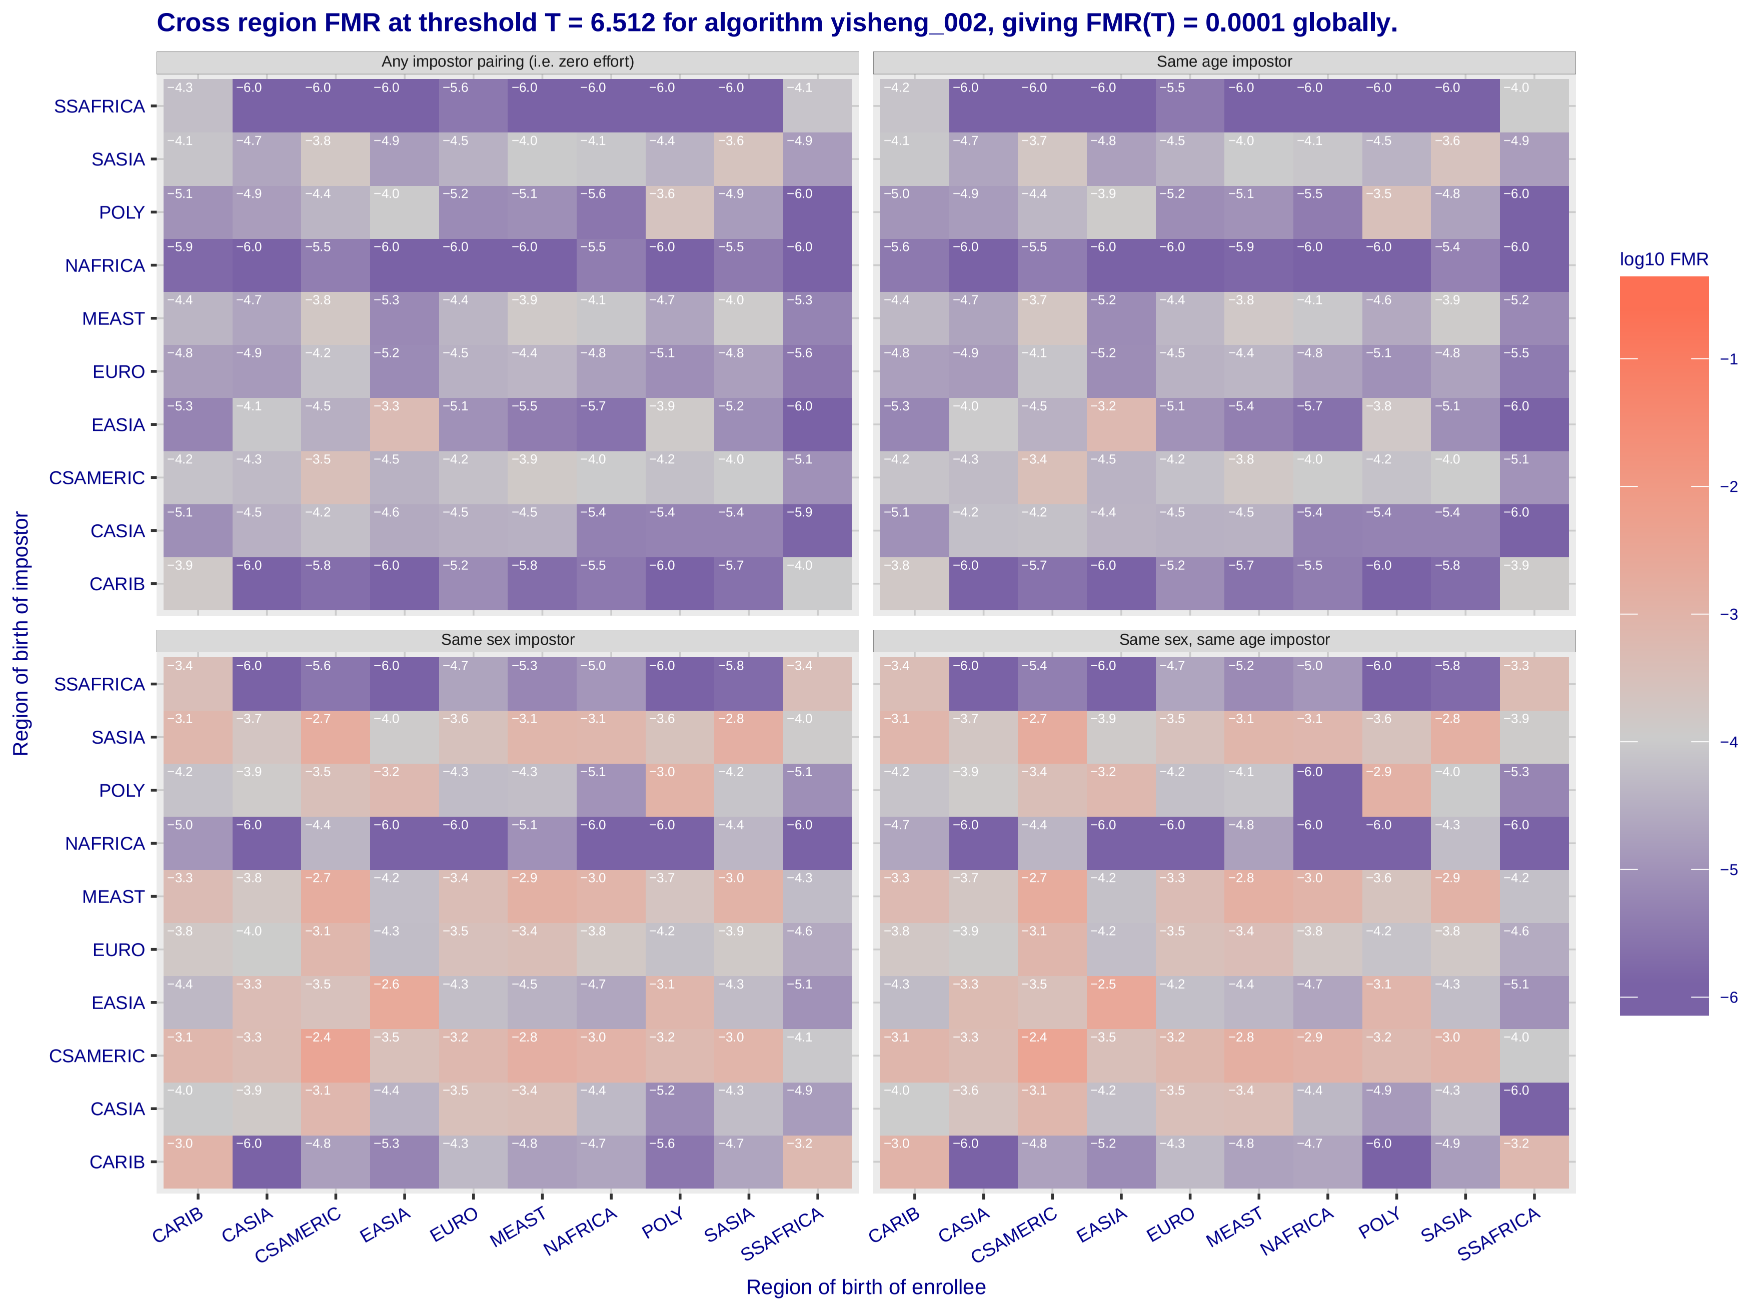

Figure 79: For algorithm yisheng-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

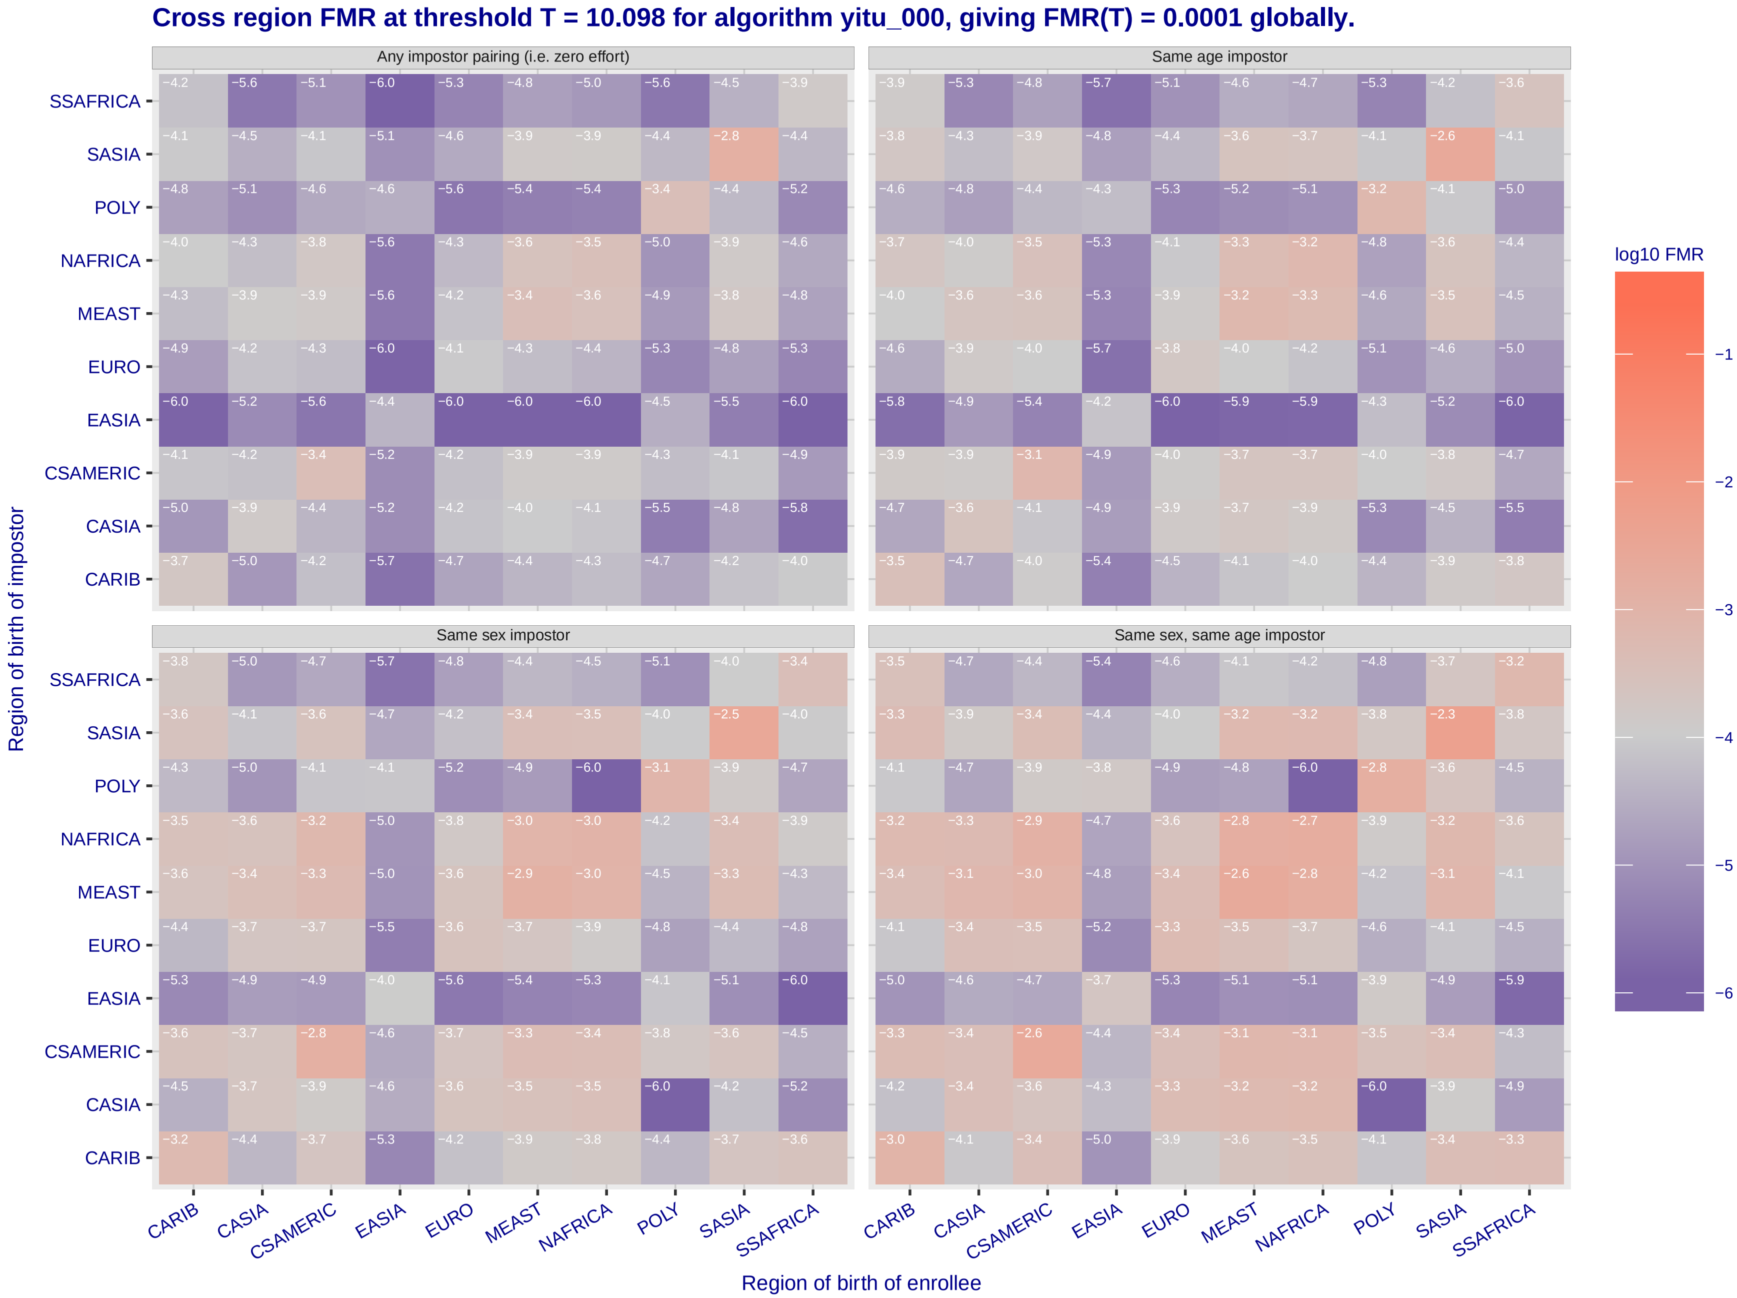

Figure 80: For algorithm yitu-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given region pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

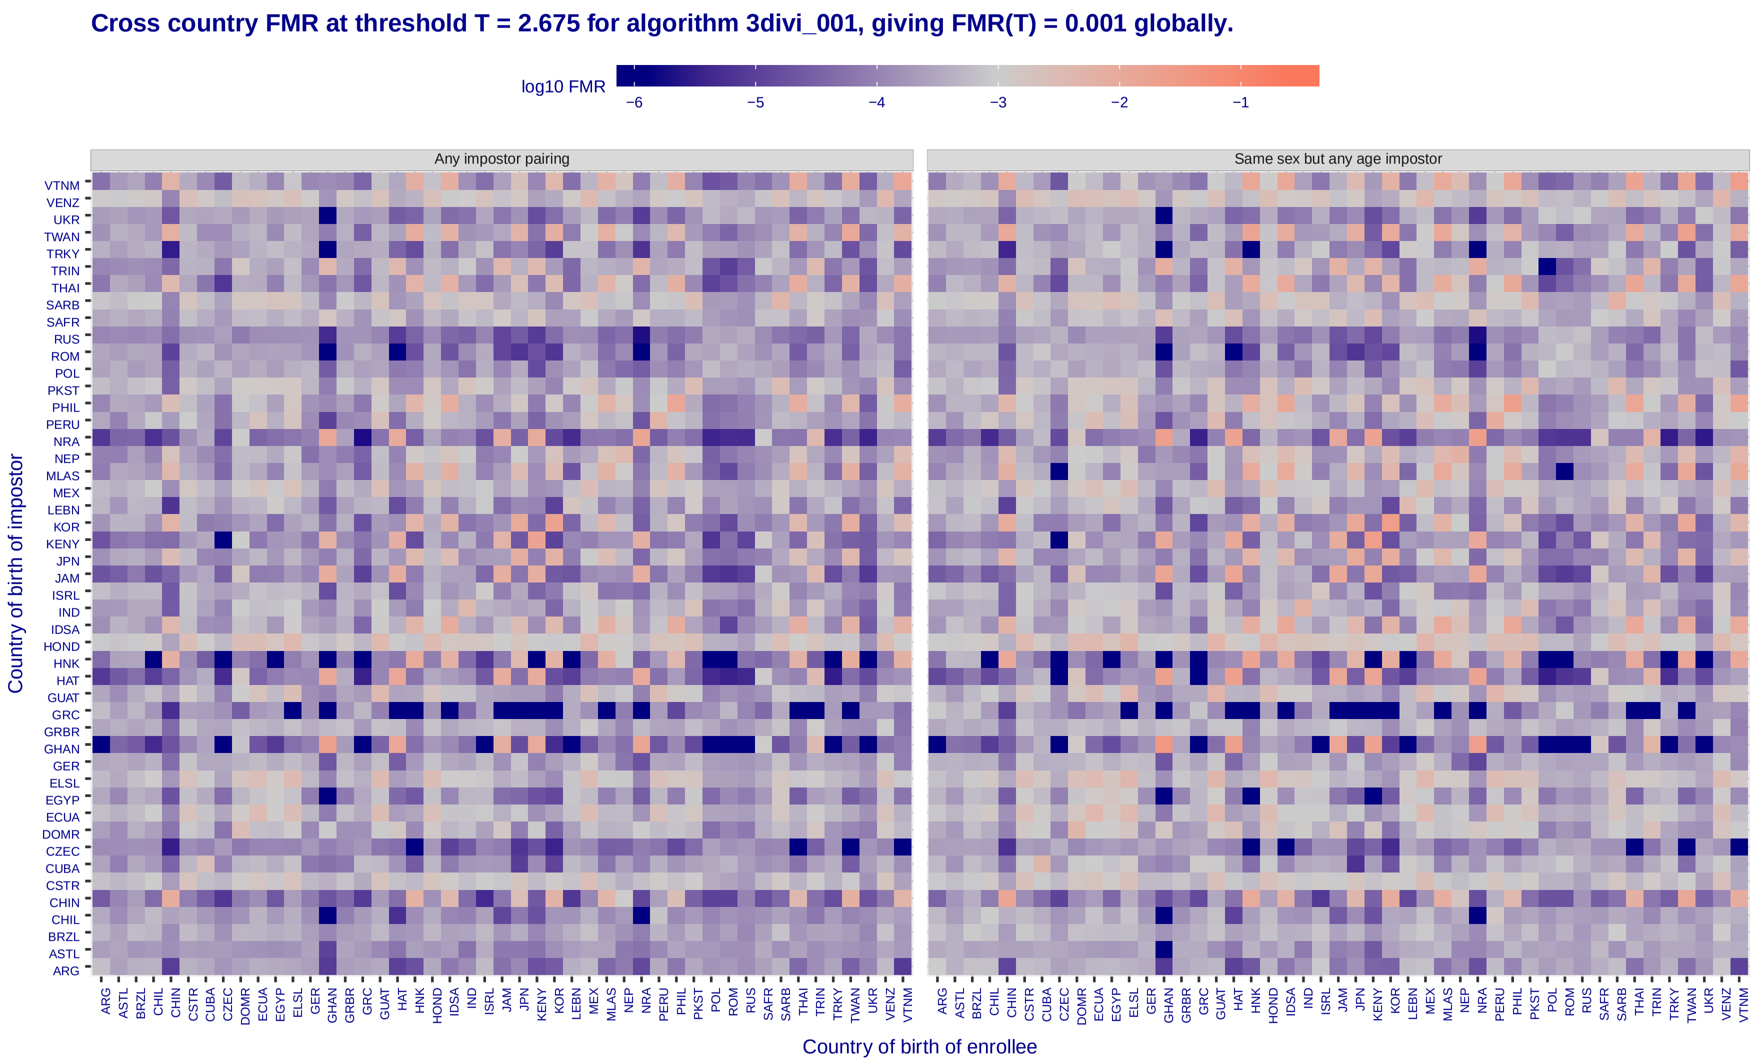

Figure 81: For algorithm 3divi-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

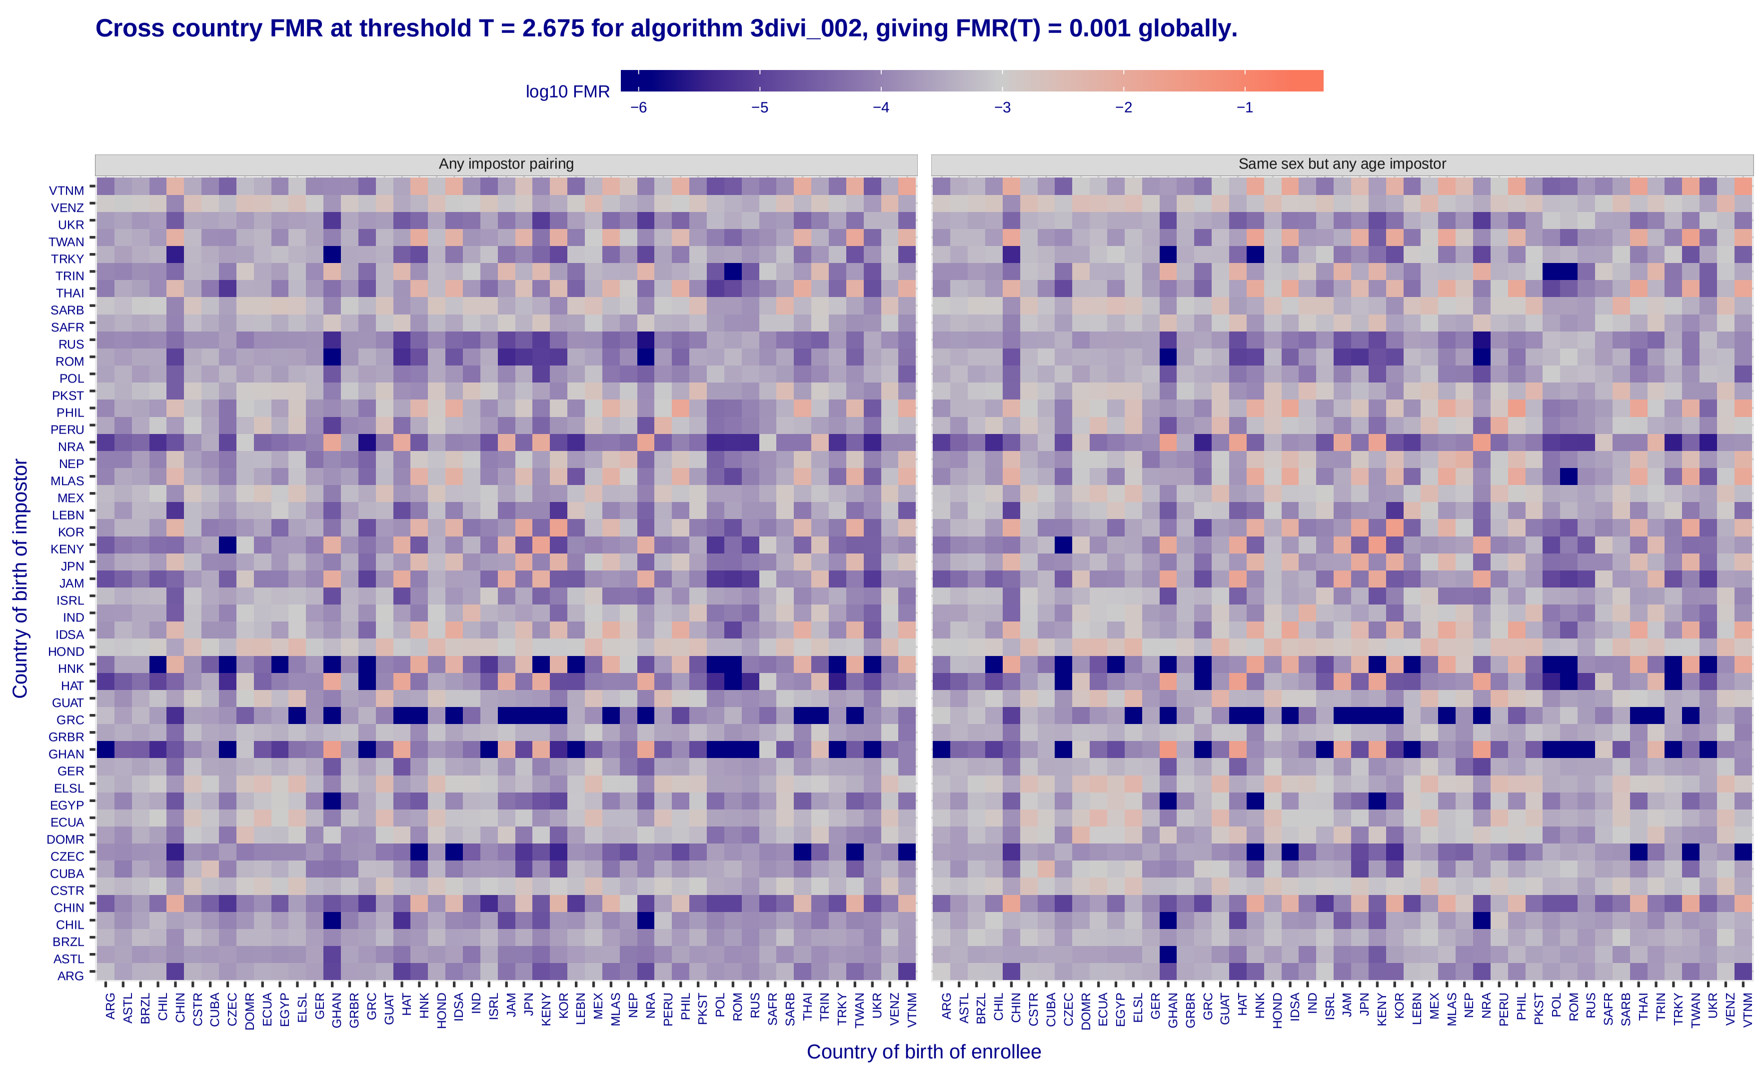

Figure 82: For algorithm 3divi-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

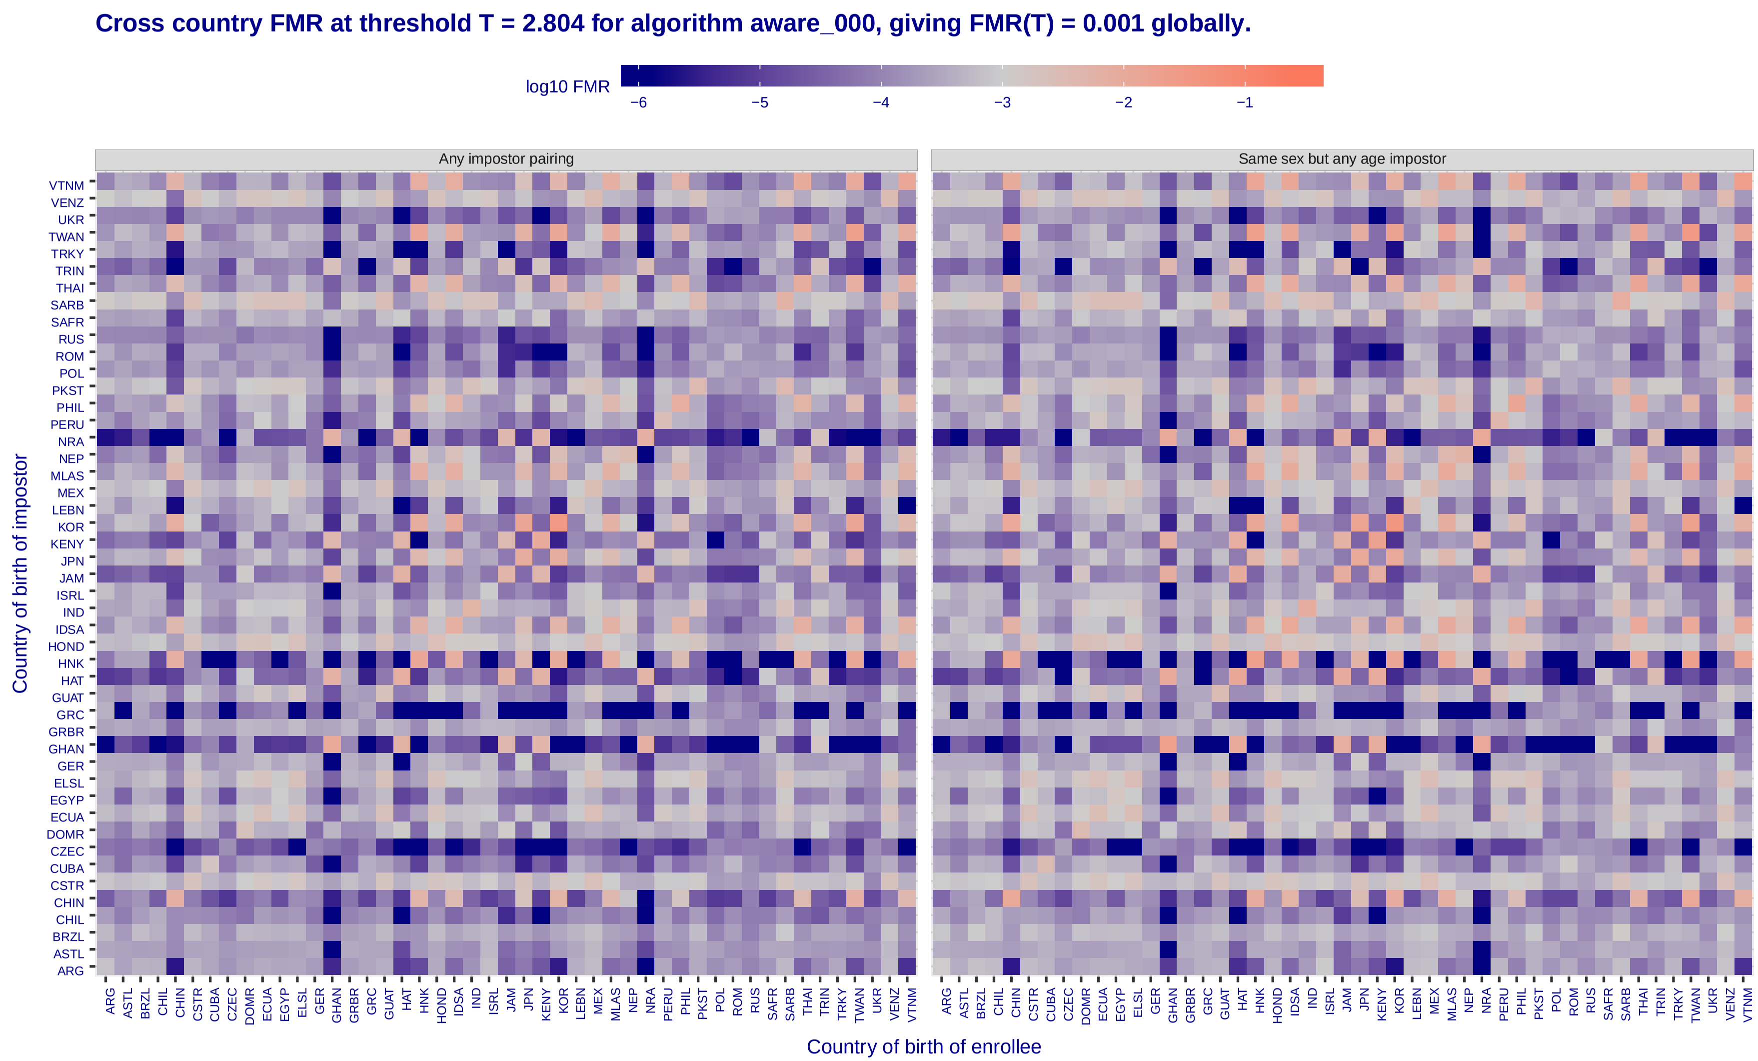

Figure 83: For algorithm aware-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

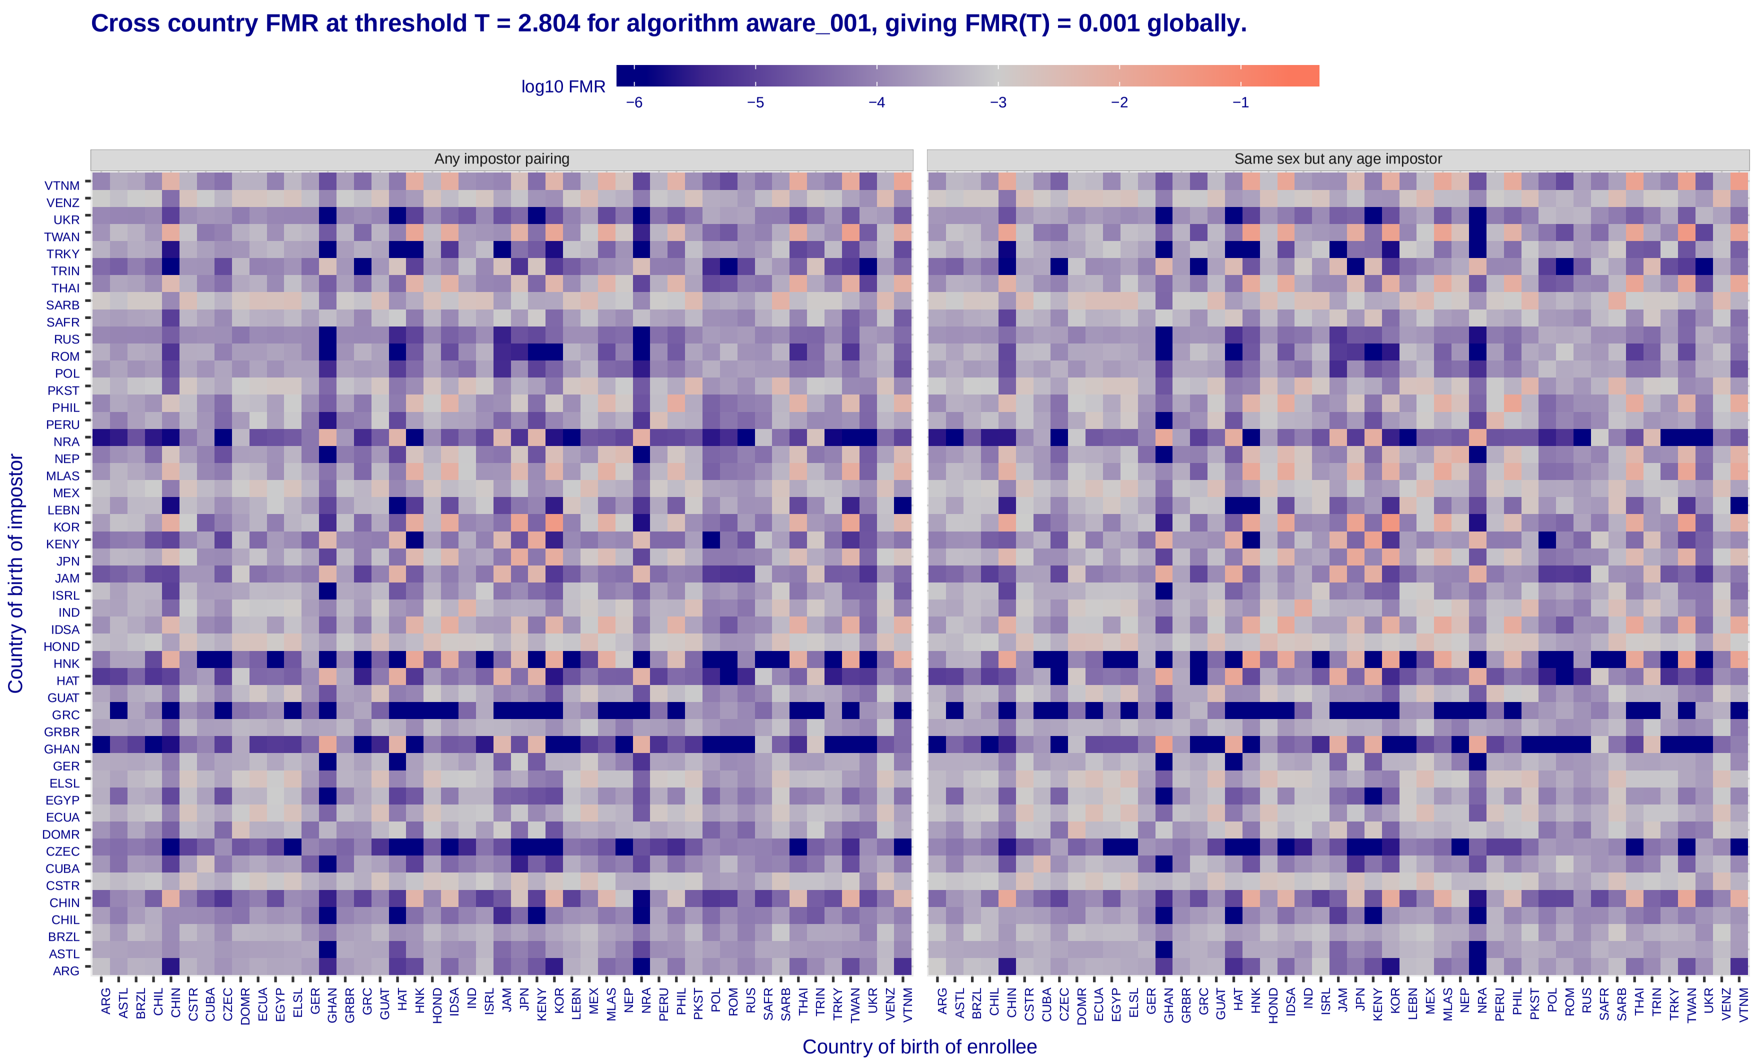

Figure 84: For algorithm aware-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

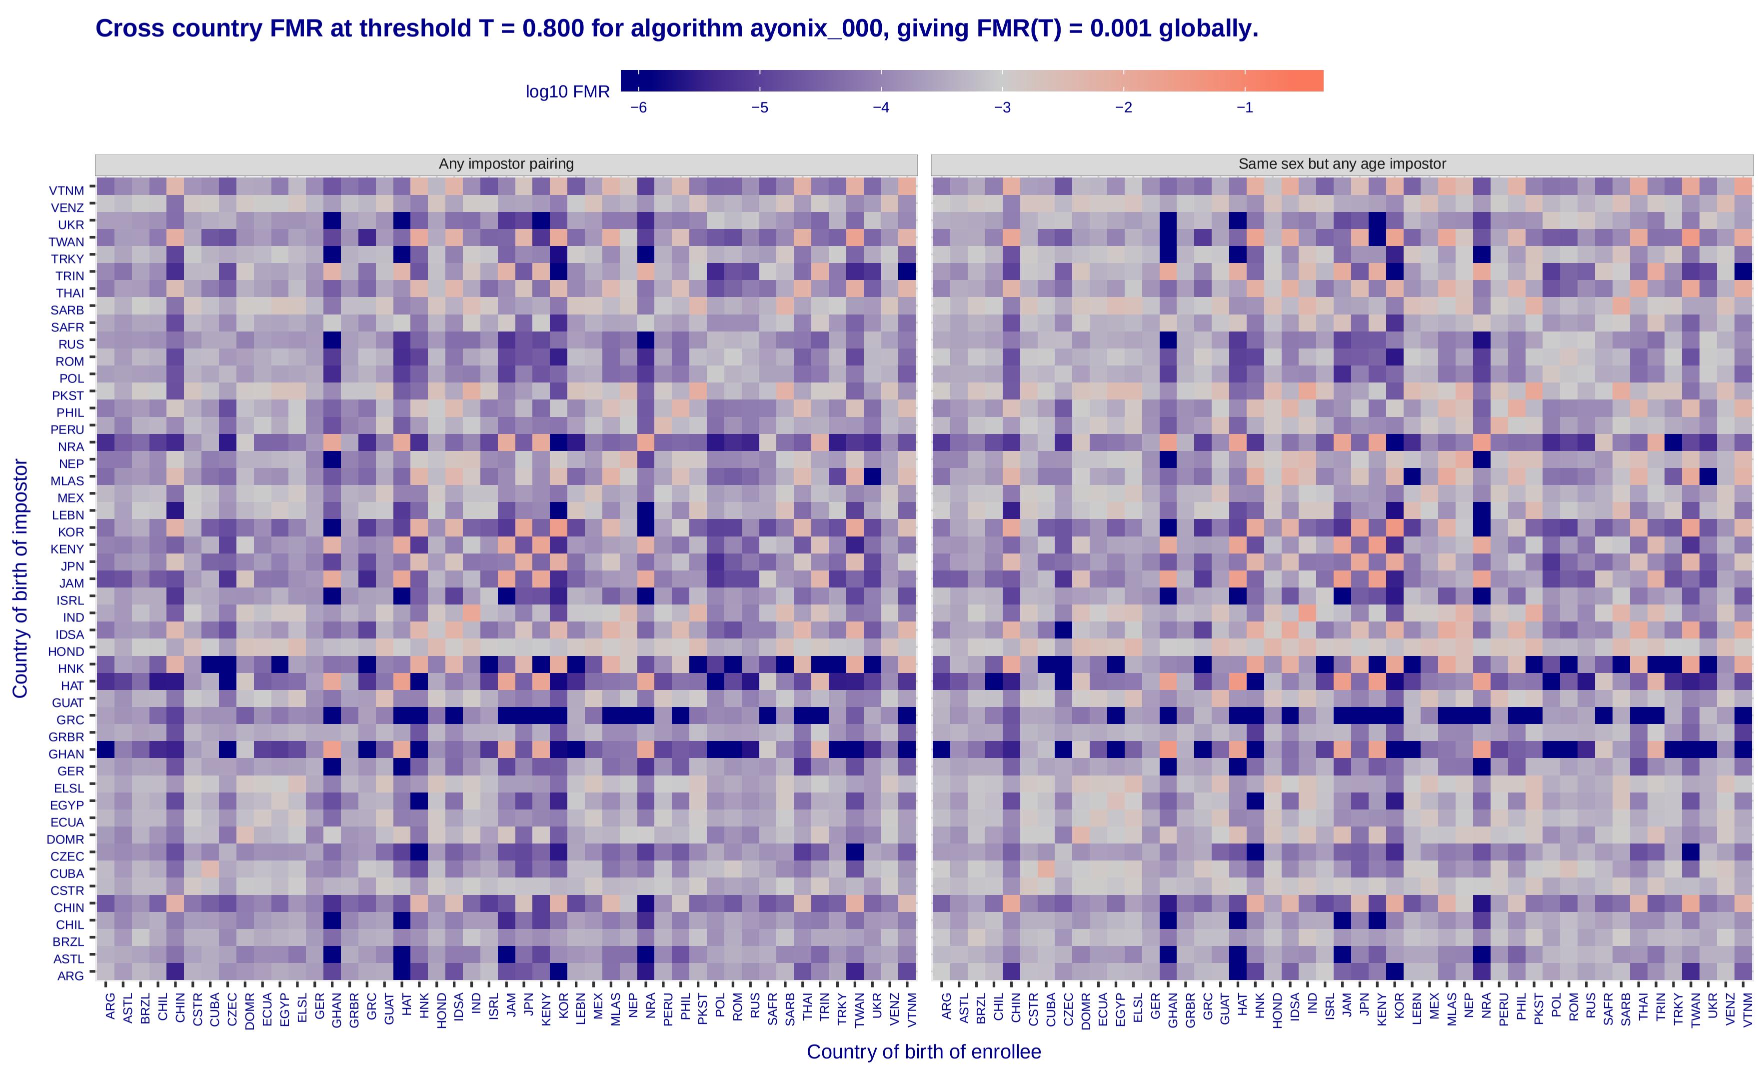

Figure 85: For algorithm ayonix-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

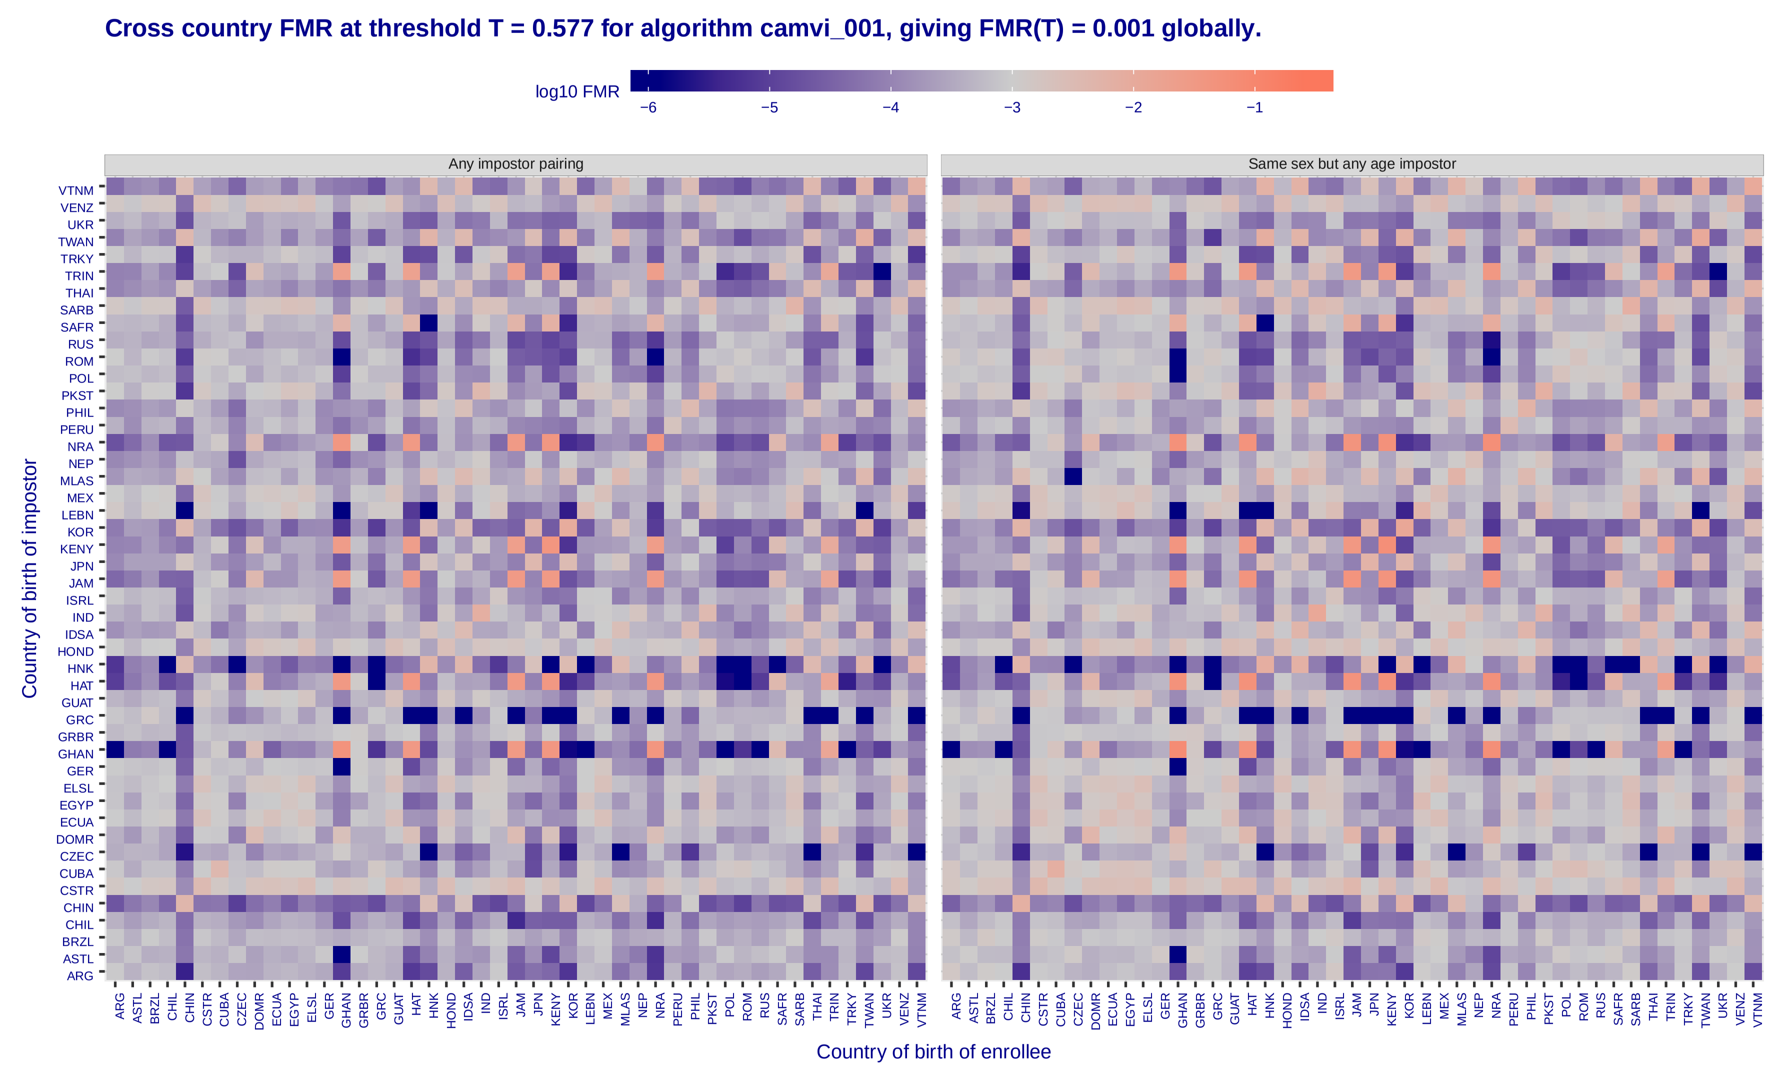

Figure 86: For algorithm camvi-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

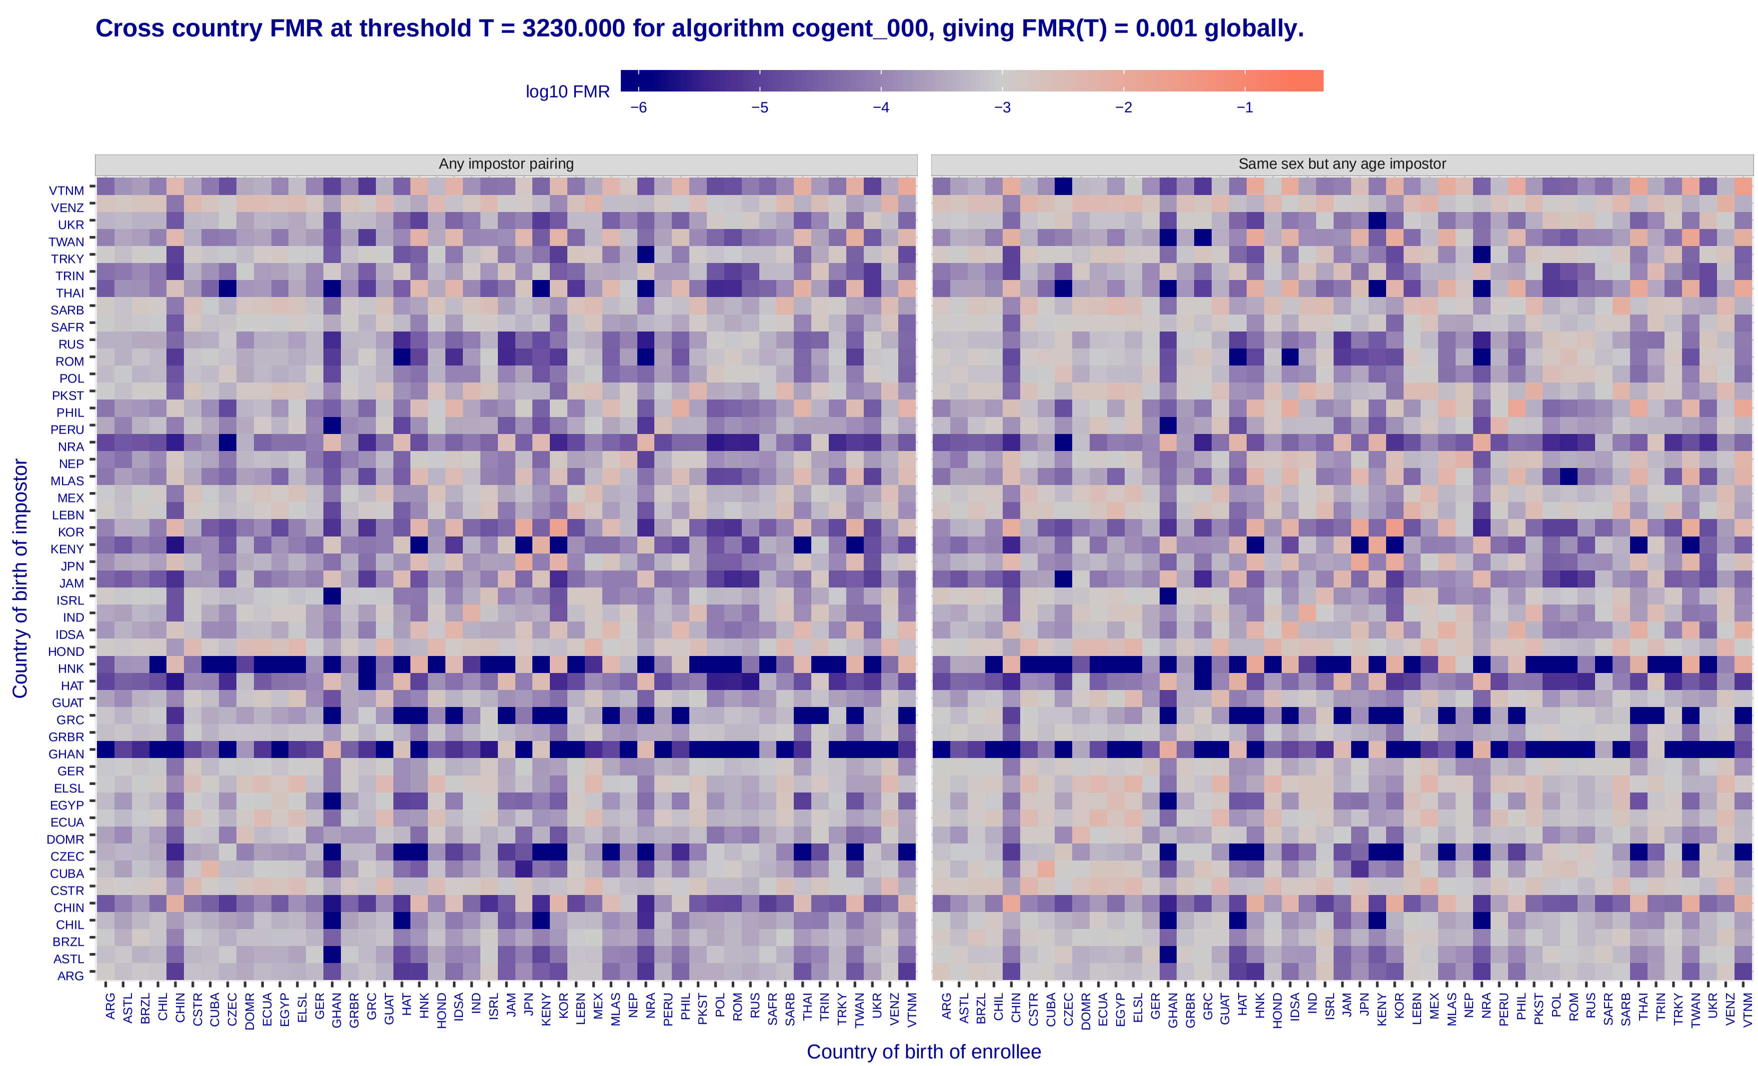

Figure 87: For algorithm cogent-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

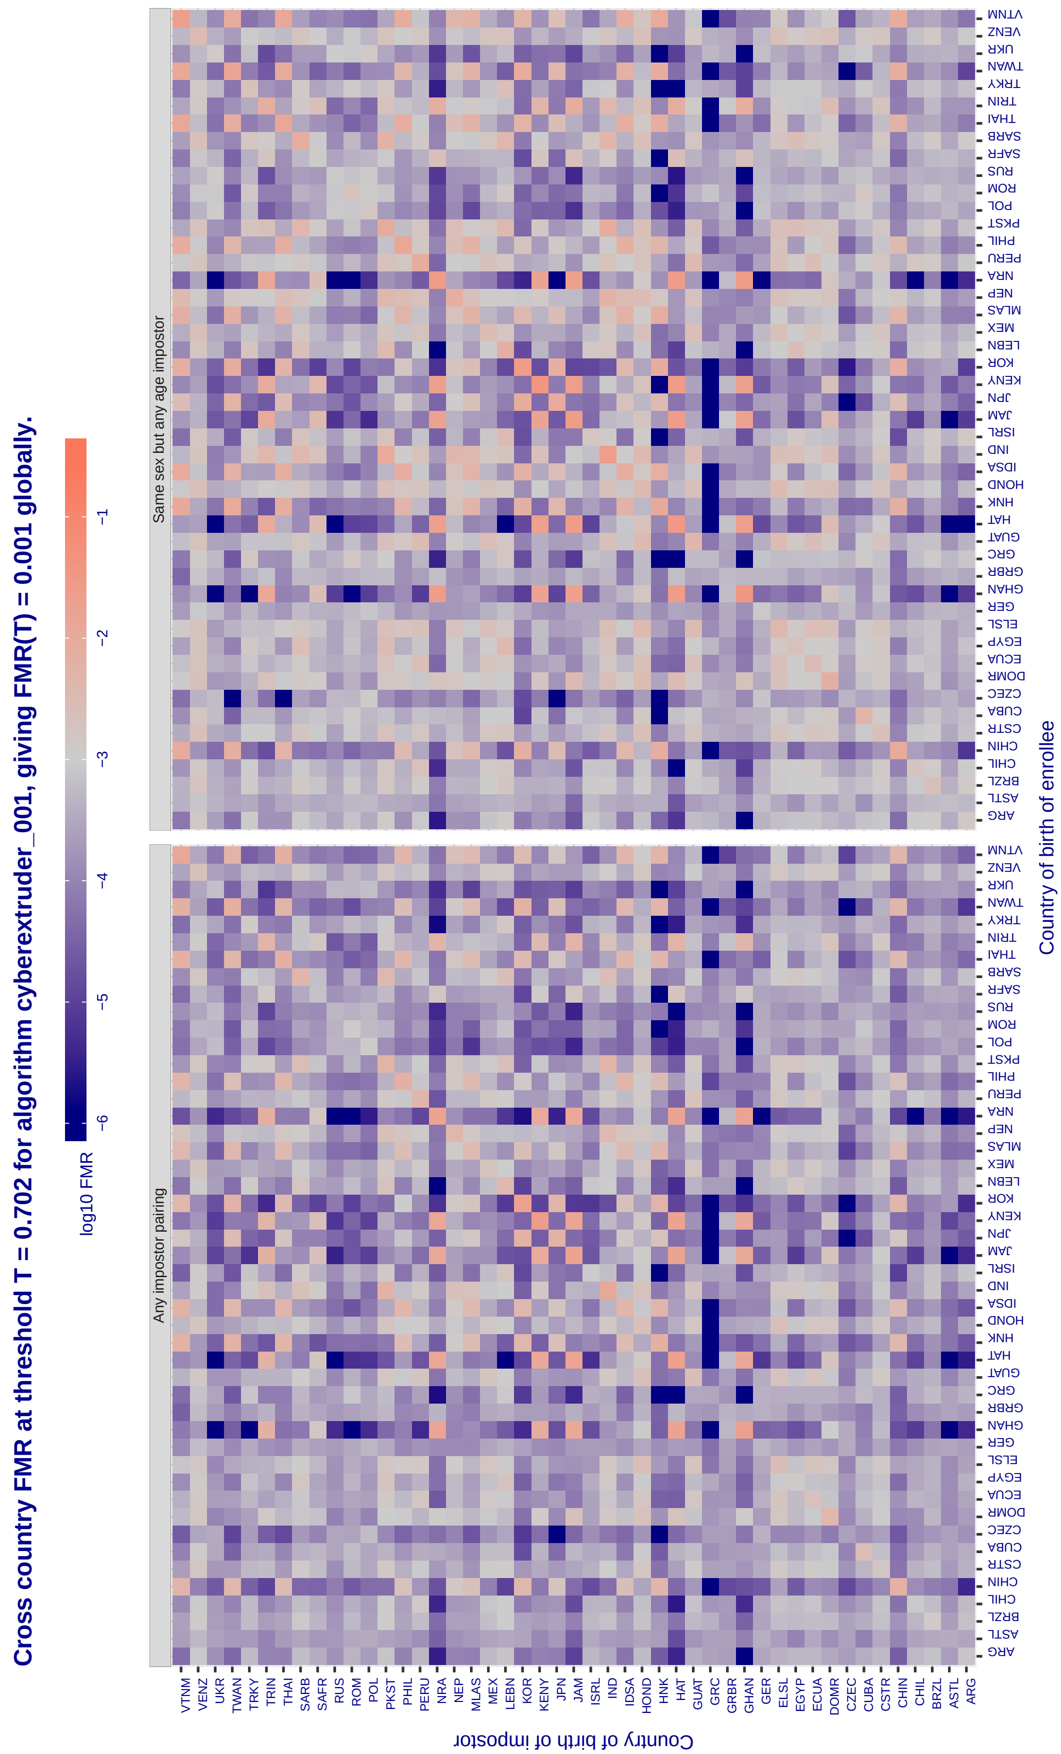

Figure 88: For algorithm cyberextruder-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

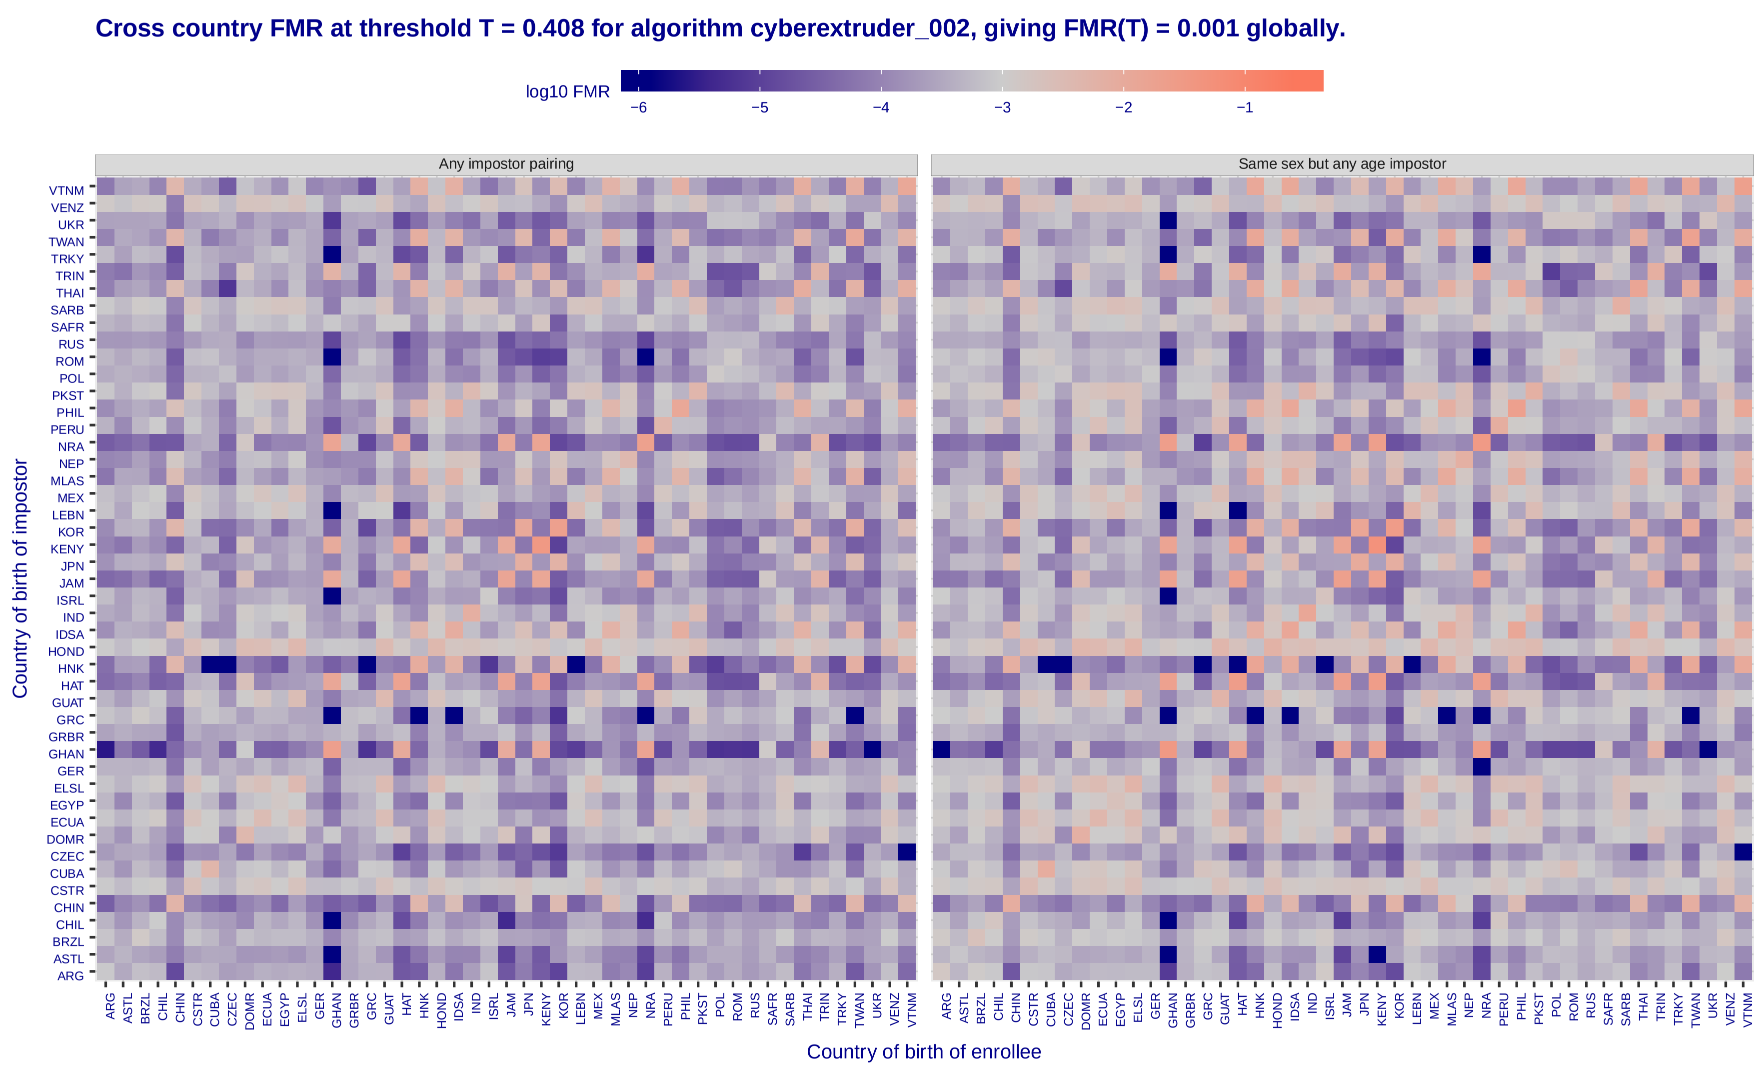

Figure 89: For algorithm cyberextruder-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

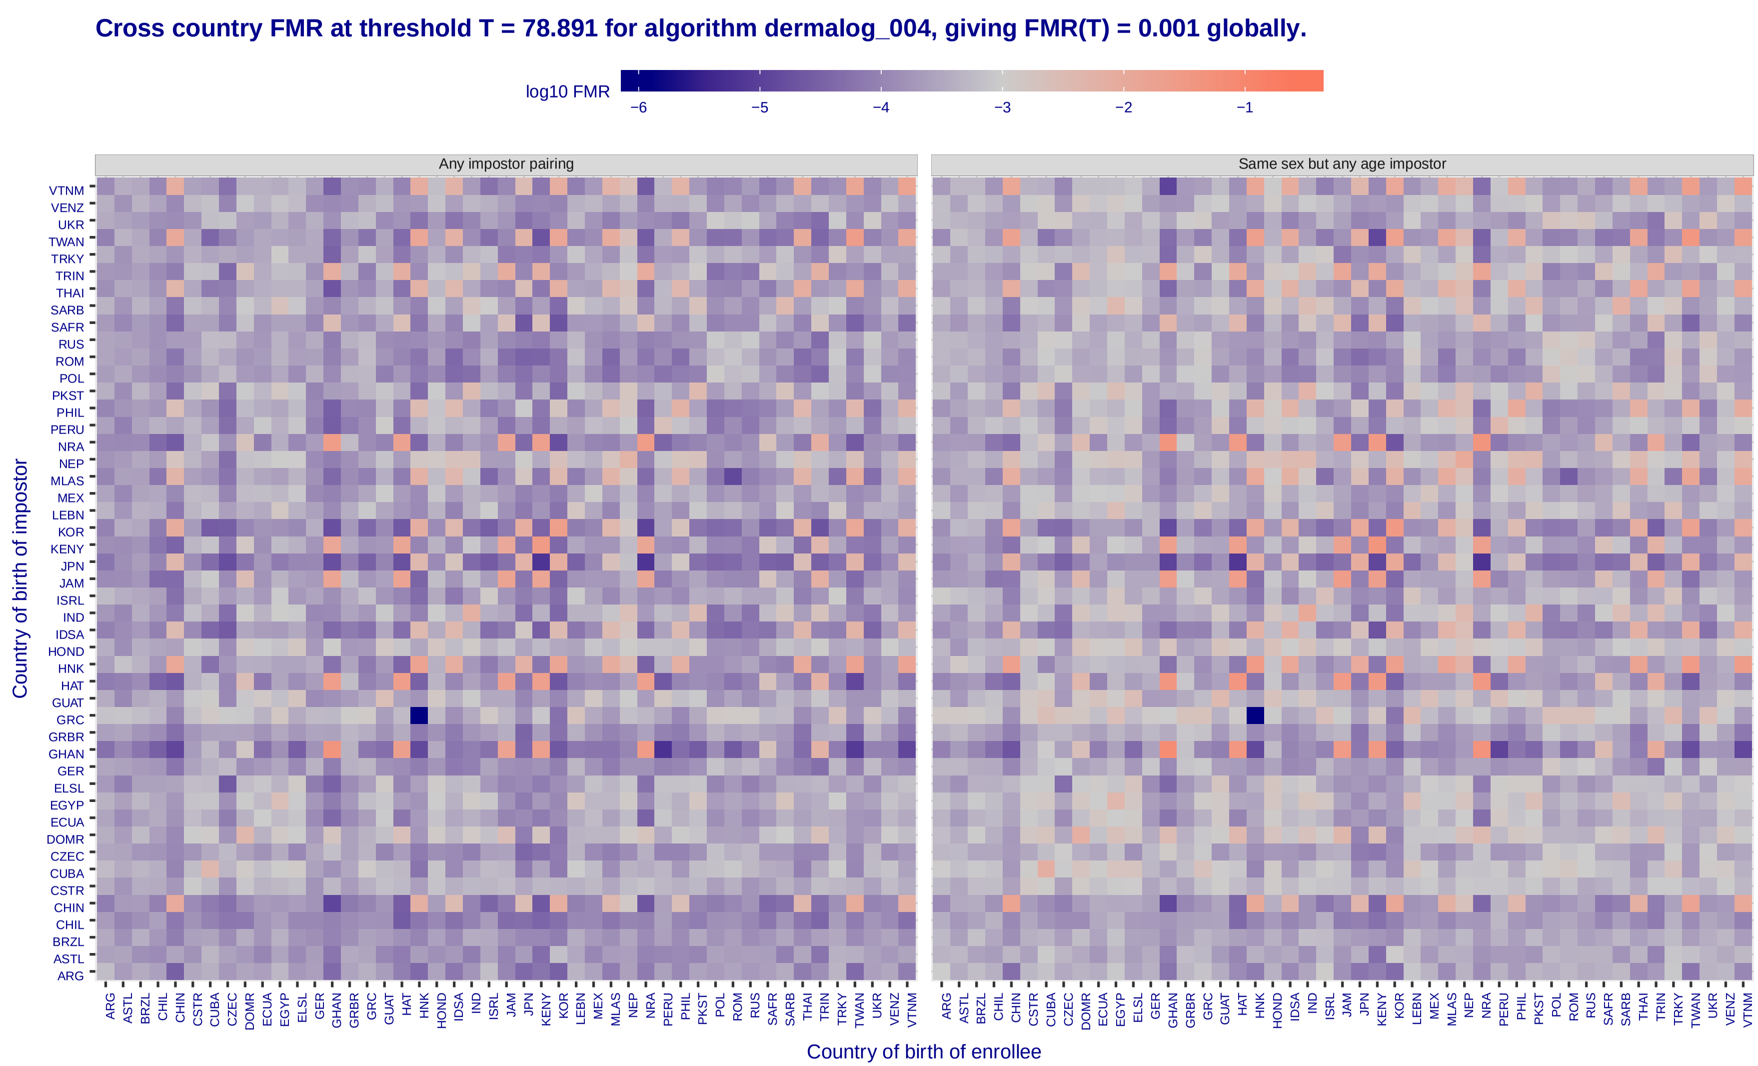

Figure 90: For algorithm dermalog-004 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

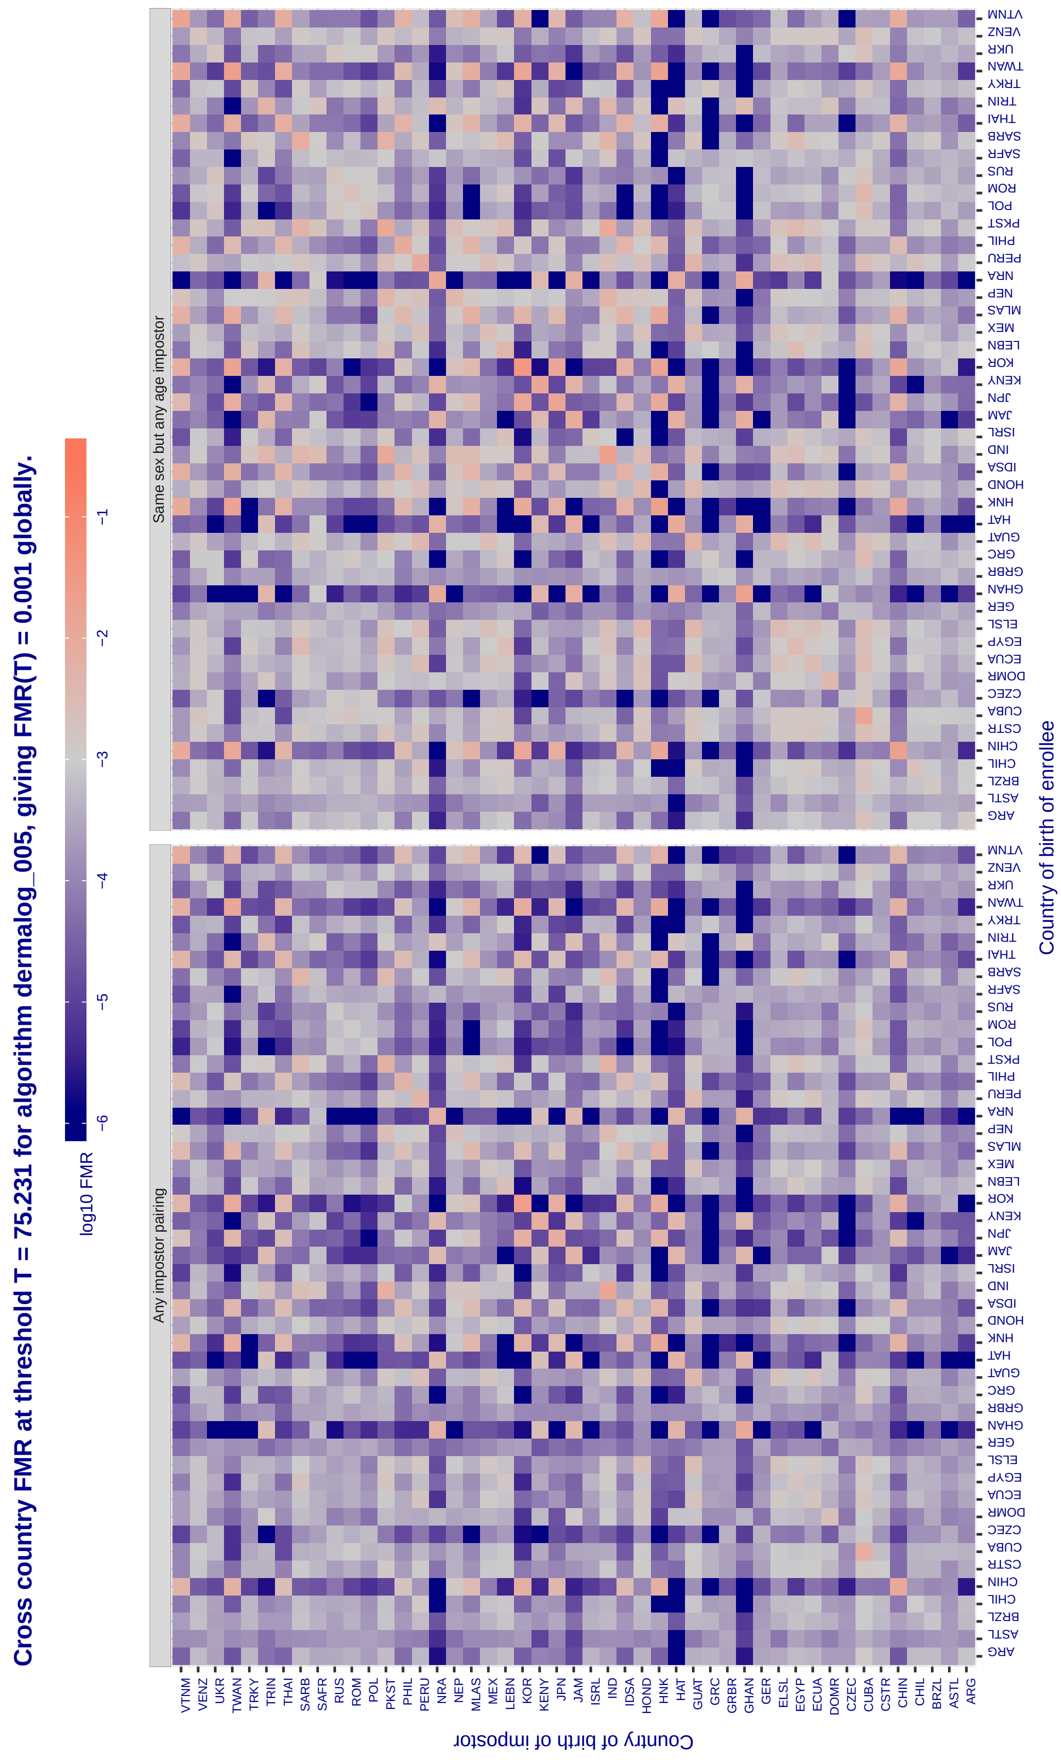

Figure 91: For algorithm dermalog-005 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

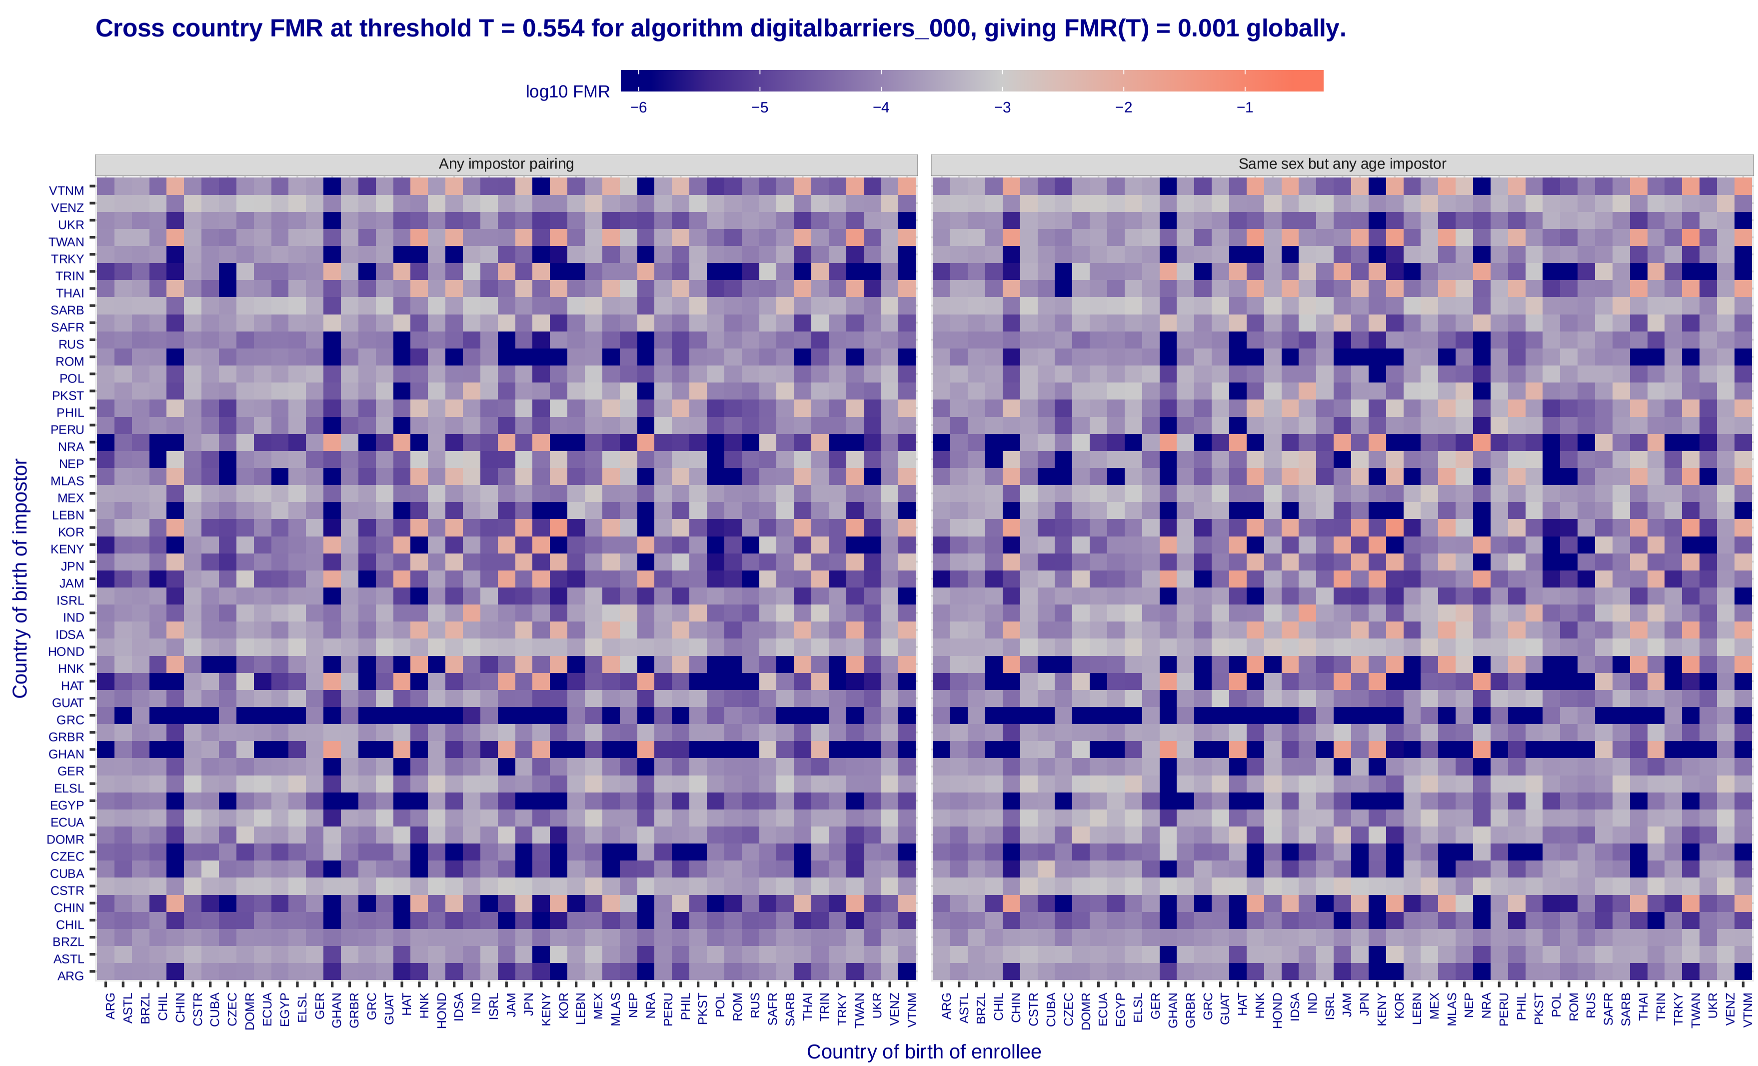

Figure 92: For algorithm digitalbarriers-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

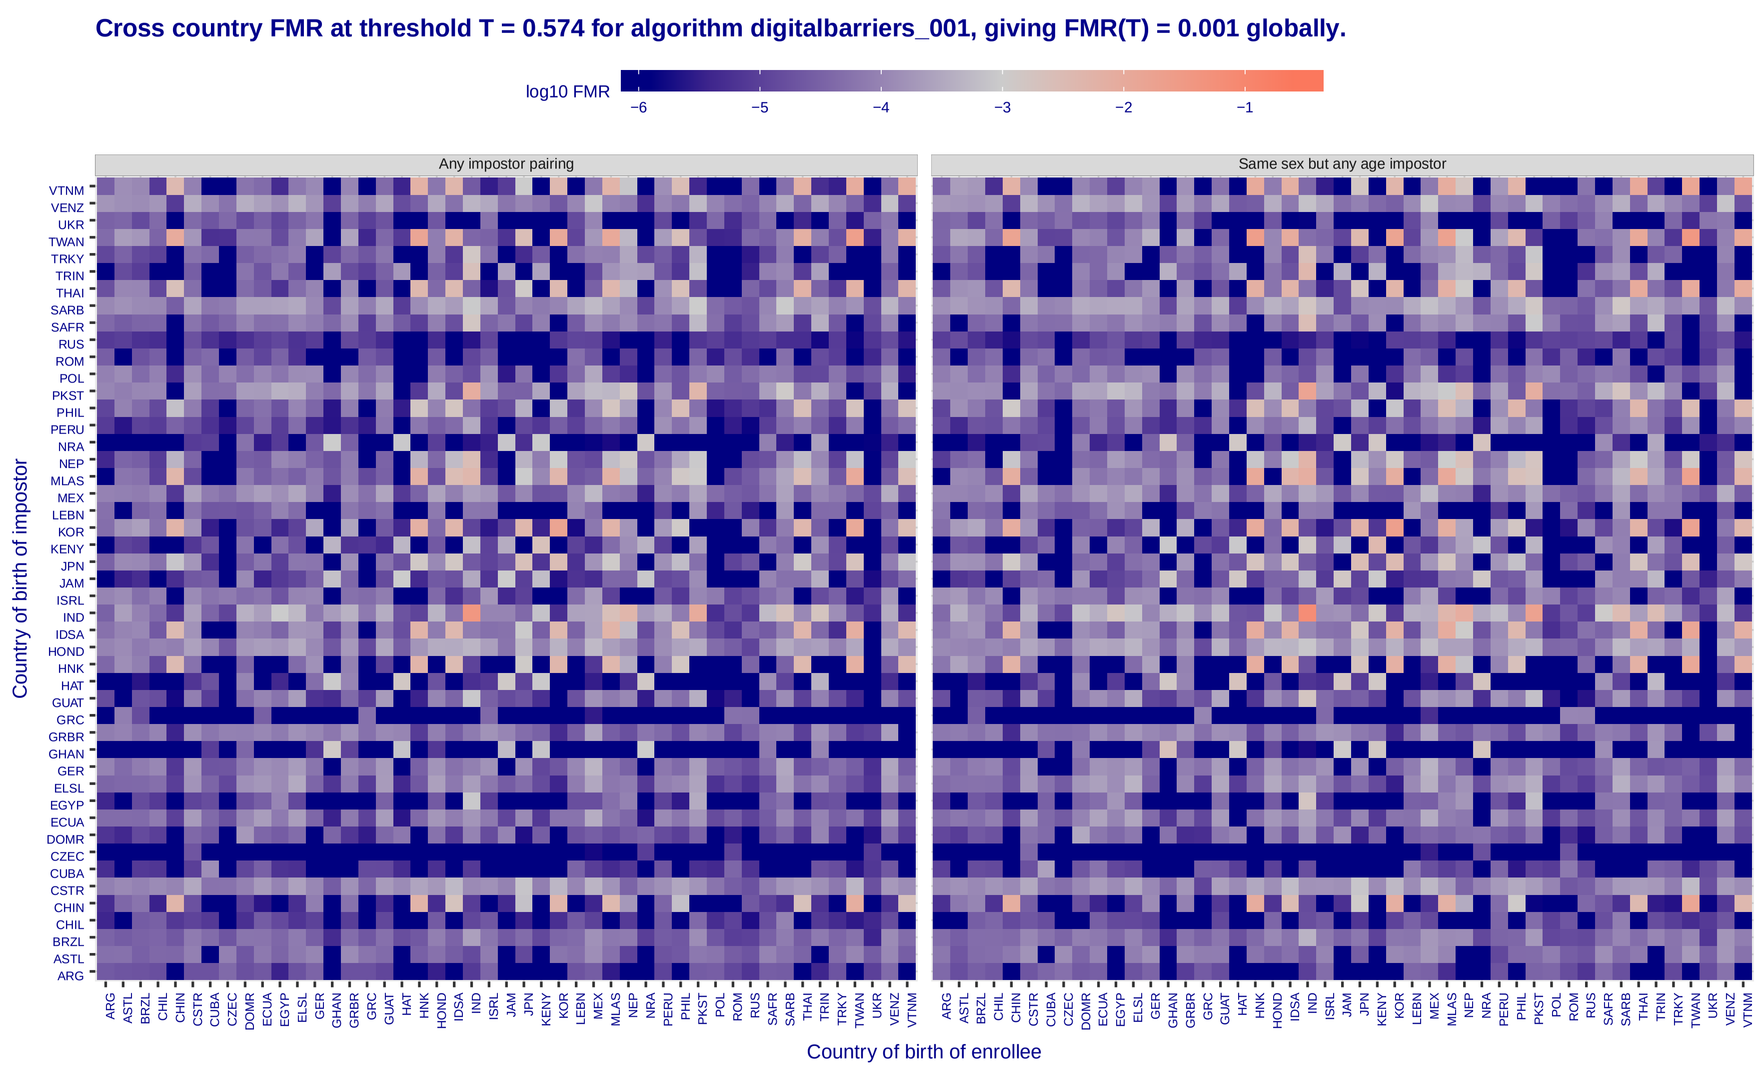

Figure 93: For algorithm digitalbarriers-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

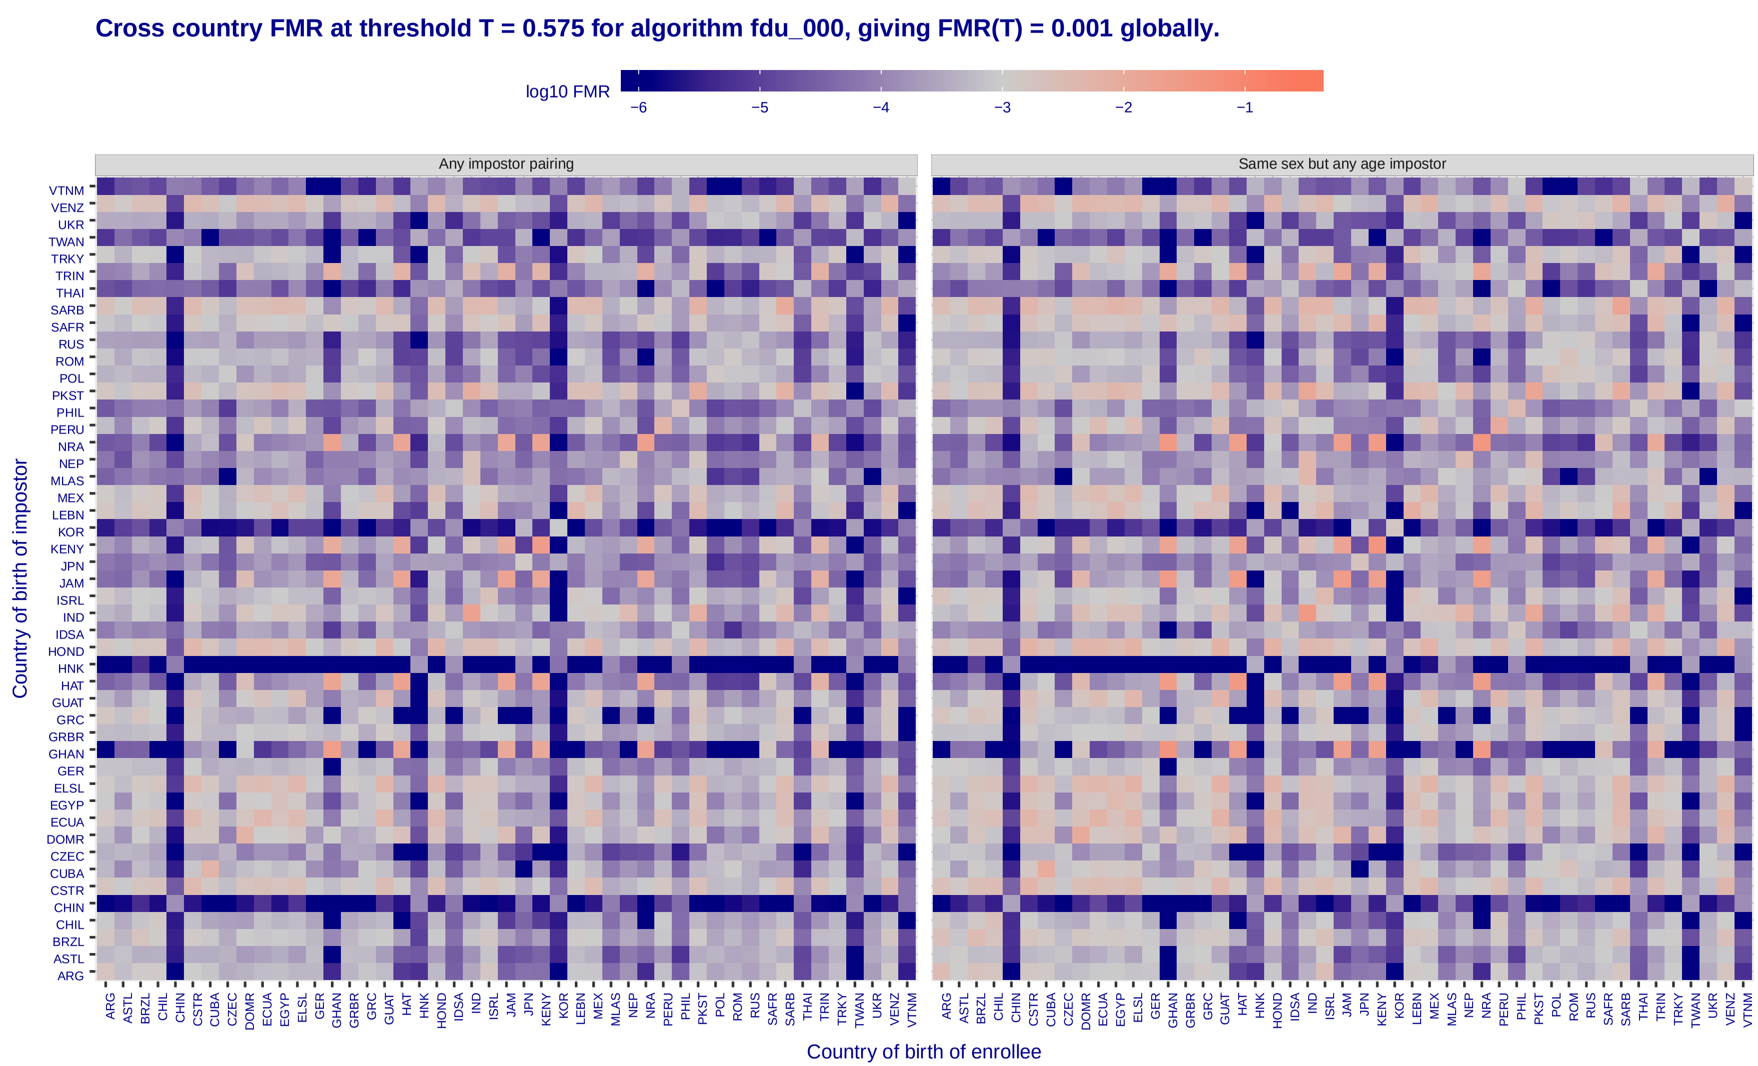

Figure 94: For algorithm fdu-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

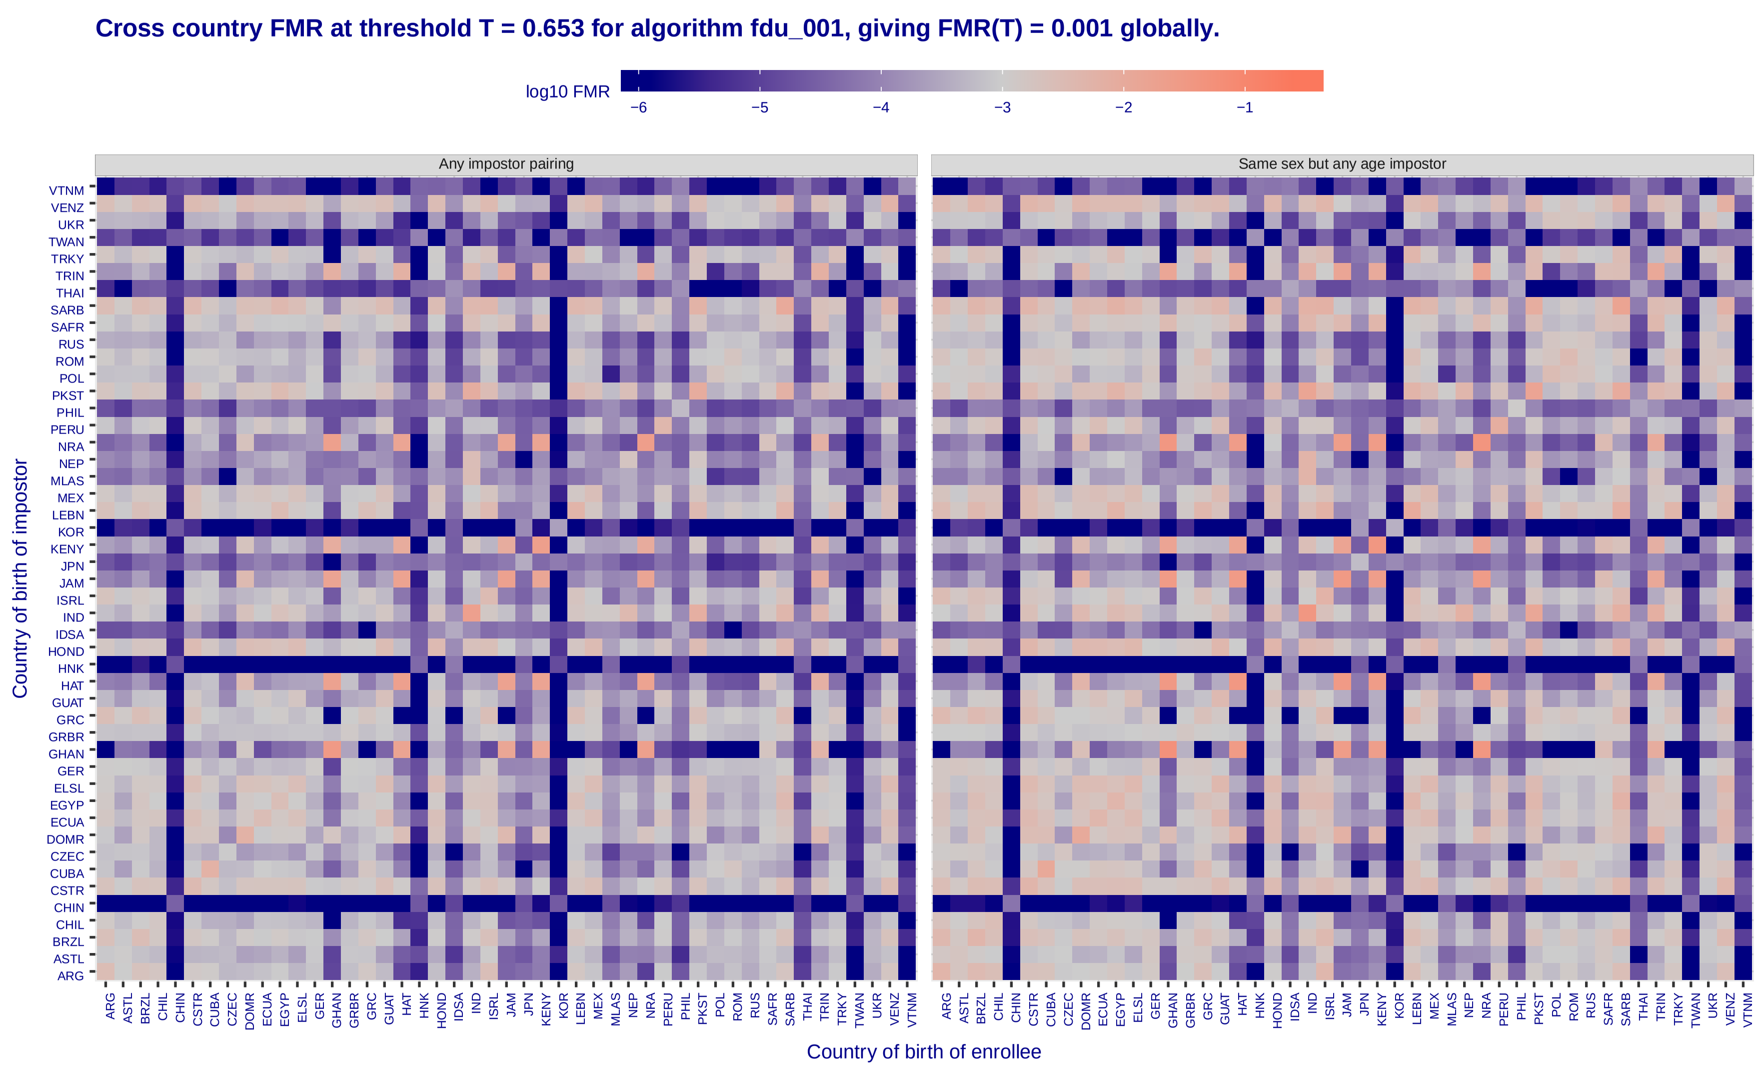

Figure 95: For algorithm fdu-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

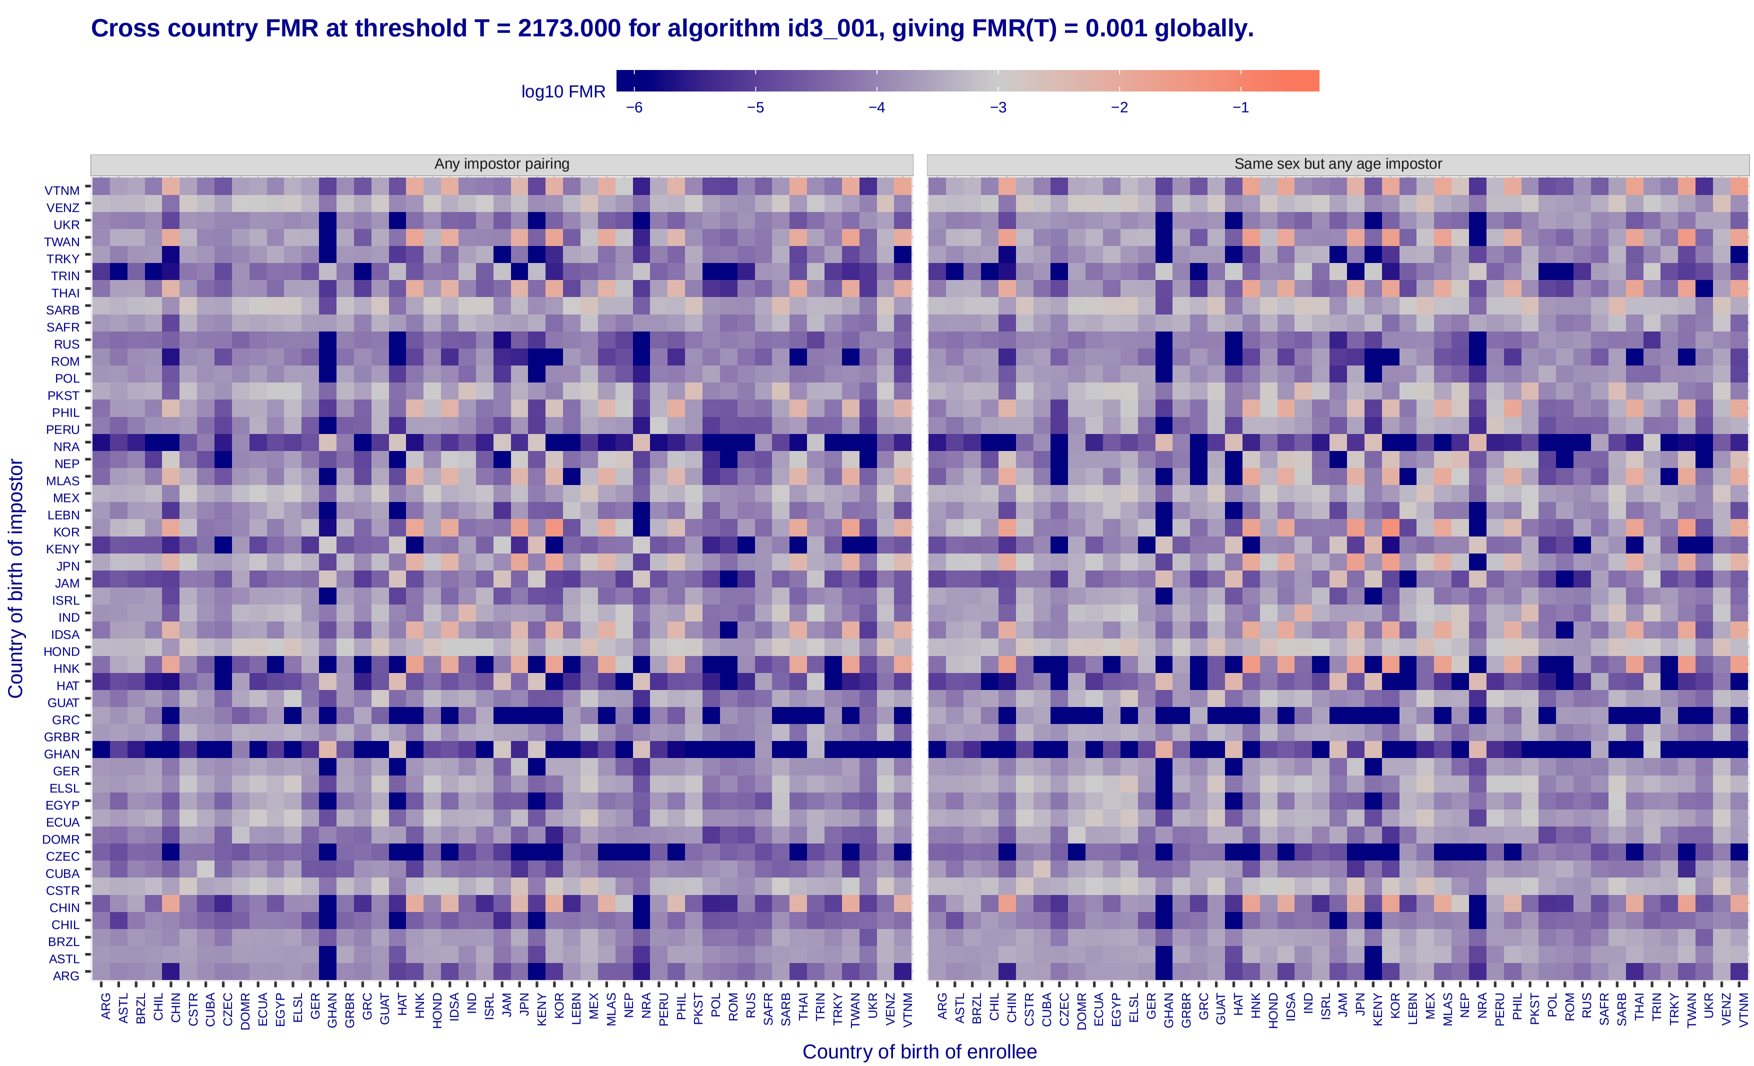

Figure 96: For algorithm id3-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

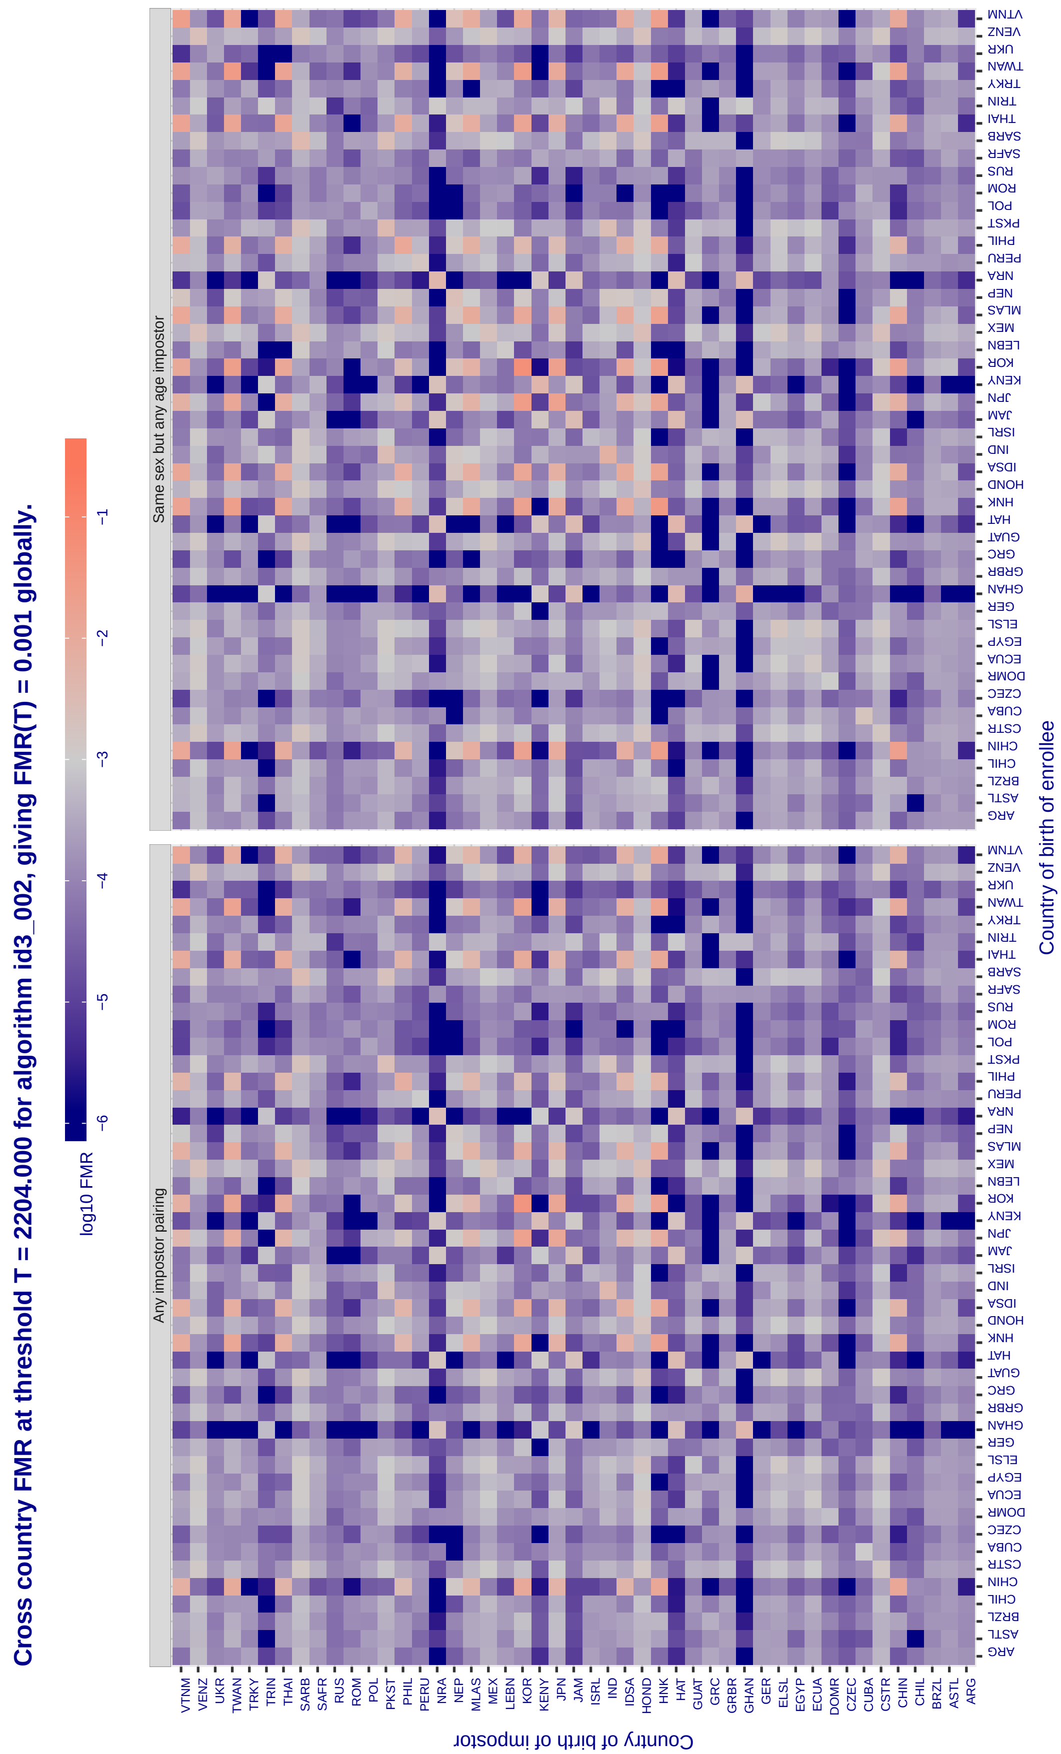

Figure 97: For algorithm id3-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

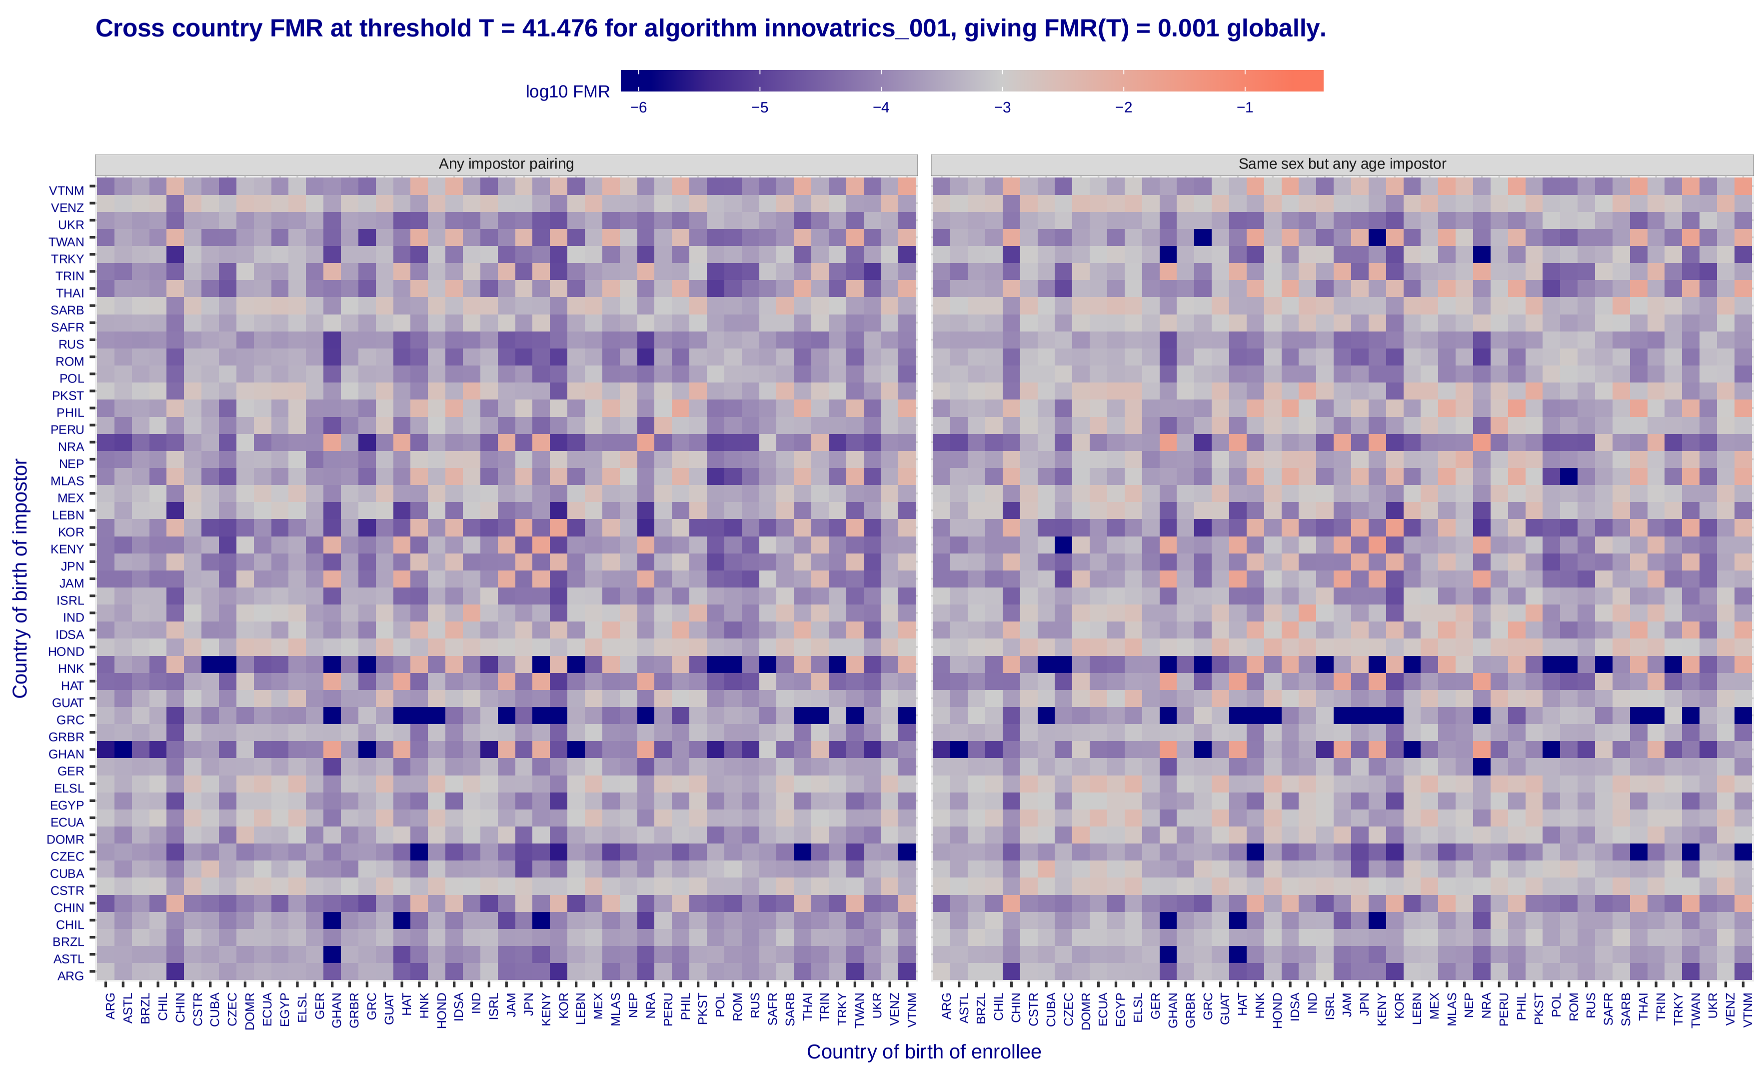

Figure 98: For algorithm innovatrics-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

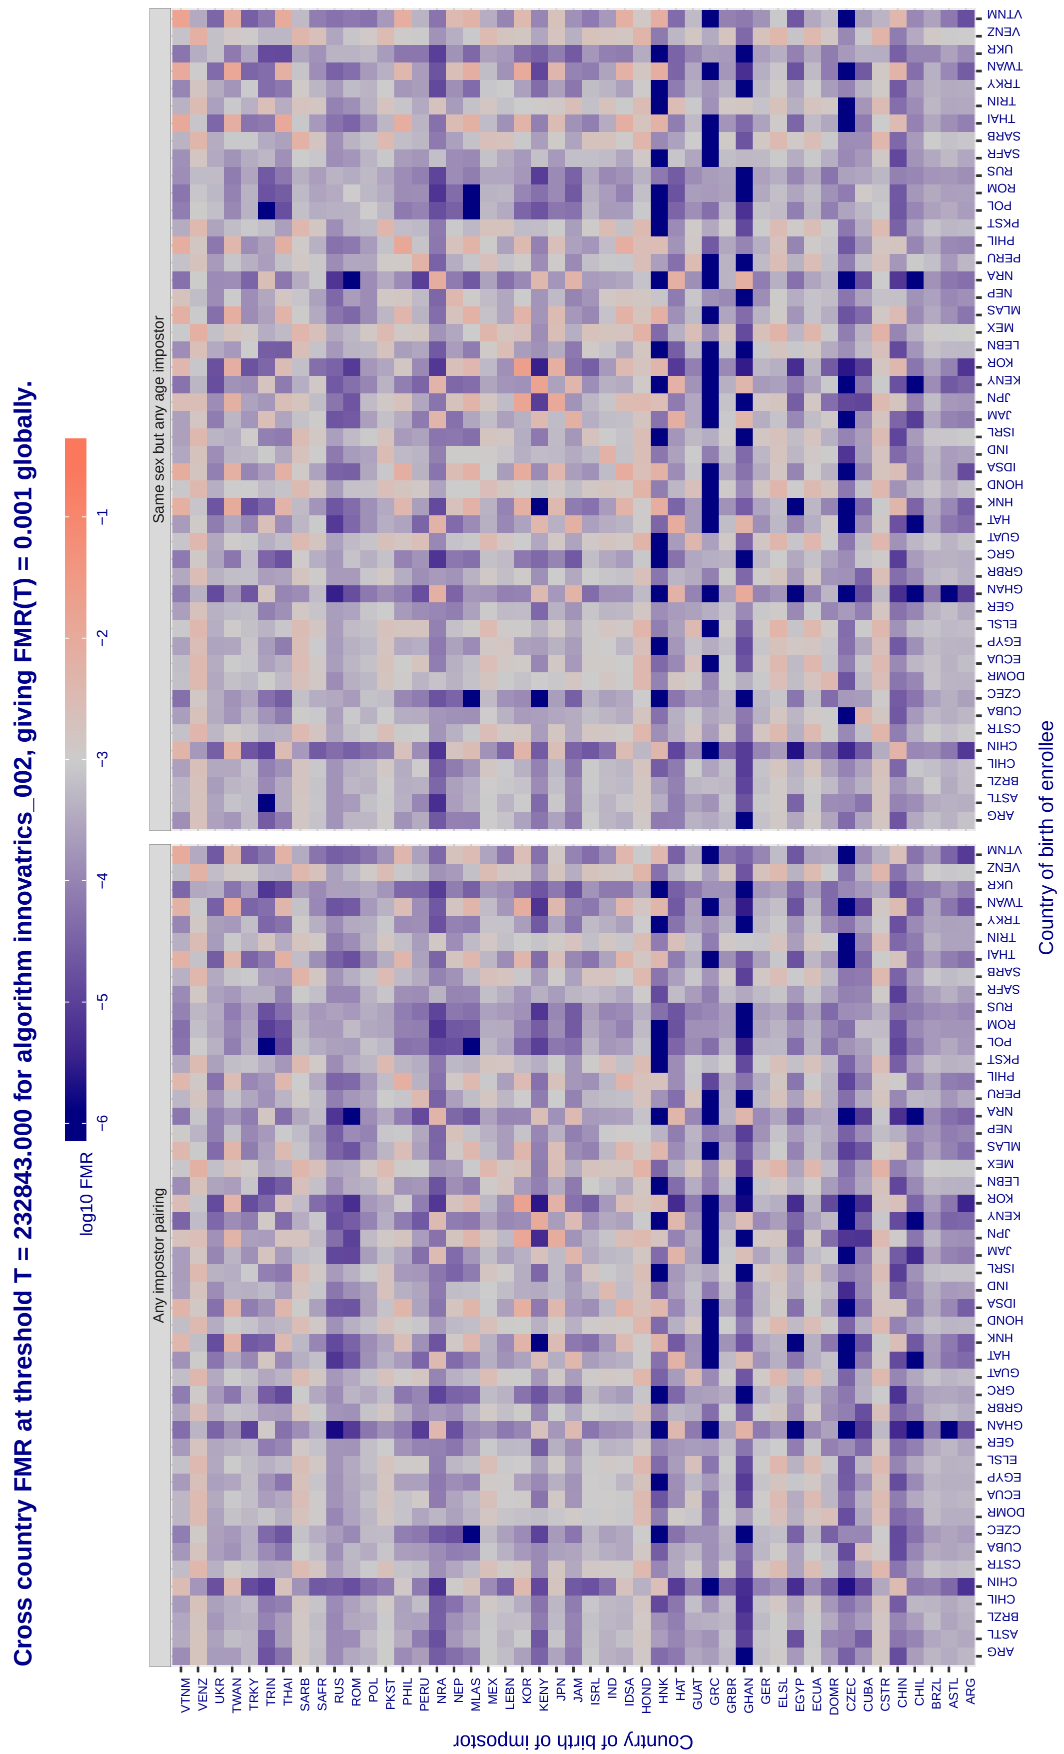

Figure 99: For algorithm innovatrics-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

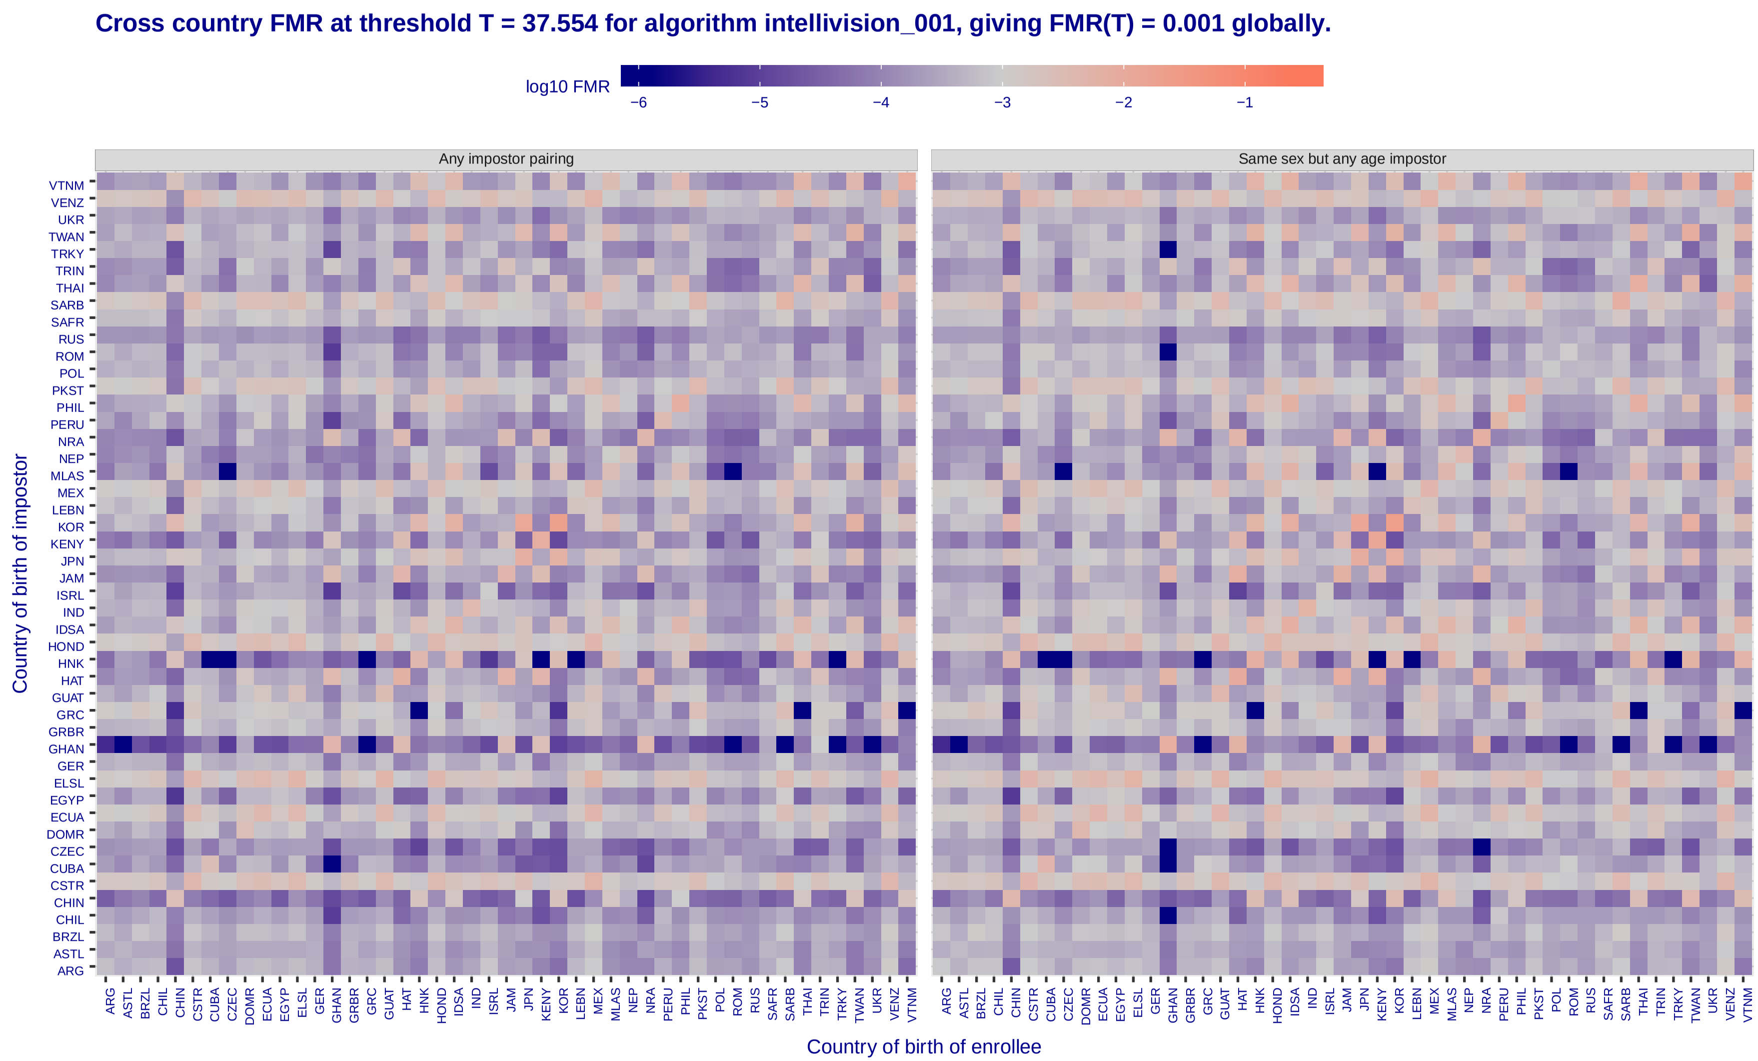

Figure 100: For algorithm intellivision-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

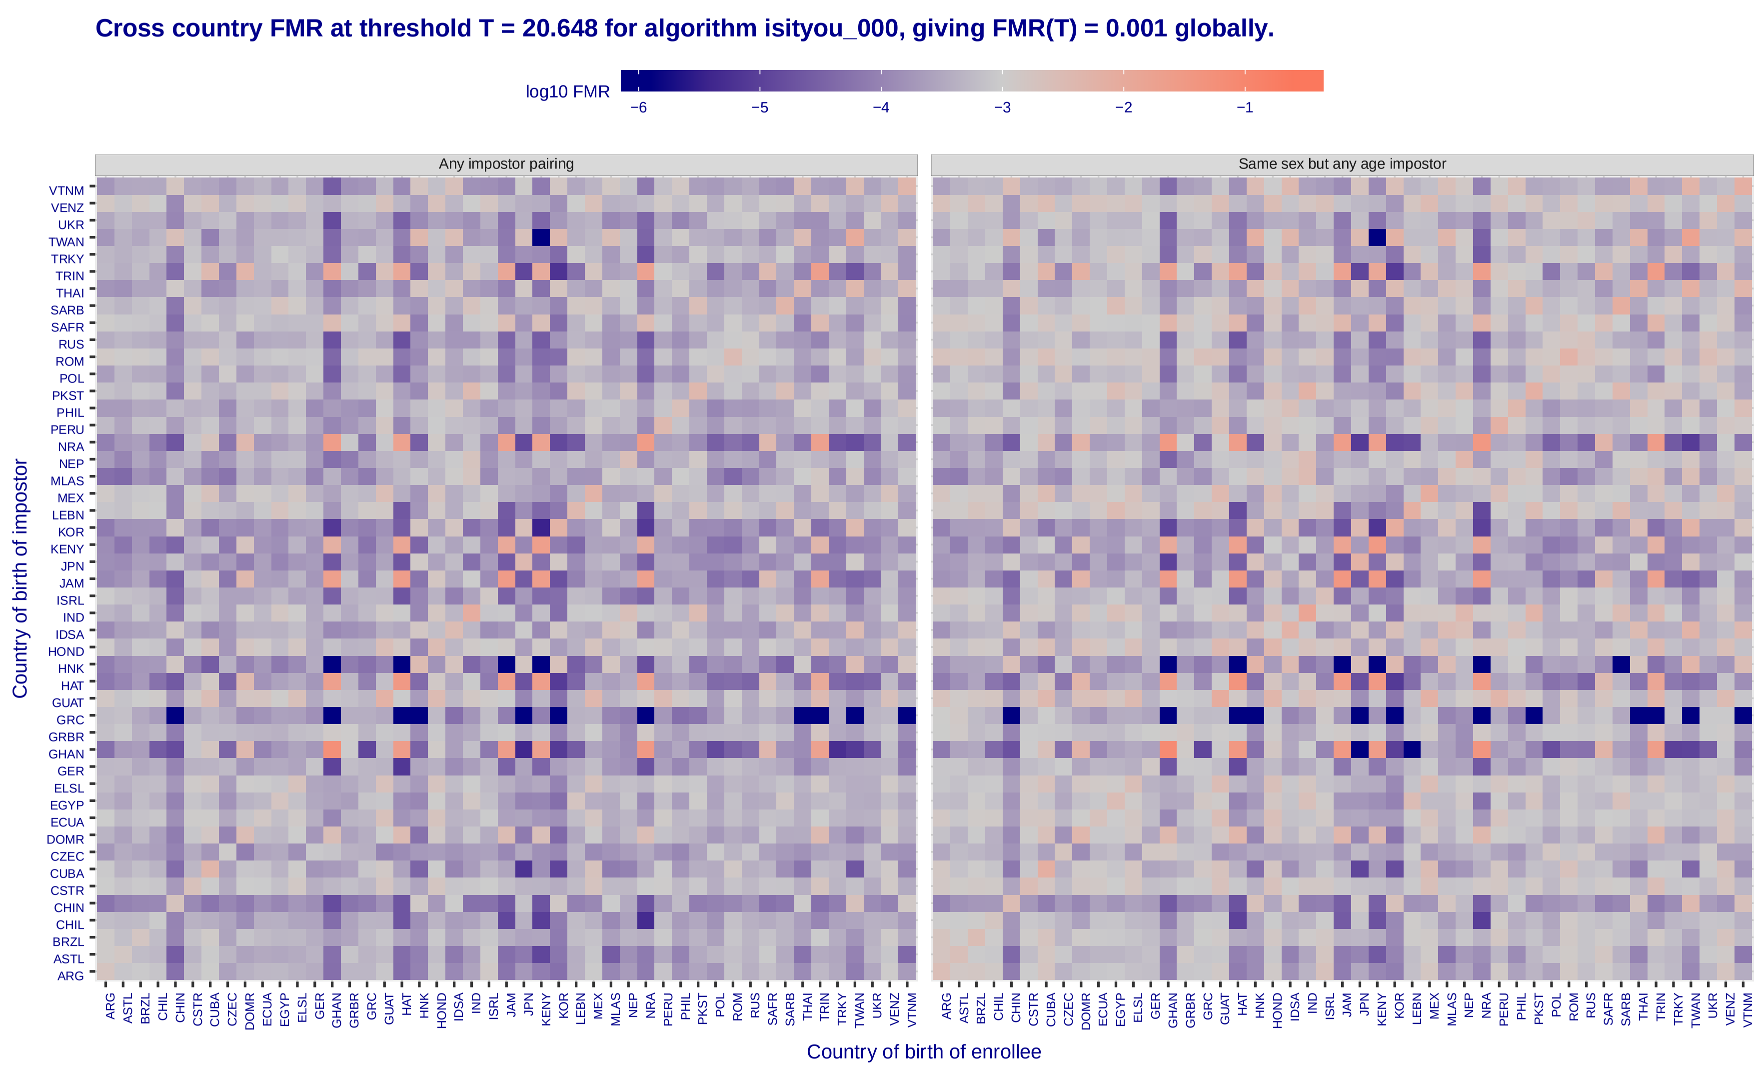

Figure 101: For algorithm isityou-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

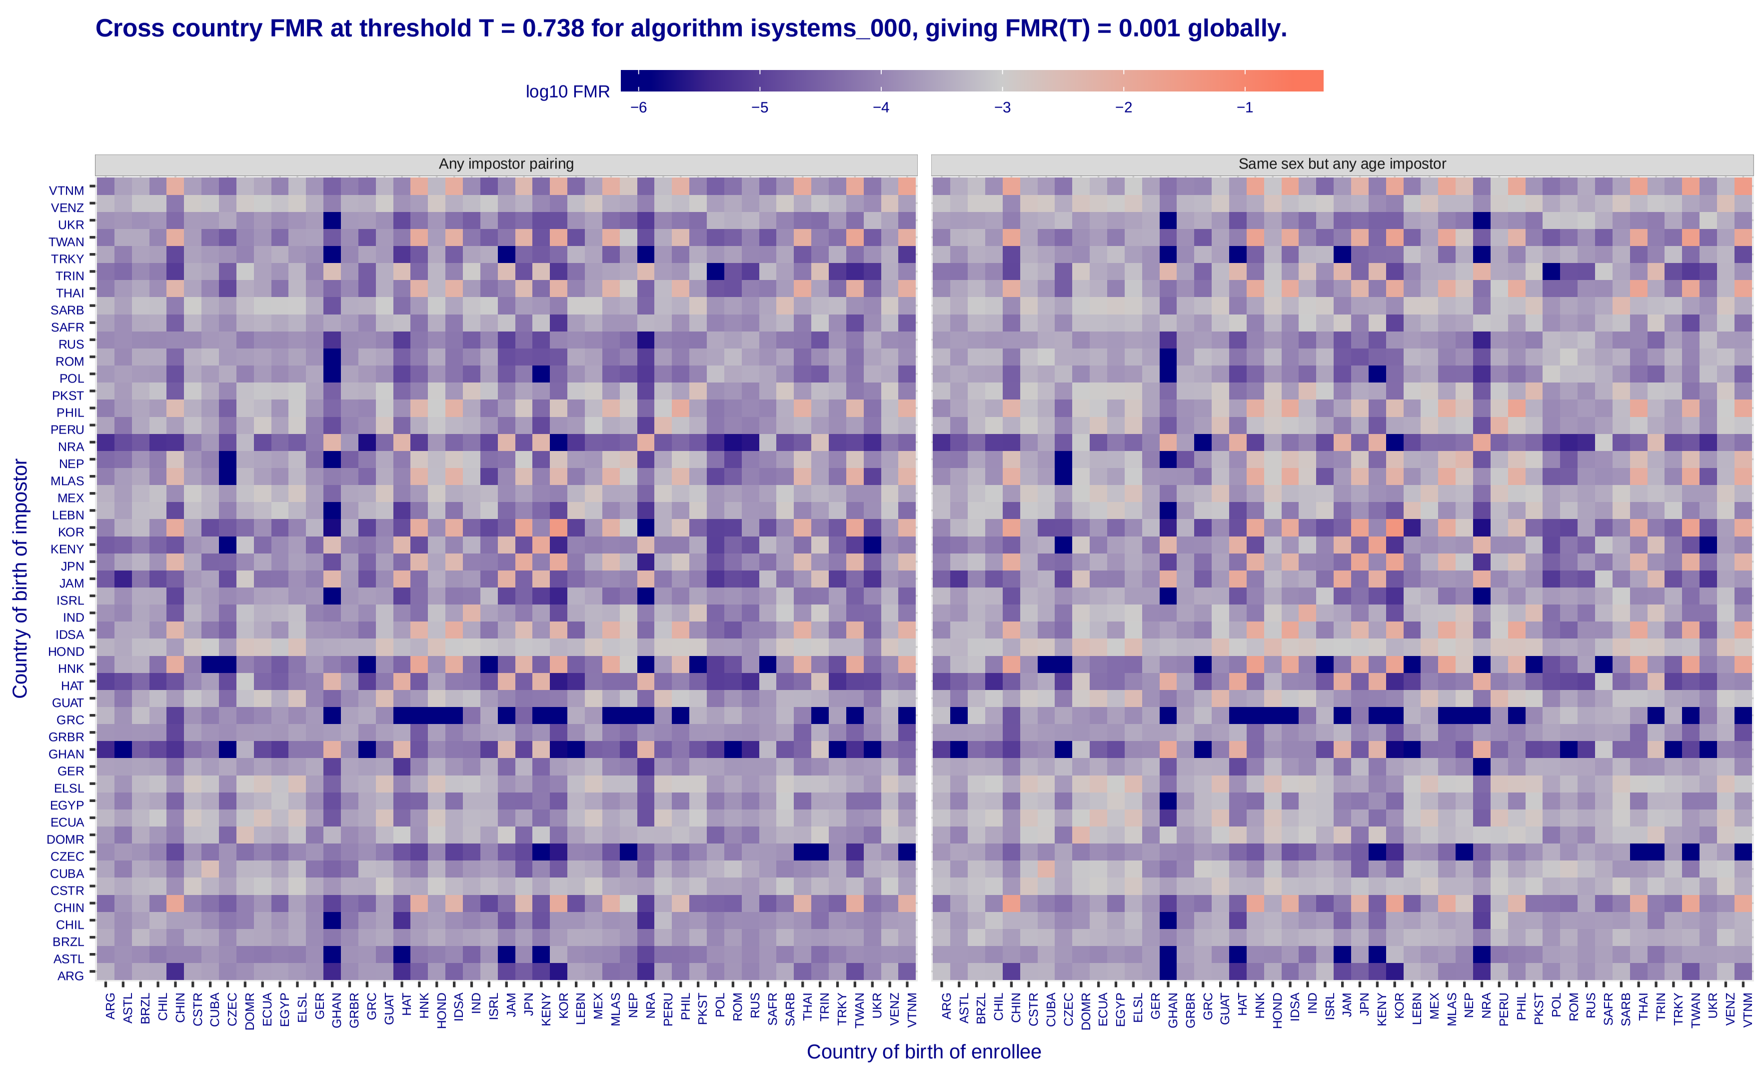

Figure 102: *For algorithm isystems-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.*

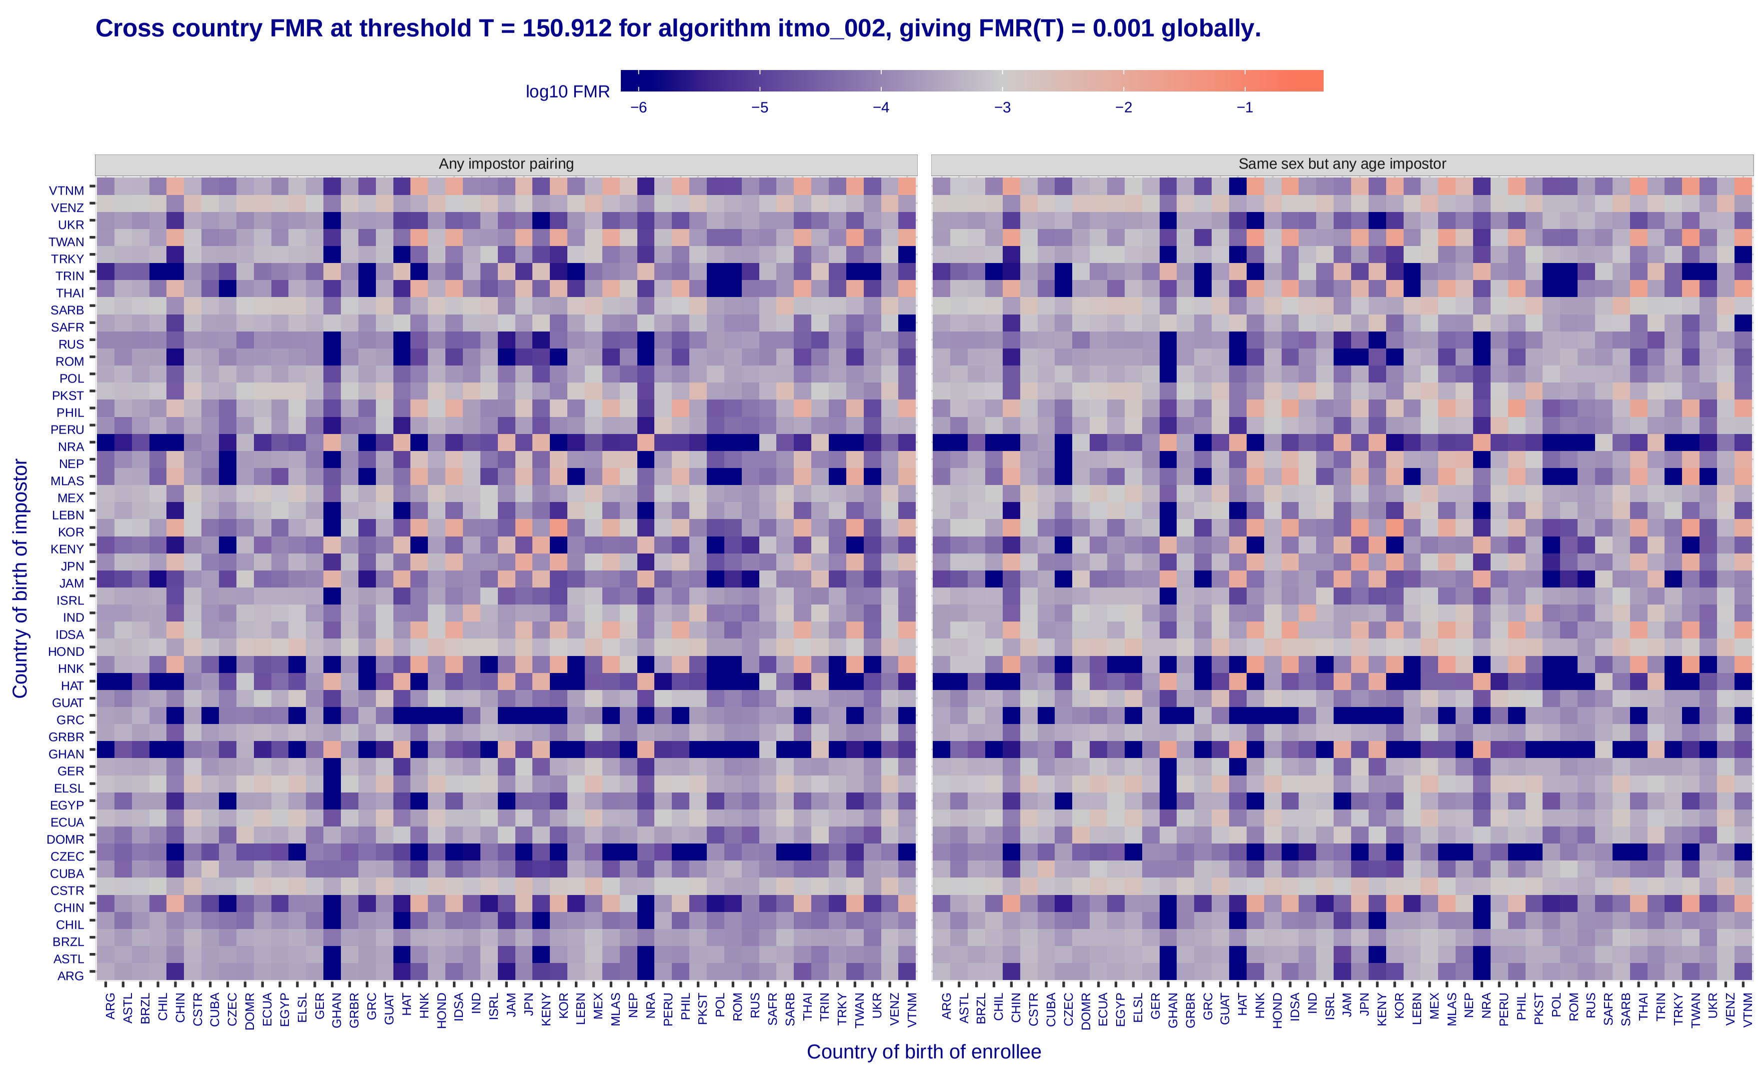

Figure 103: For algorithm itmo-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

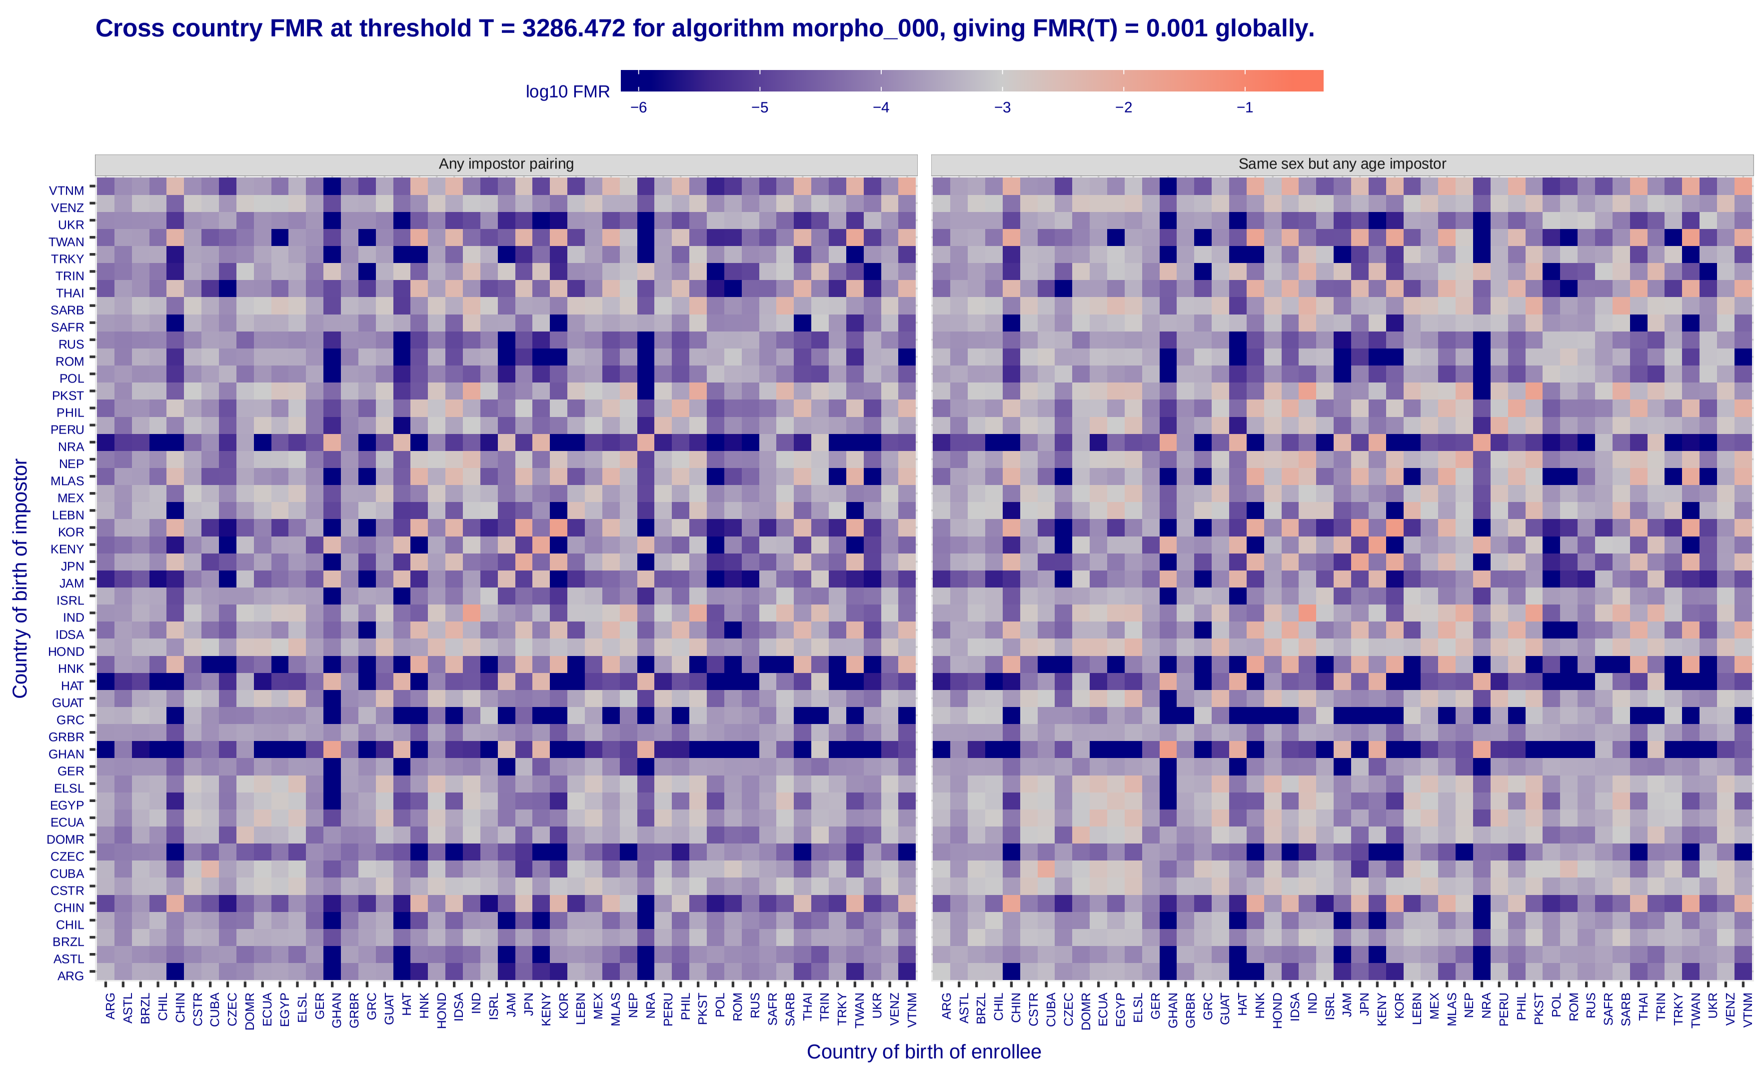

Figure 104: For algorithm morpho-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

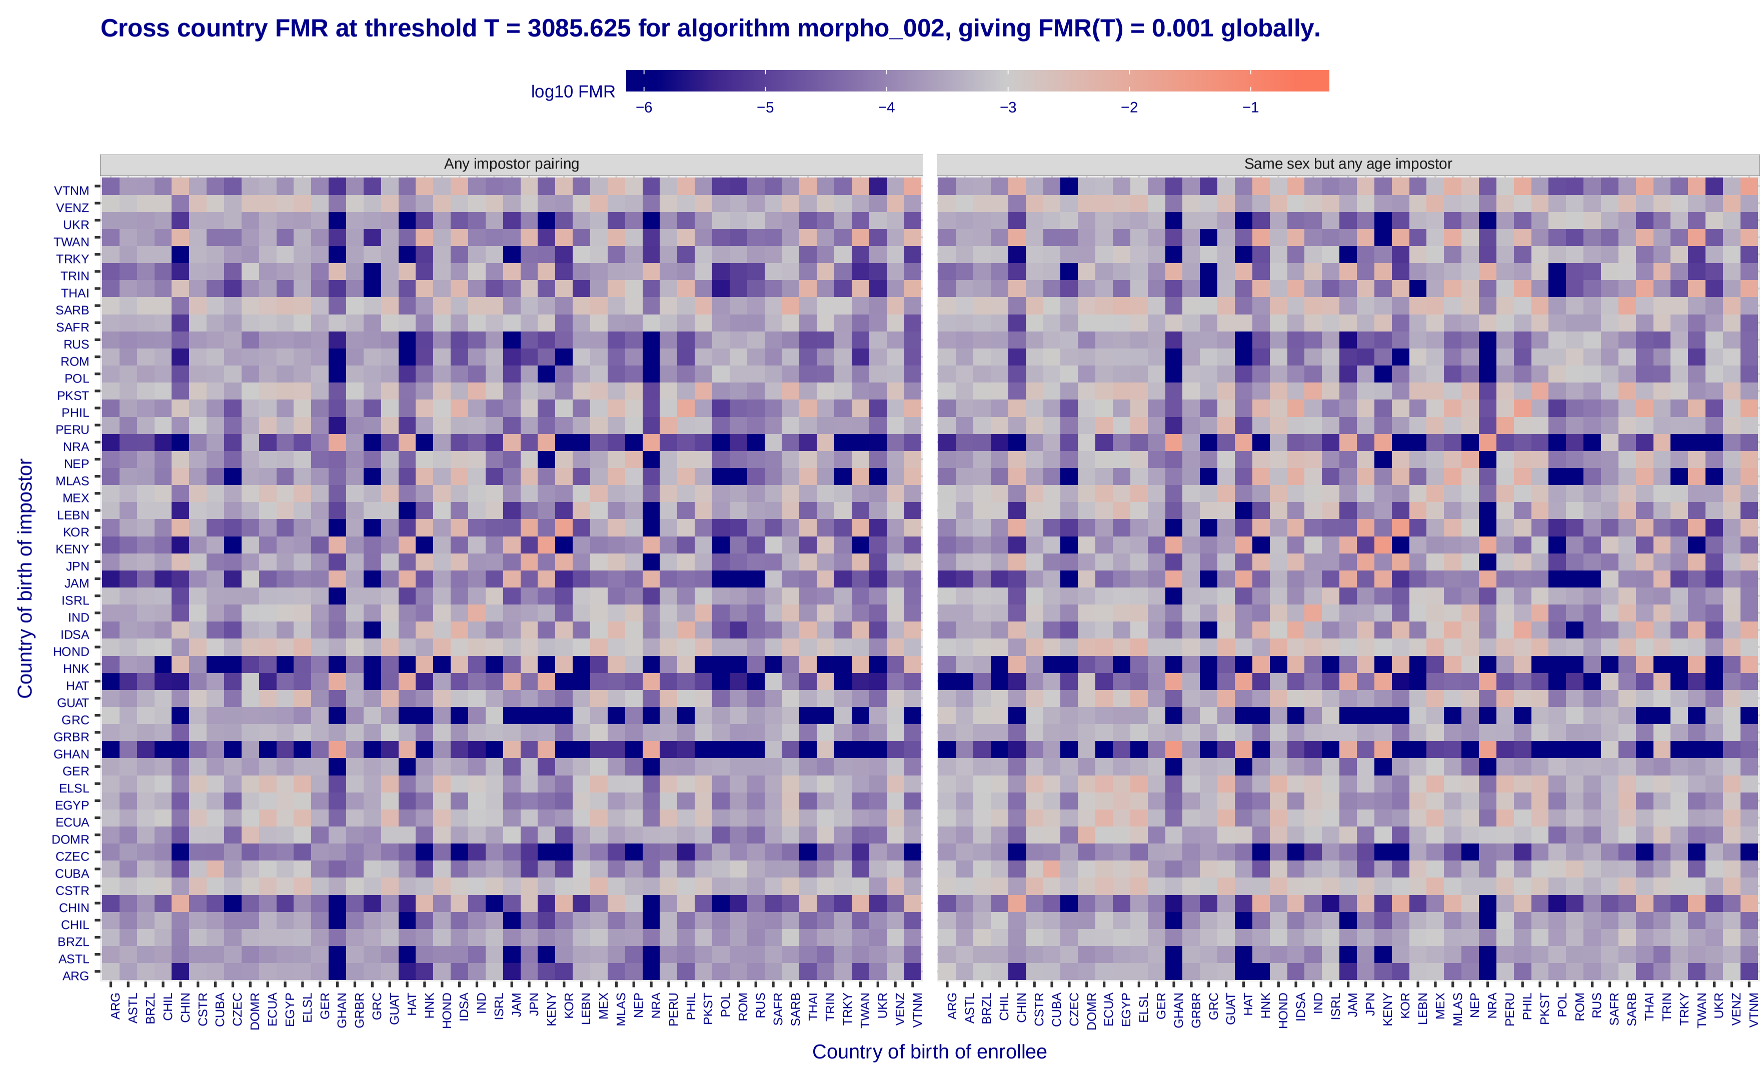

Figure 105: For algorithm morpho-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

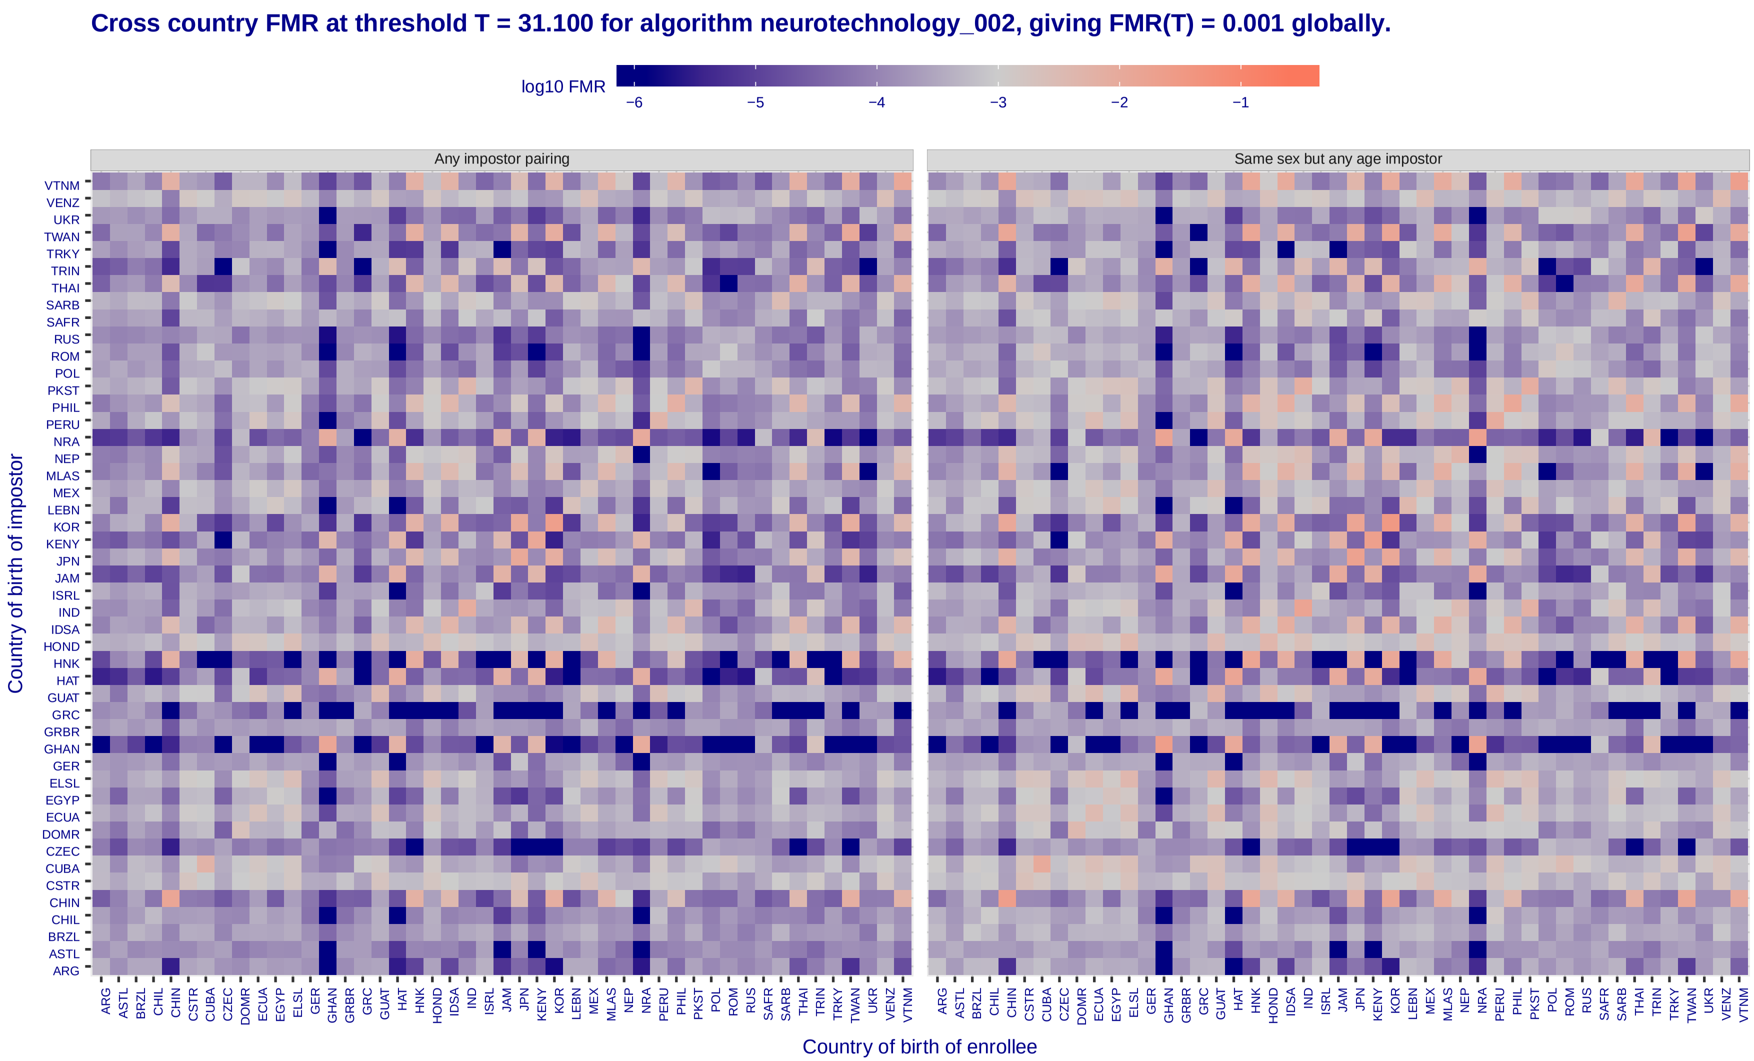

Figure 106: For algorithm neurotechnology-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

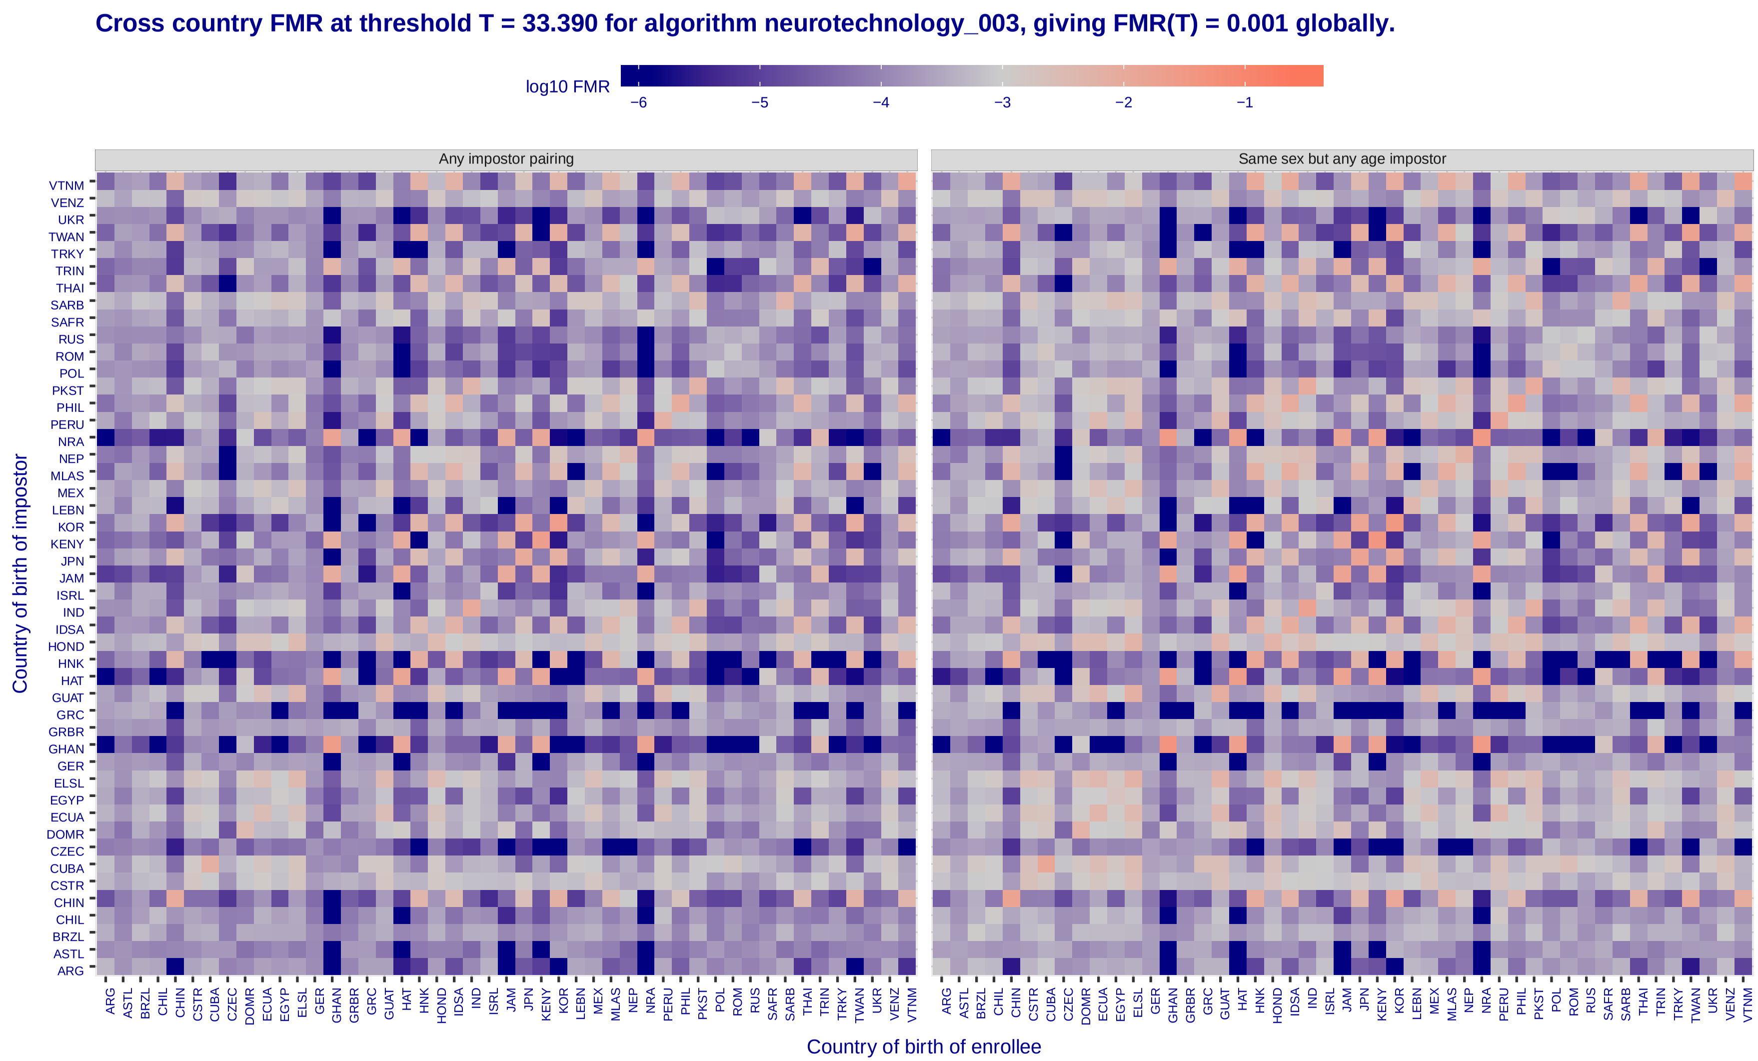

Figure 107: For algorithm neurotechnology-003 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

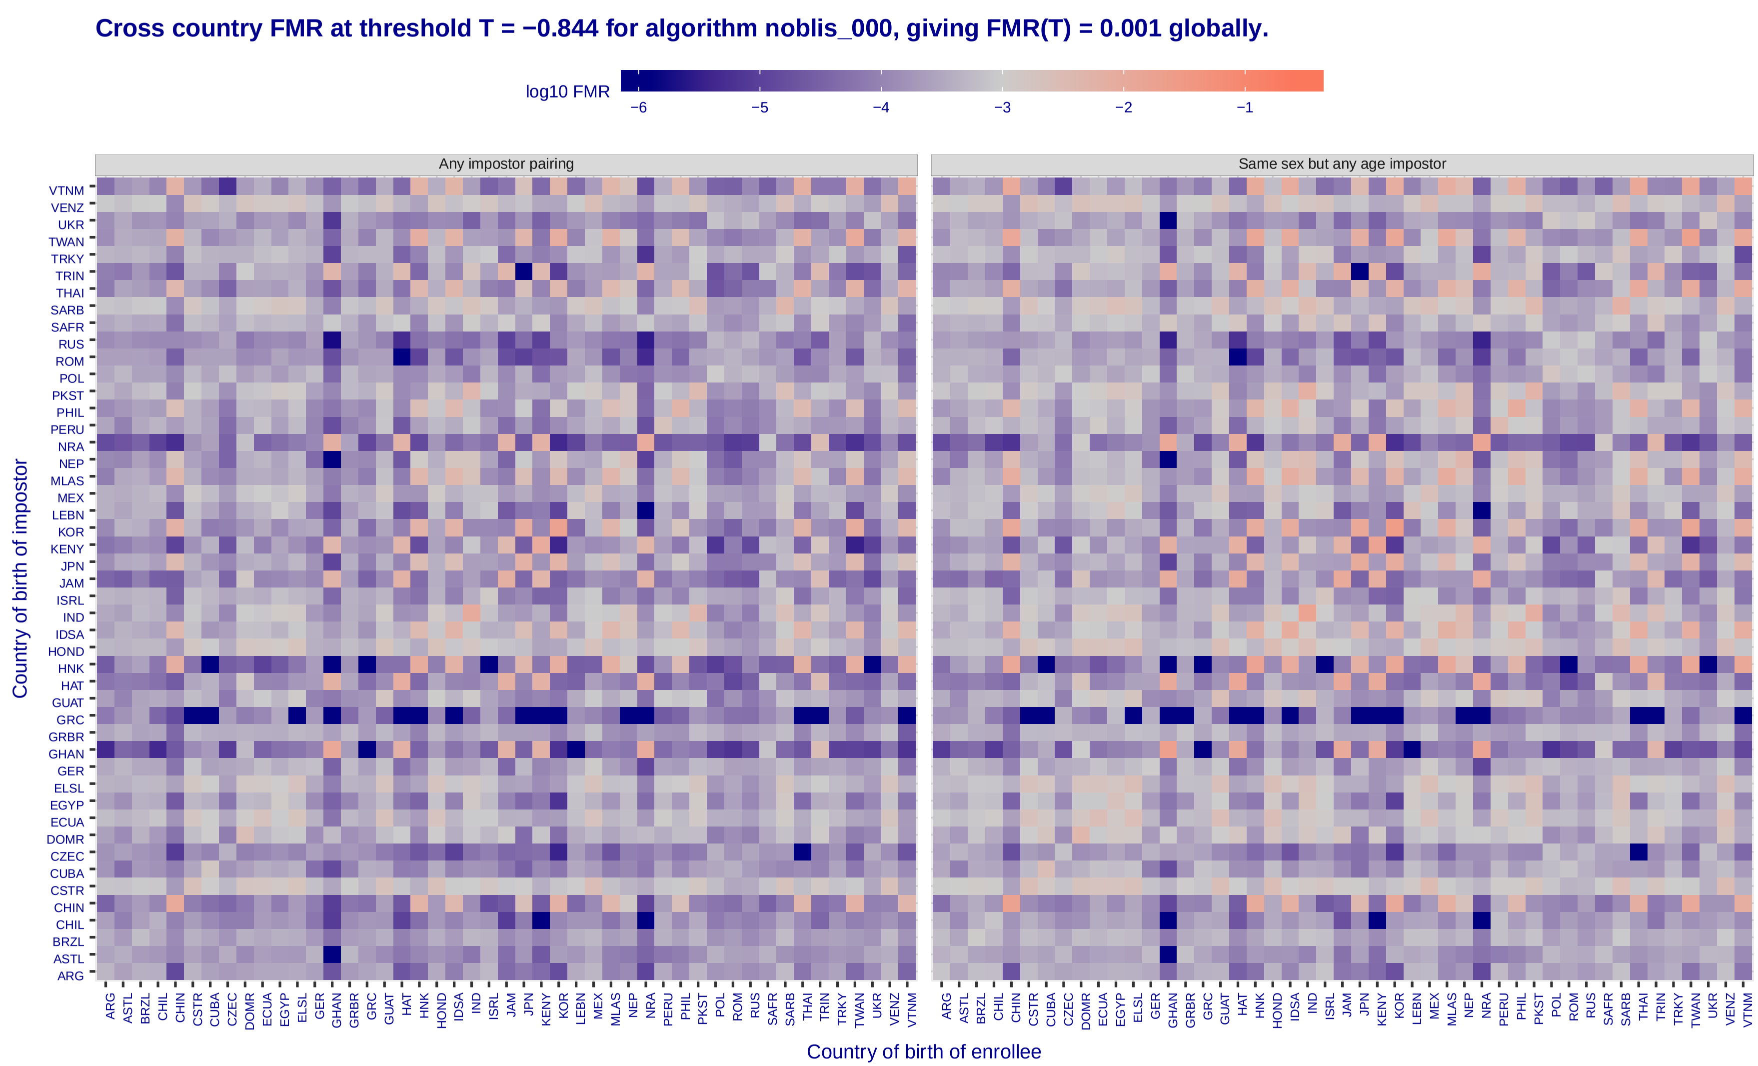

Figure 108: For algorithm noblis-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

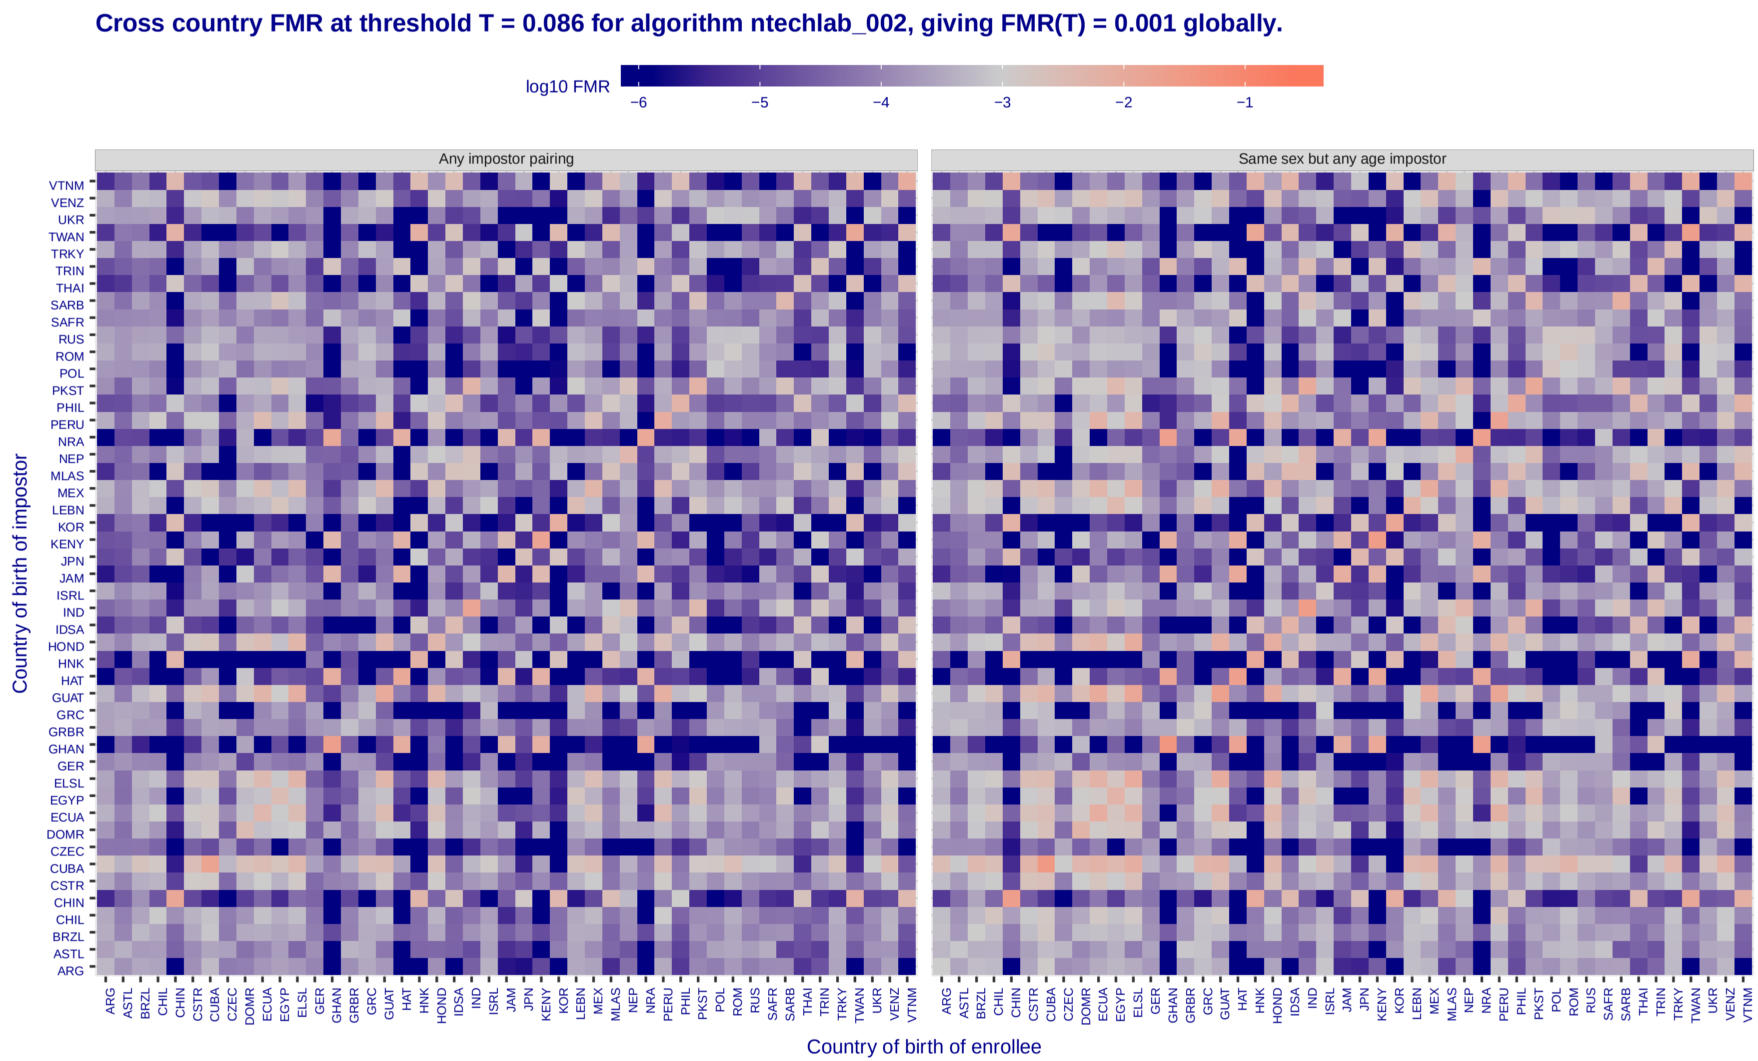

Figure 109: For algorithm ntechlab-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

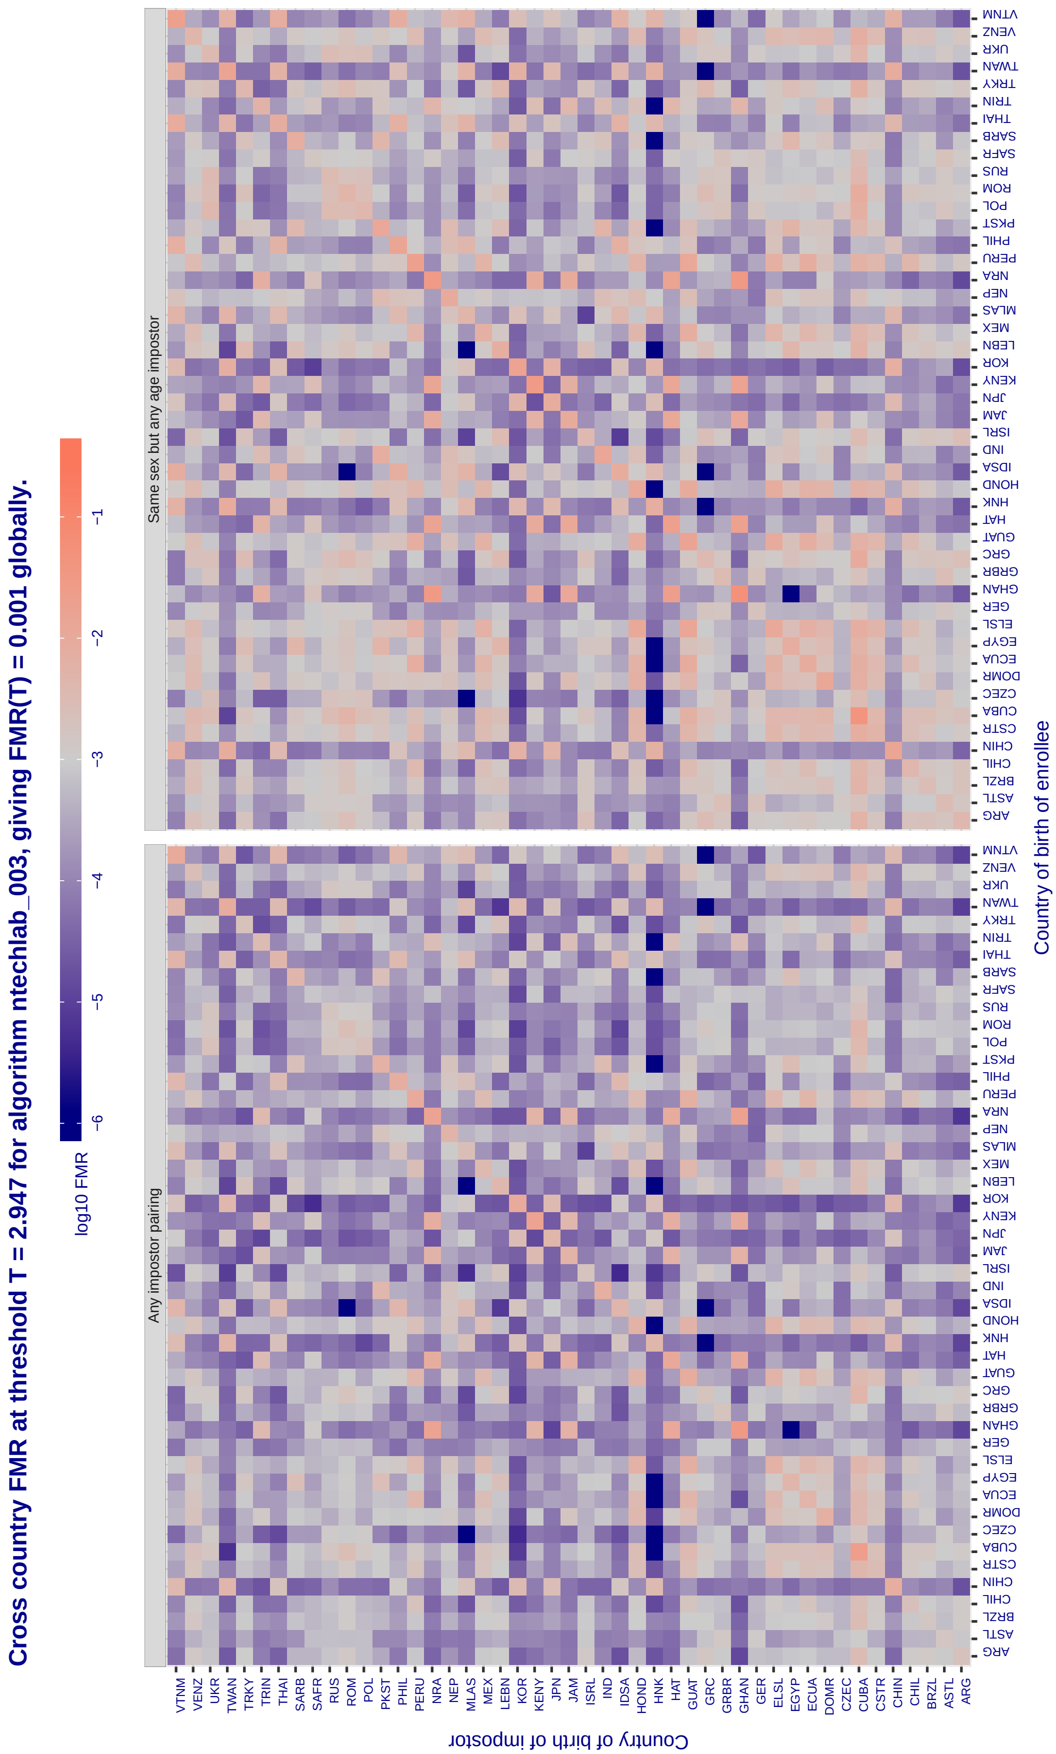

Figure 110: For algorithm ntechlab-003 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

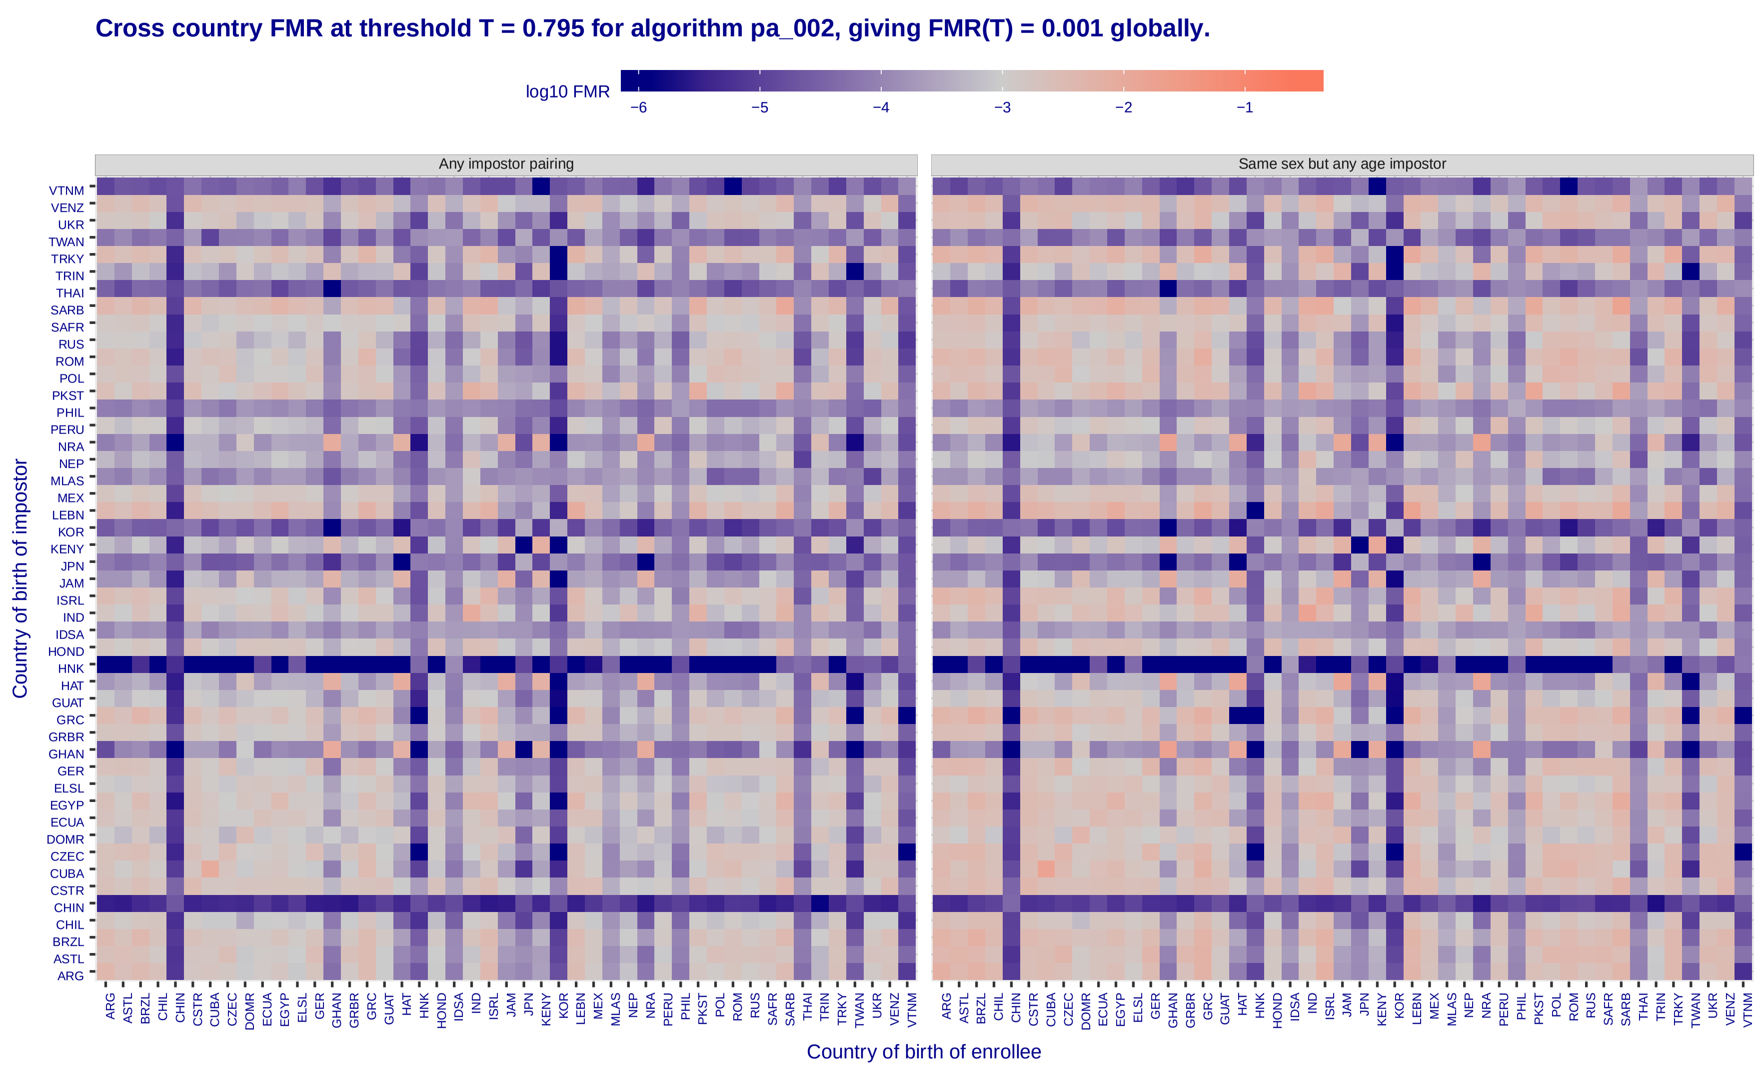

Figure 111: For algorithm pa-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

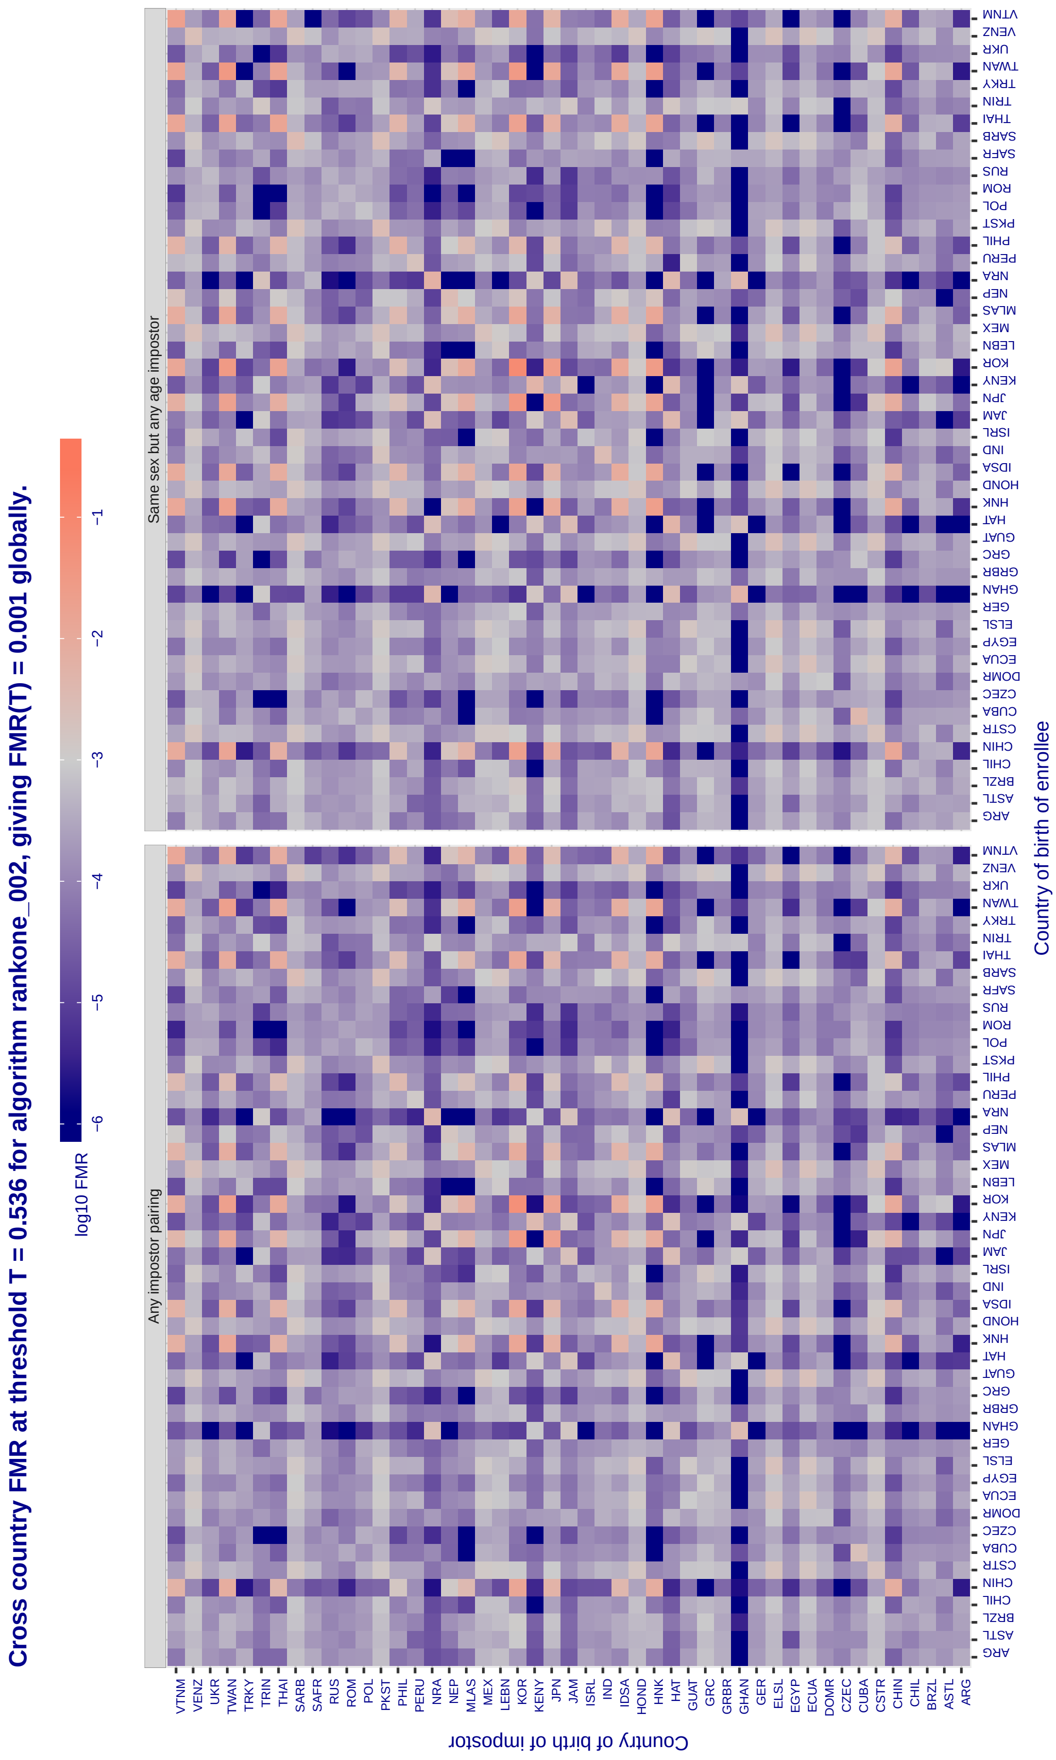

Figure 112: For algorithm rankone-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

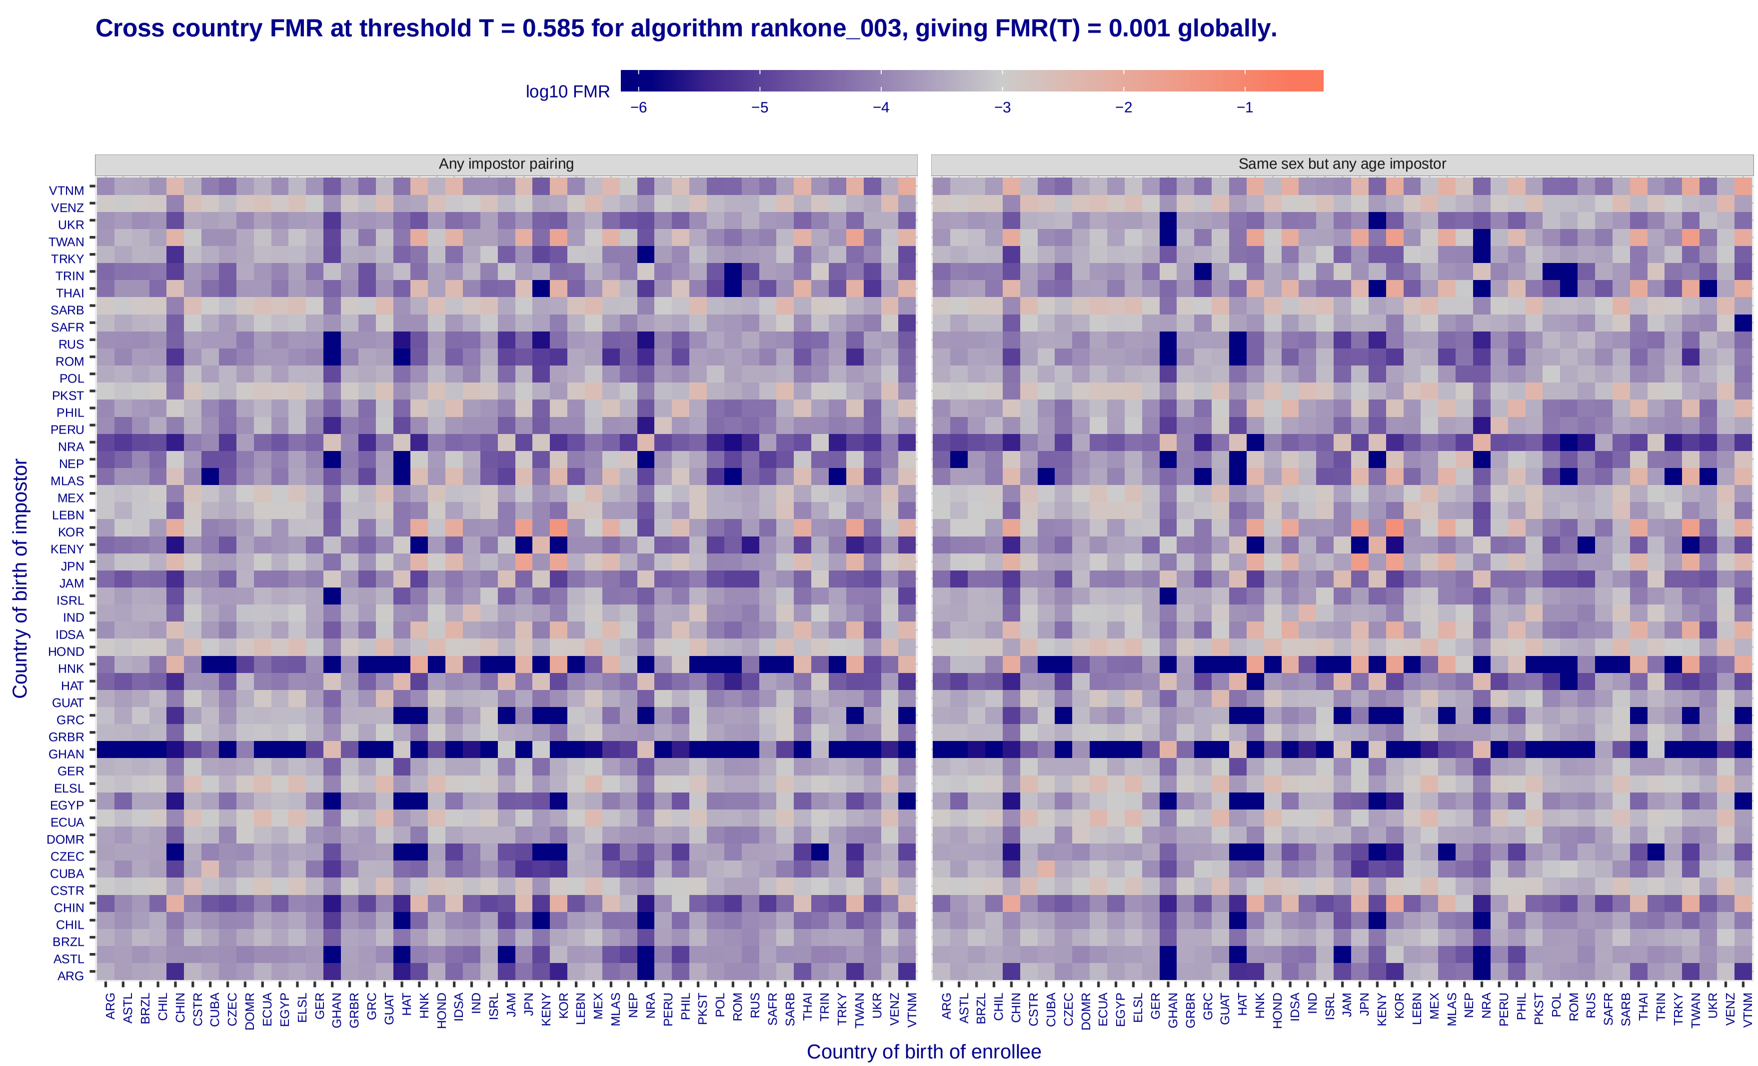

Figure 113: For algorithm rankone-003 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

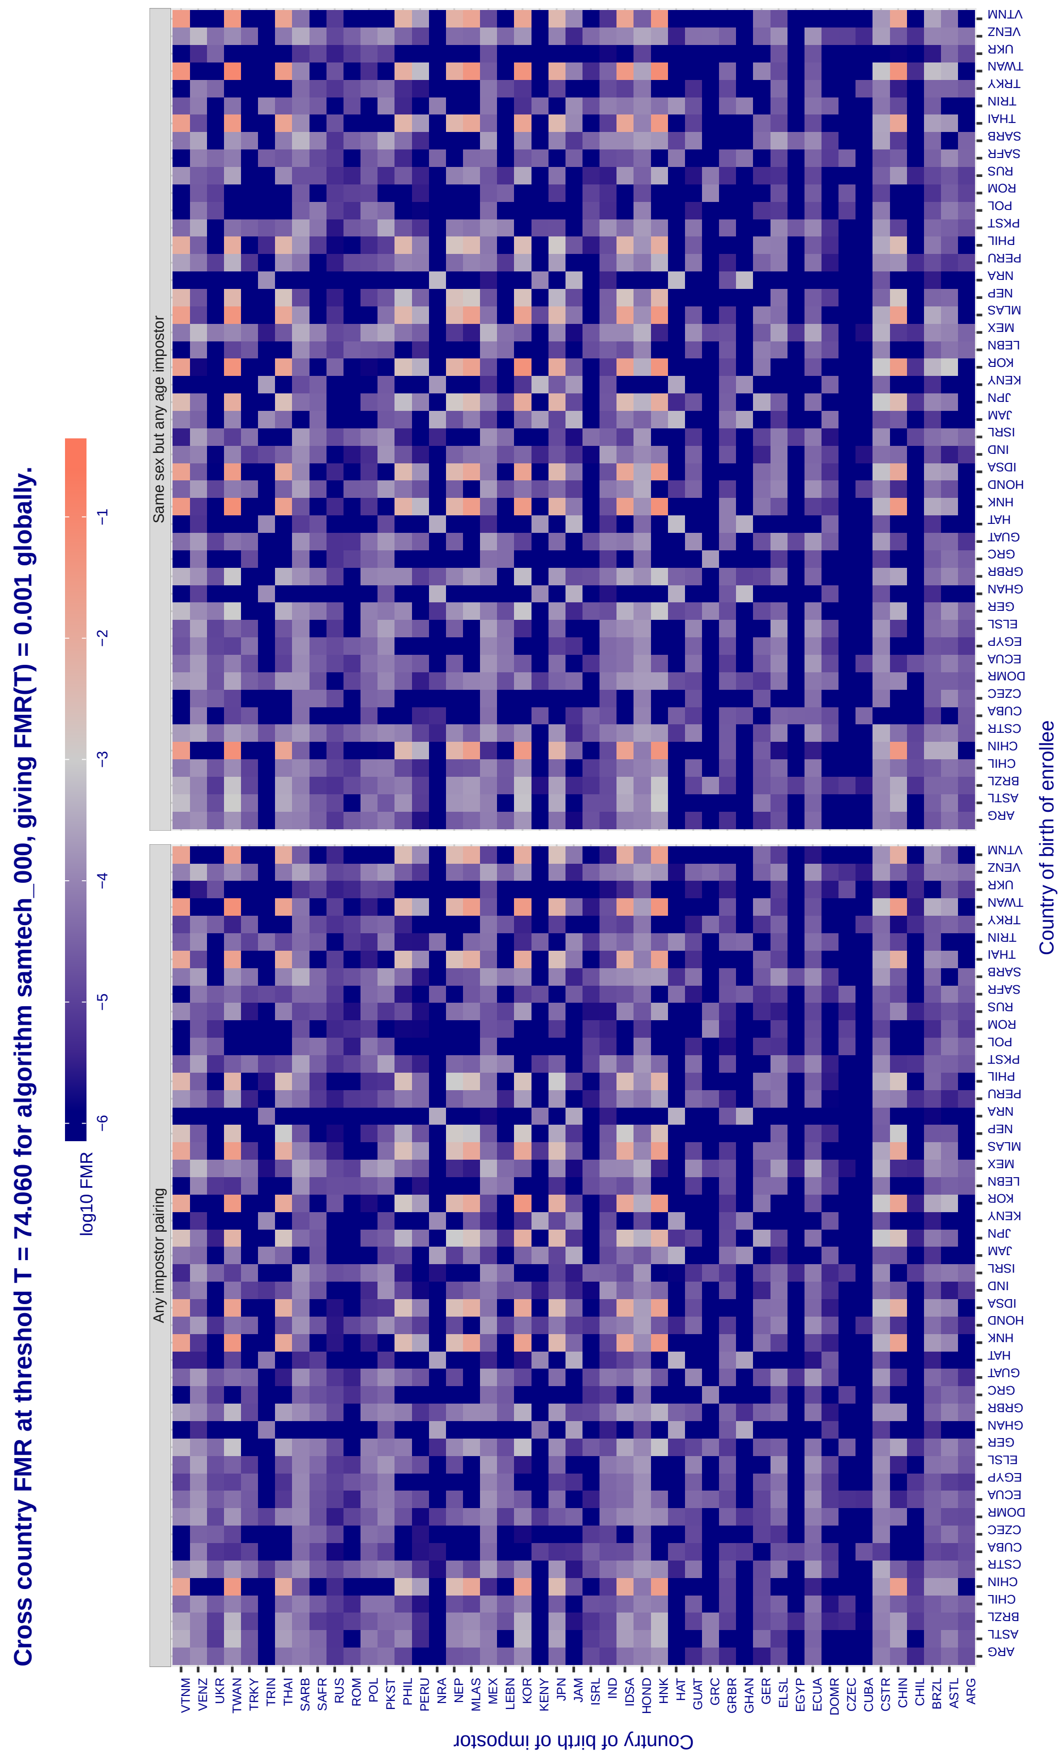

Figure 114: For algorithm samtech-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

Figure 115: For algorithm shaman-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

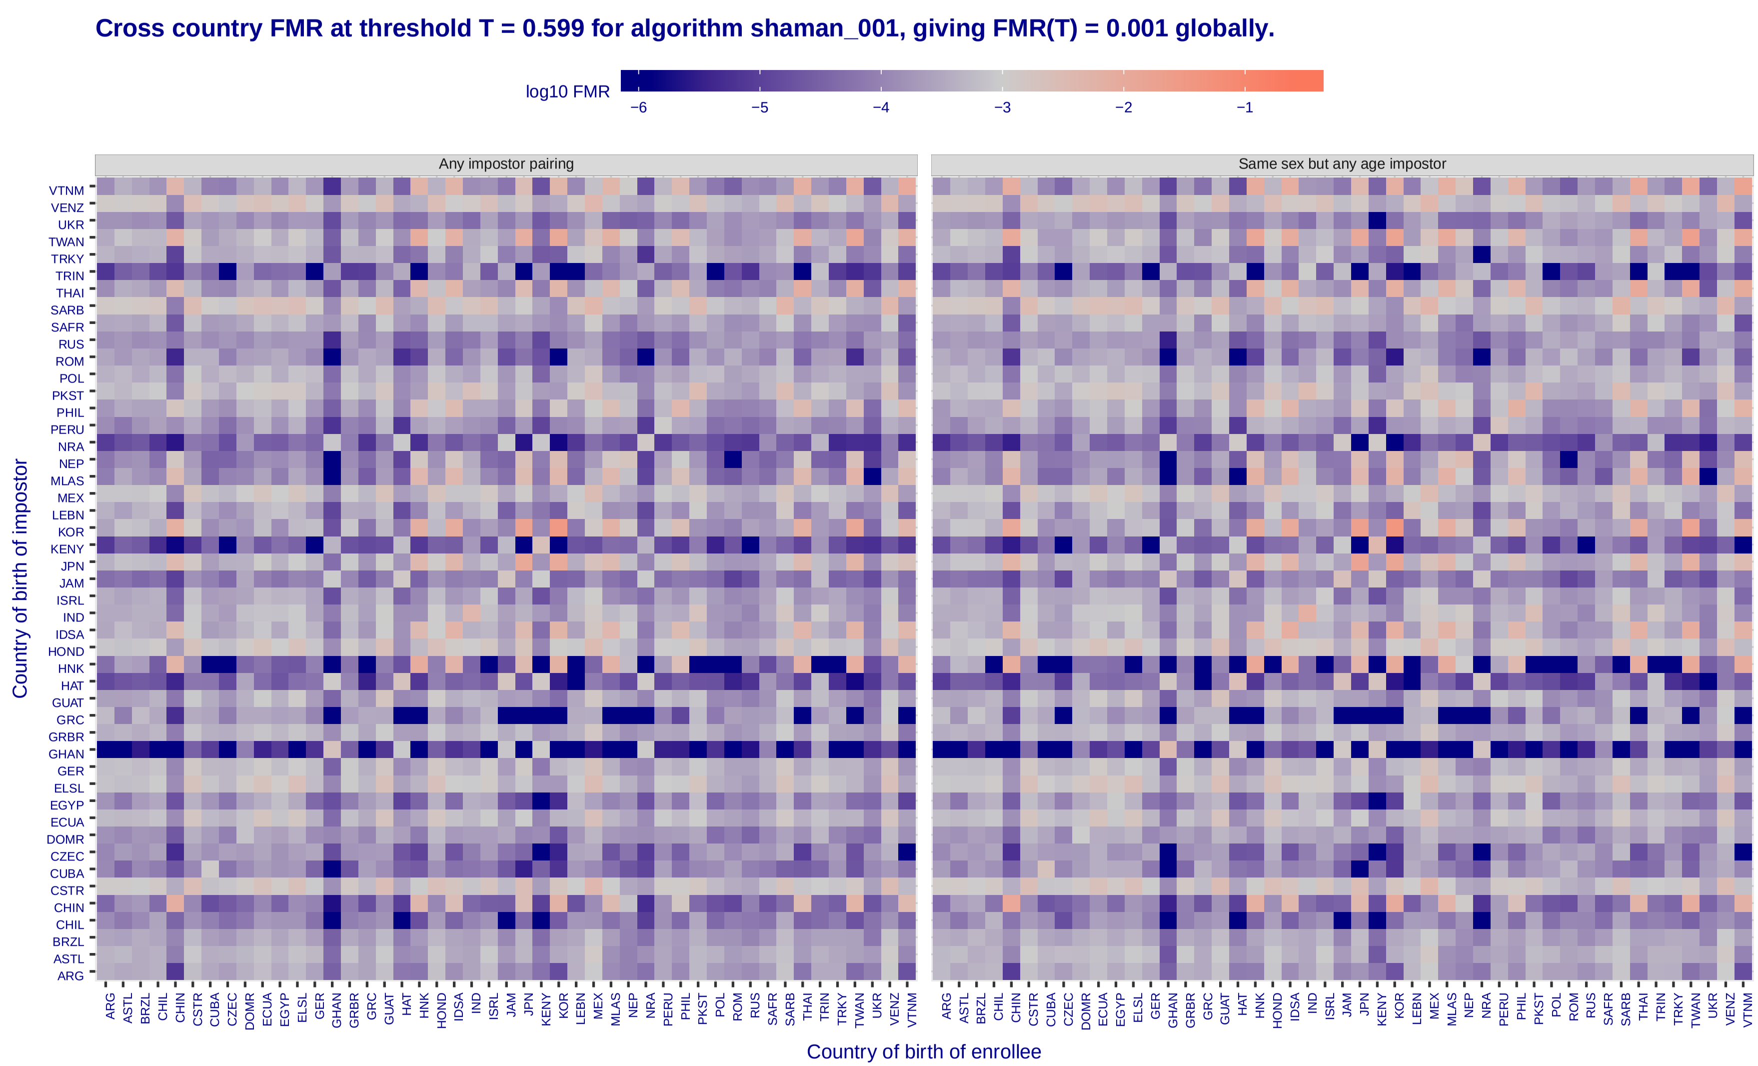

Figure 116: For algorithm shaman-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

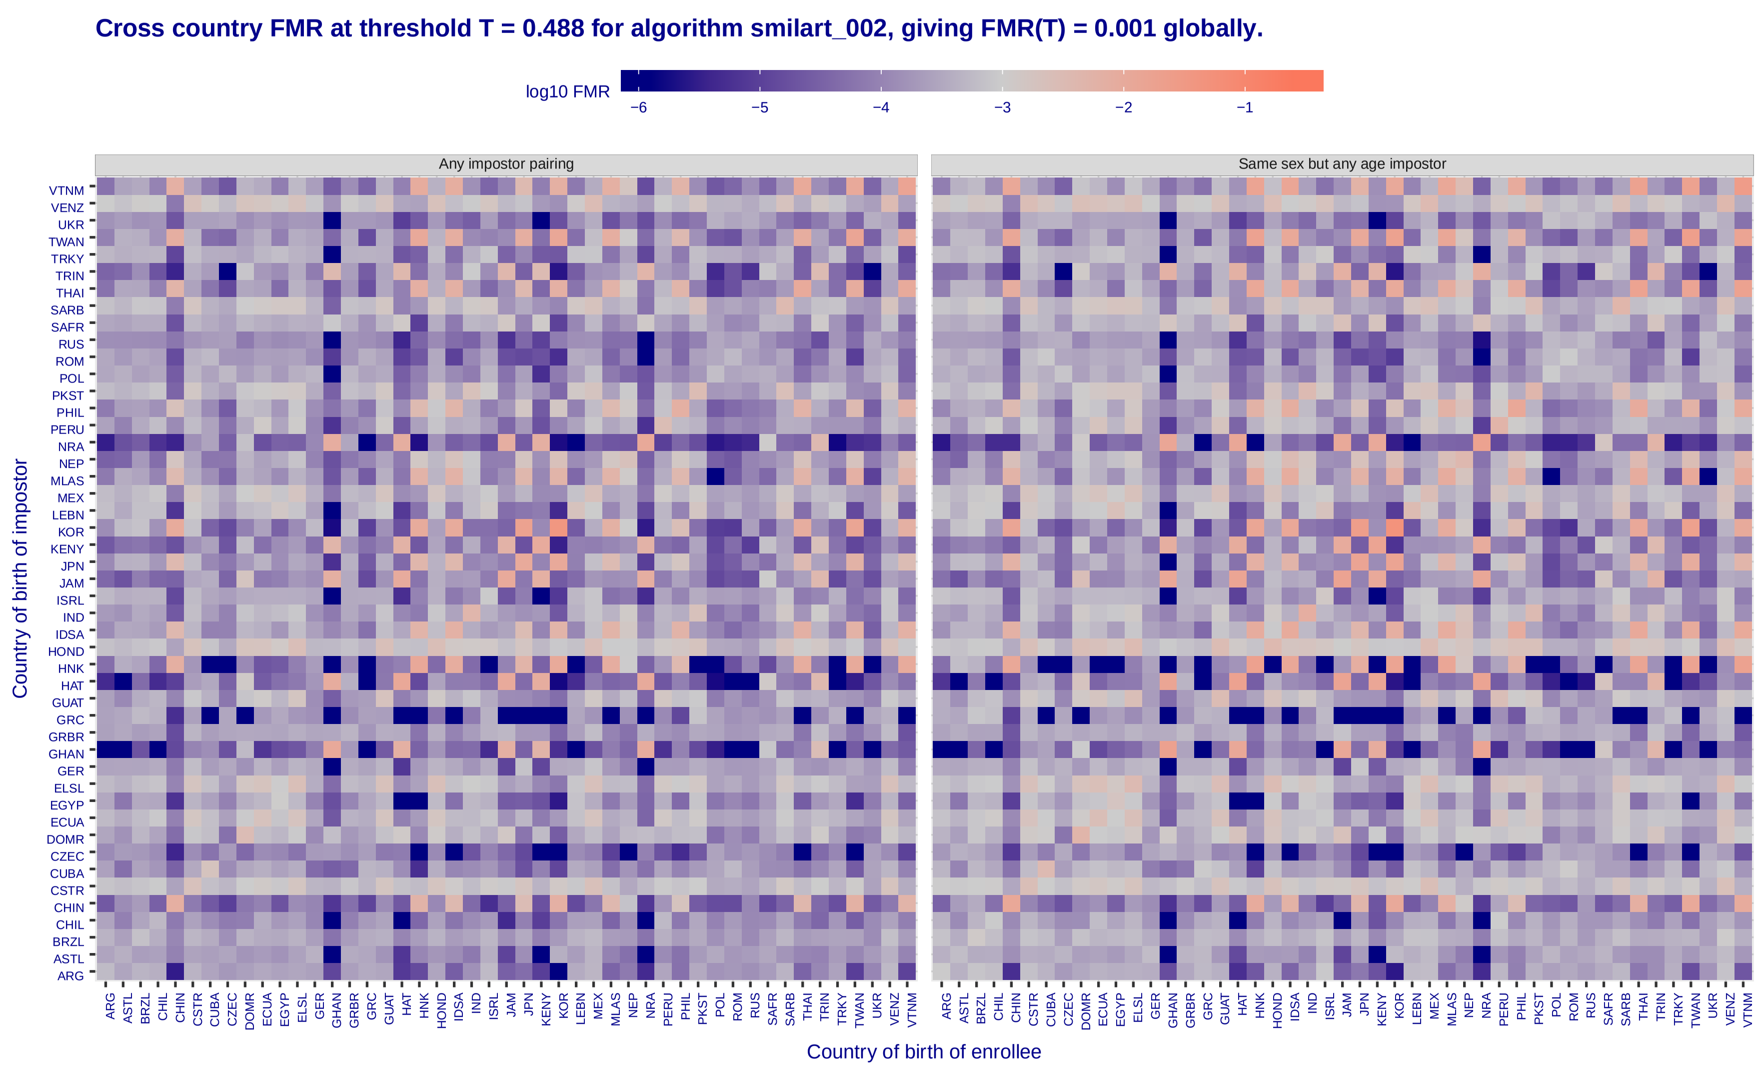

Figure 117: For algorithm smilart-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

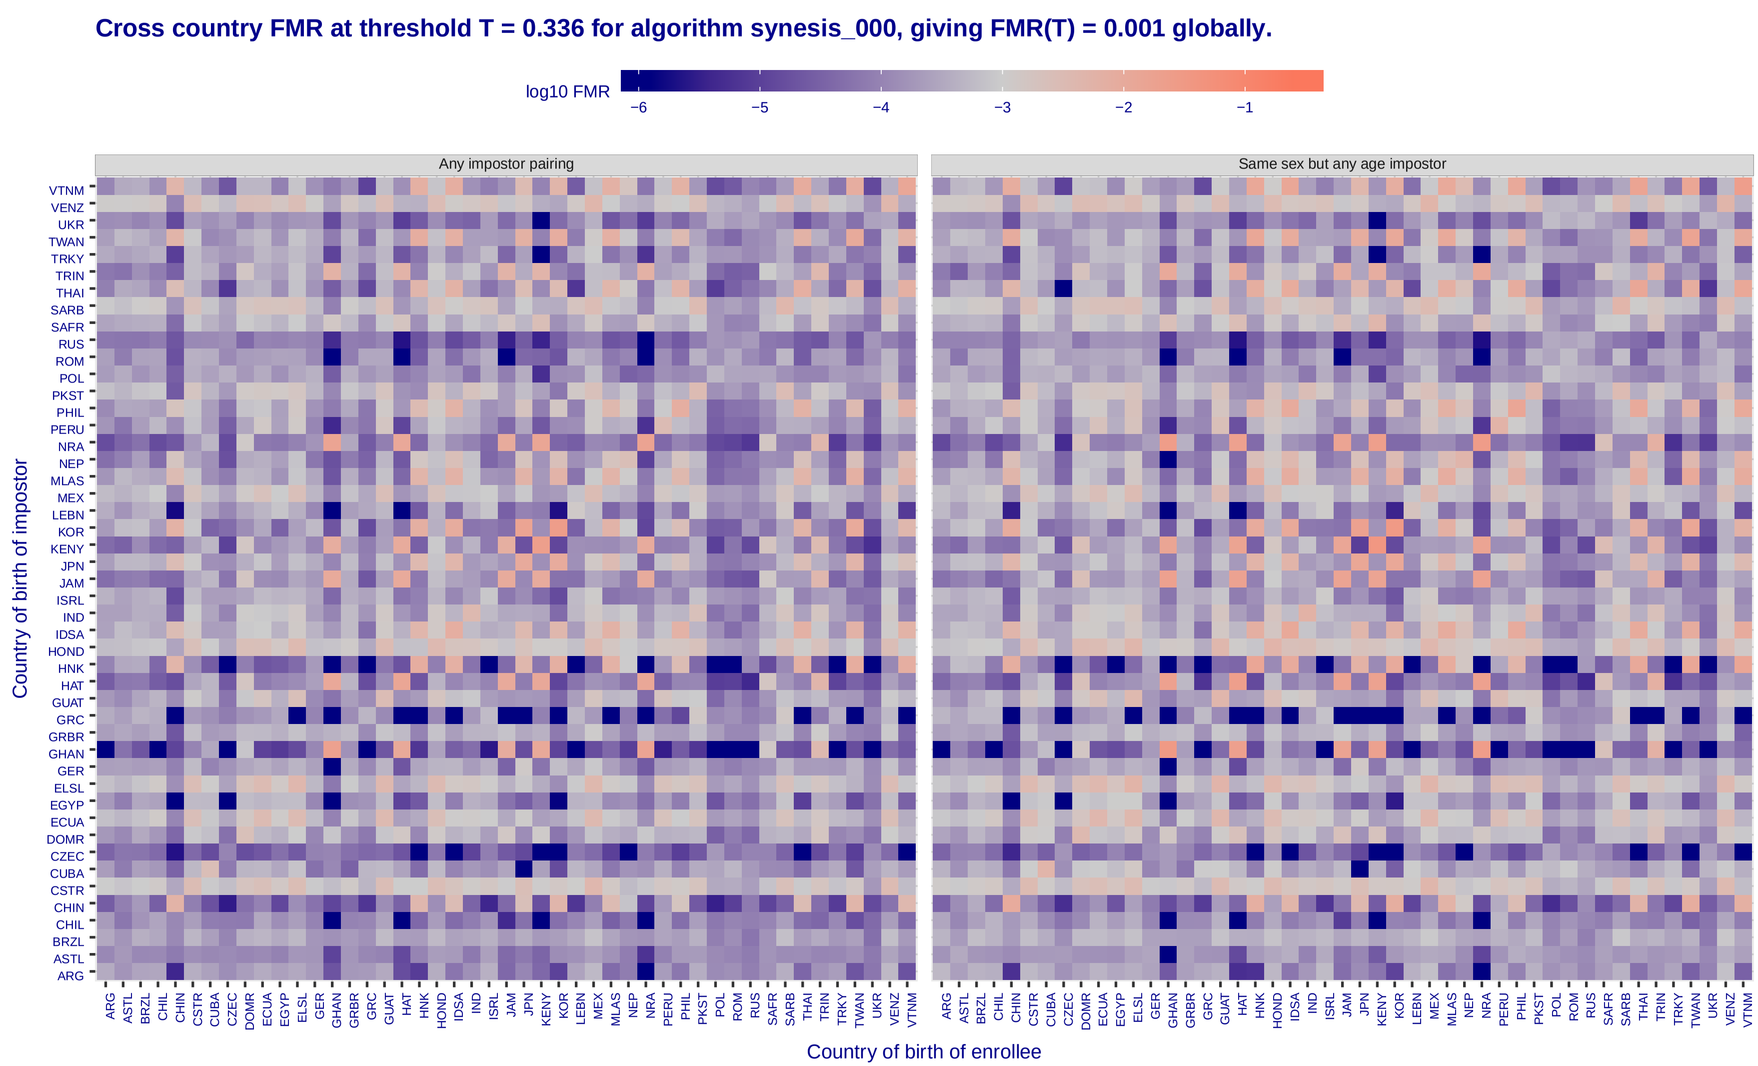

Figure 118: For algorithm synesis-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

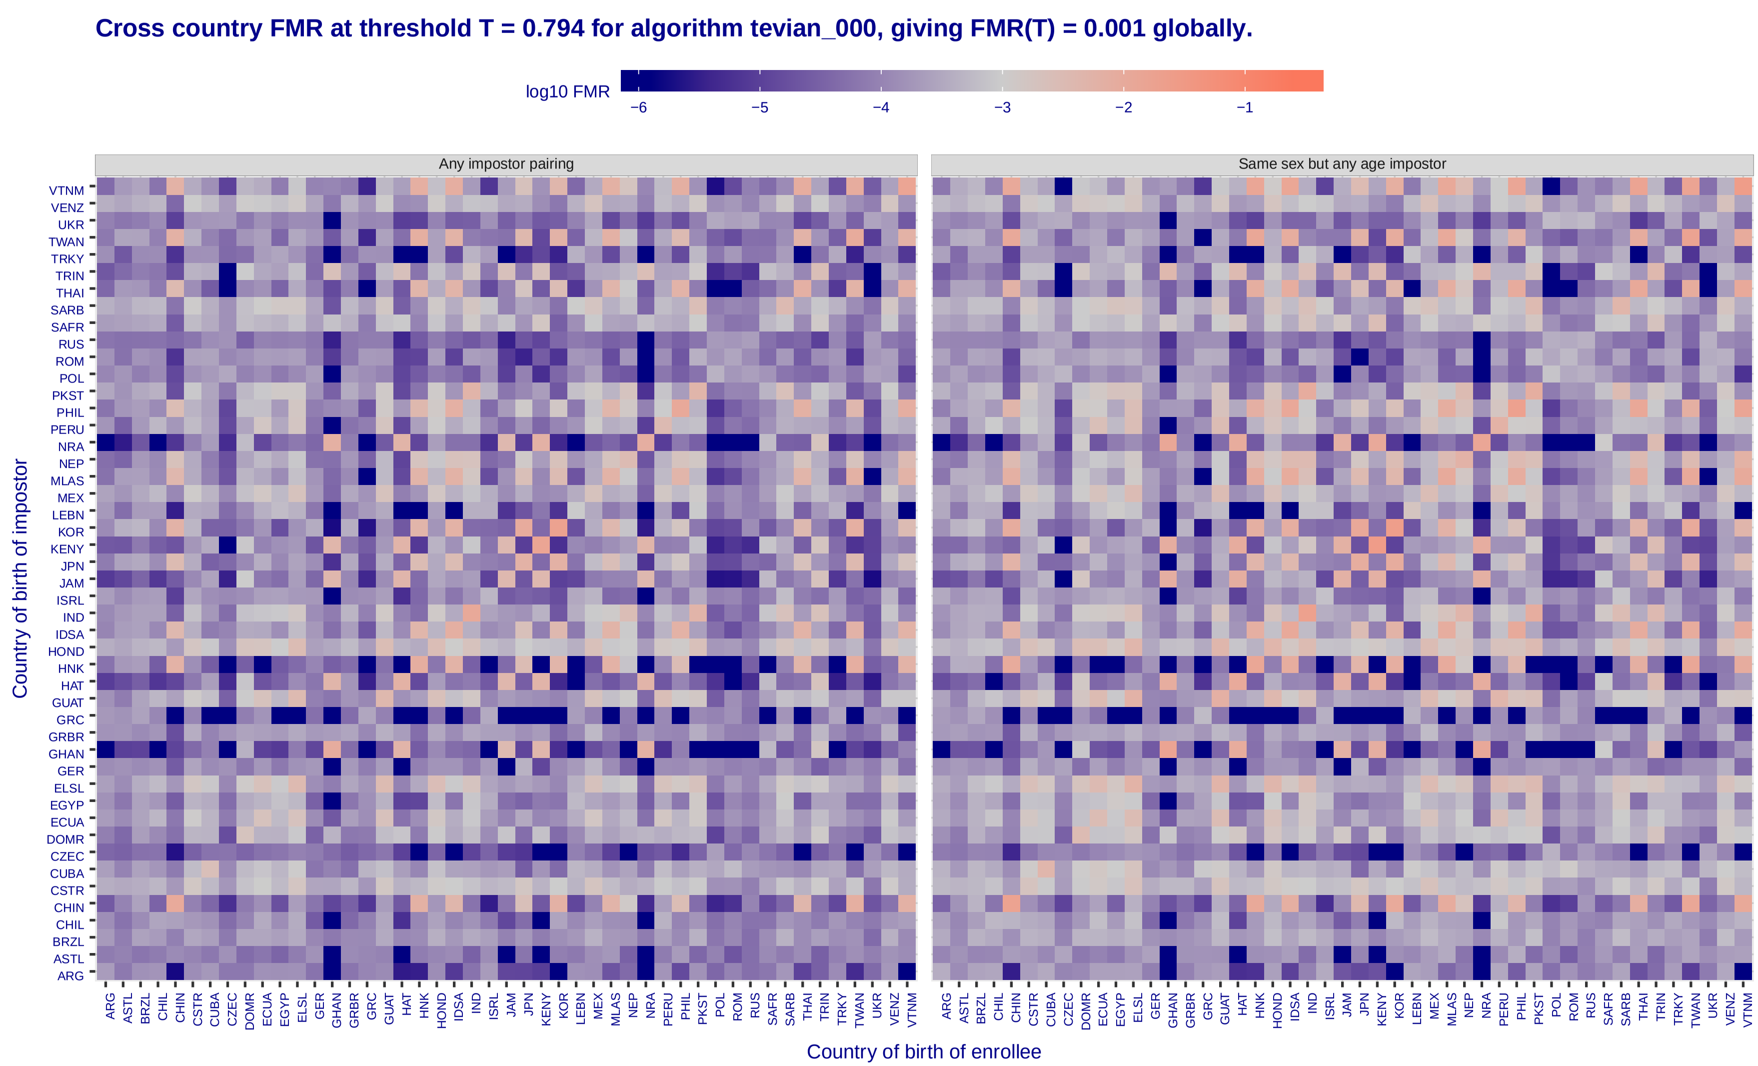

Figure 119: For algorithm tevian-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

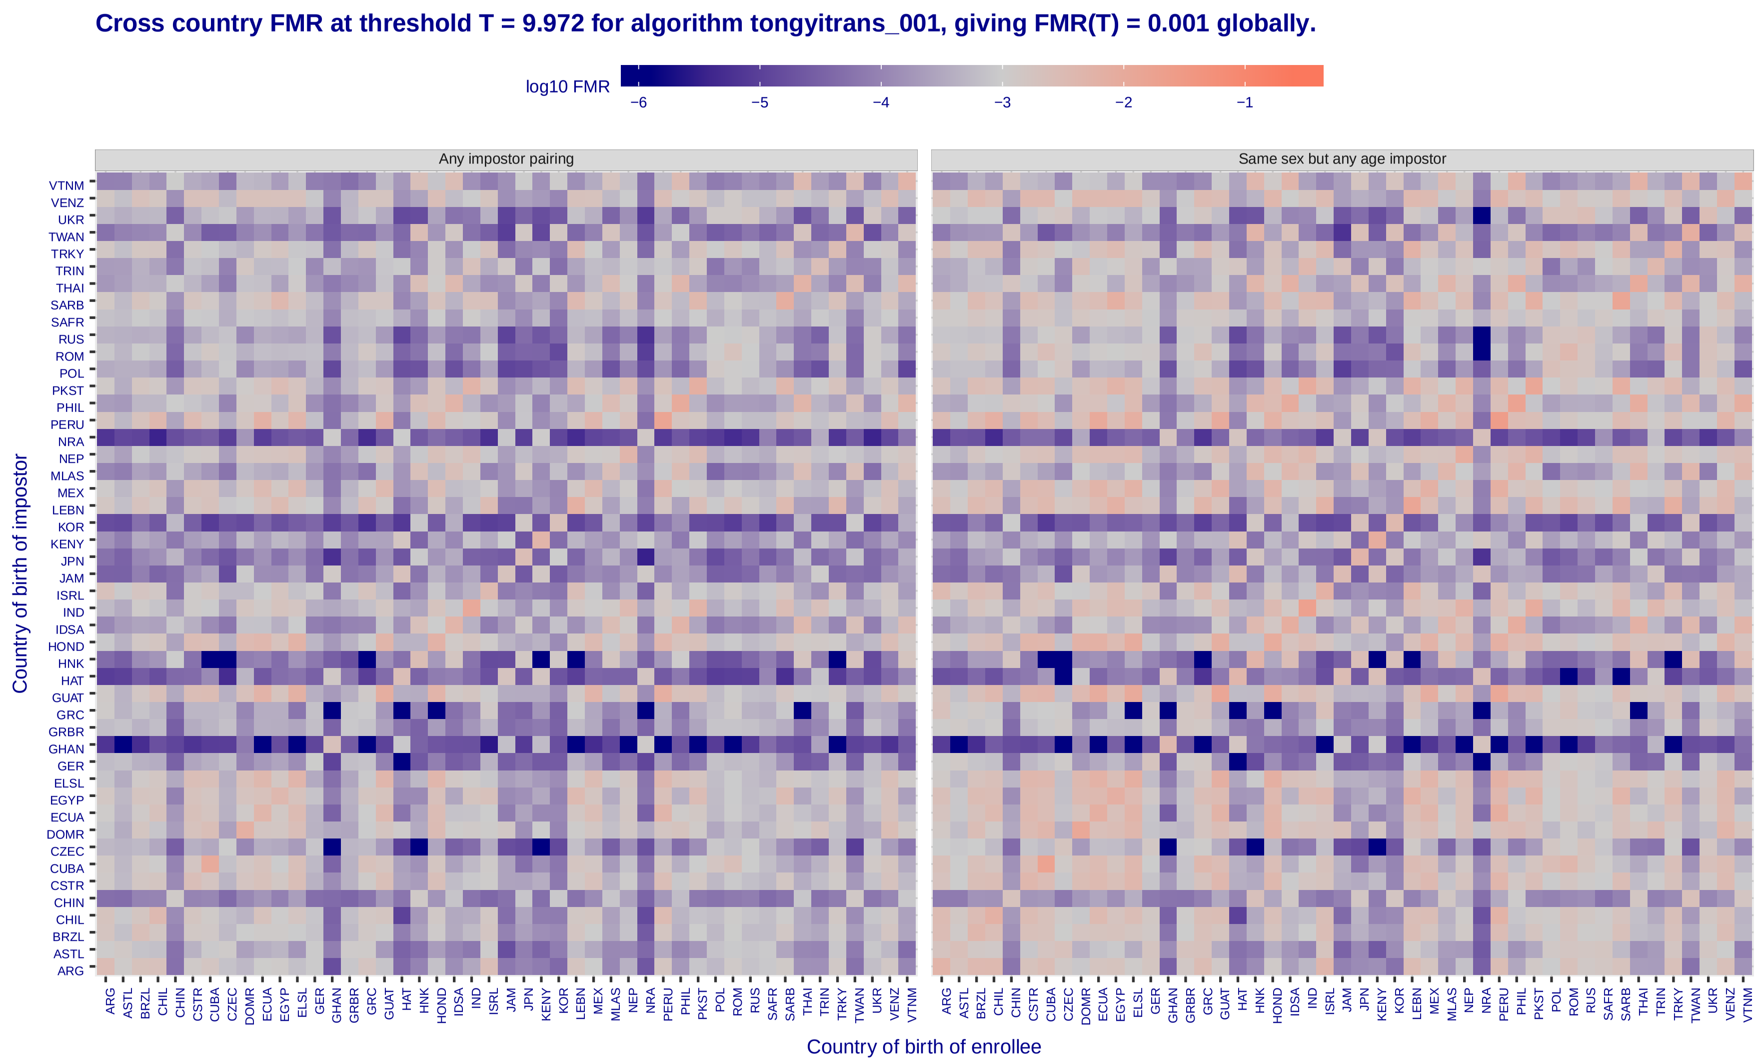

Figure 120: For algorithm tongyitrans-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

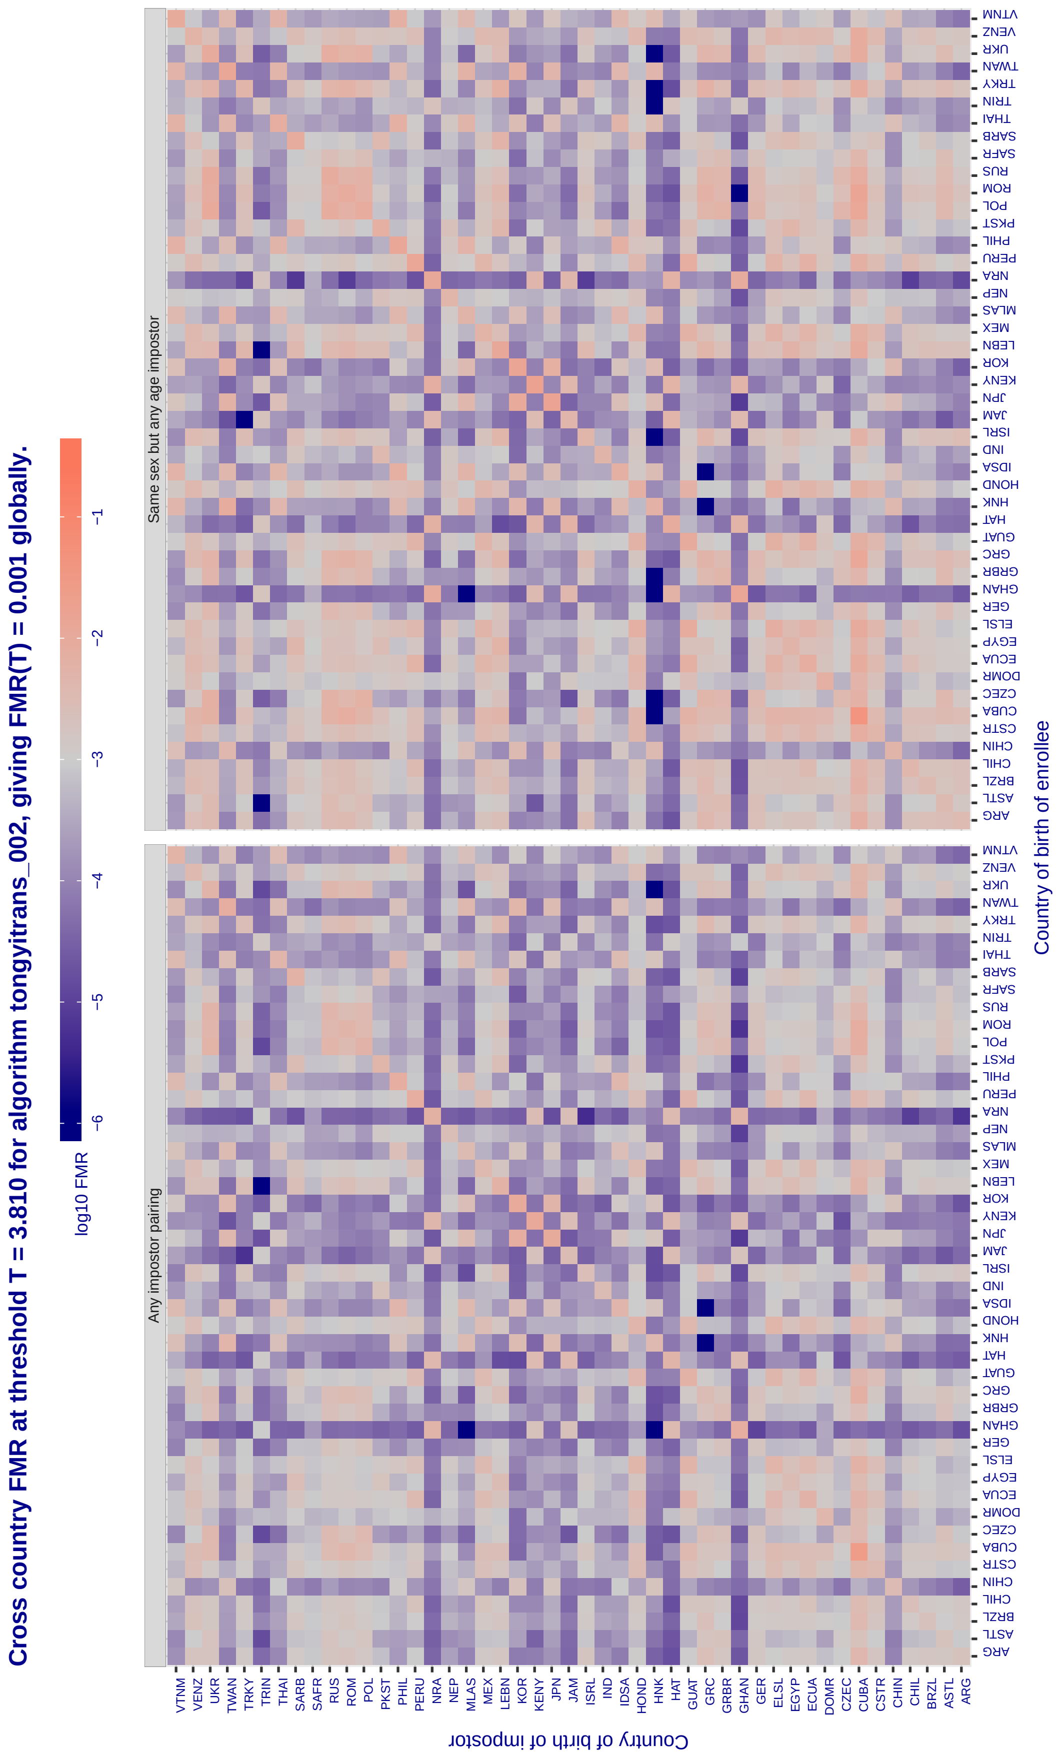

Figure 121: For algorithm tongyitrans-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

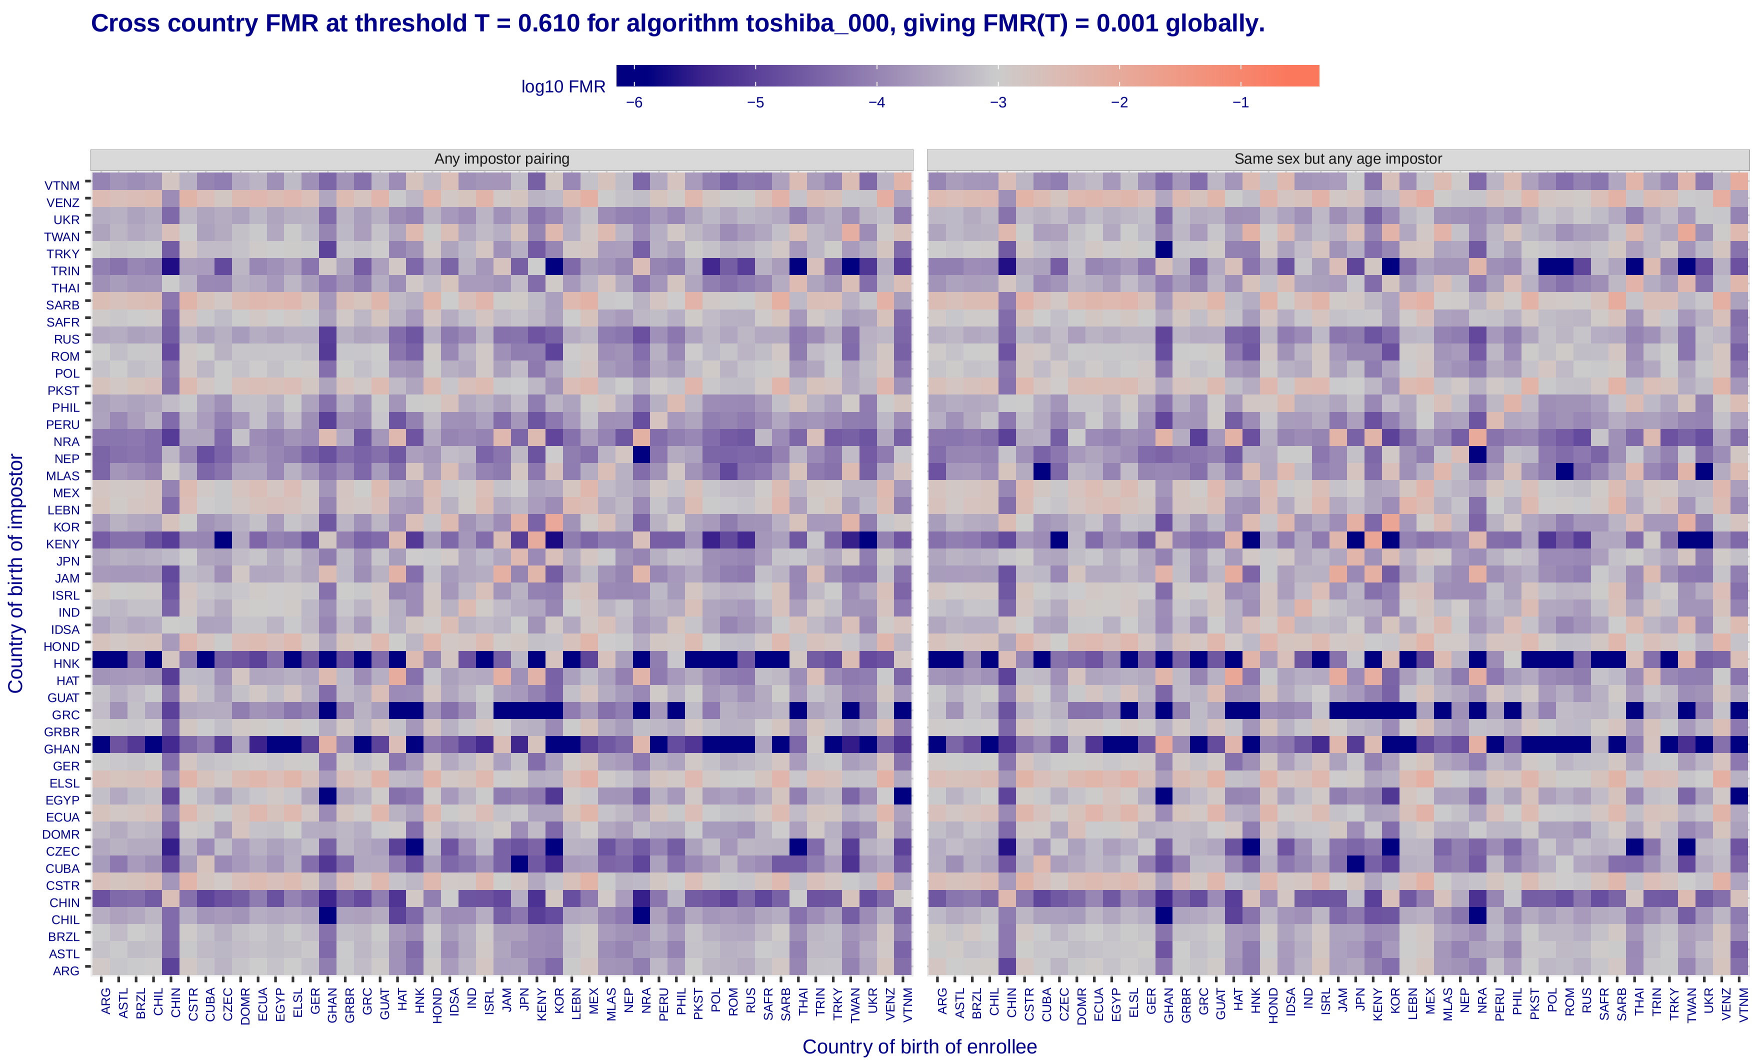

Figure 122: For algorithm toshiba-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

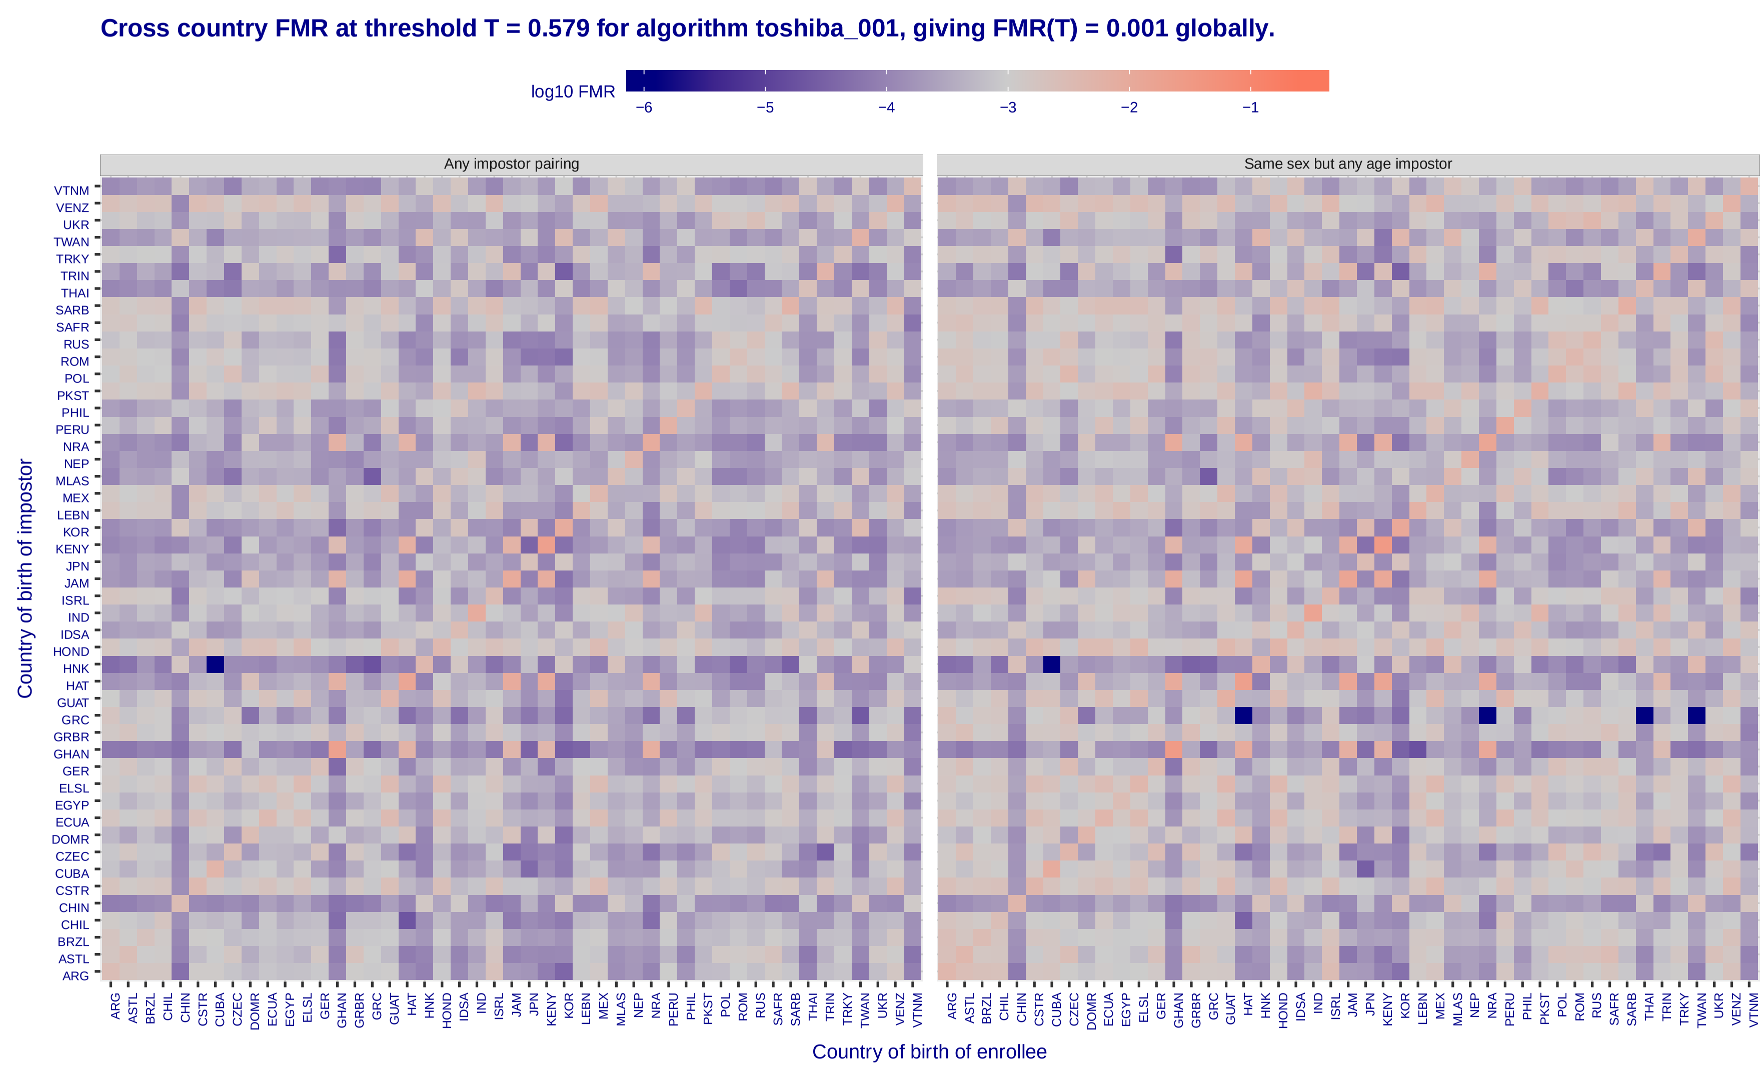

Figure 123: *For algorithm toshiba-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.*

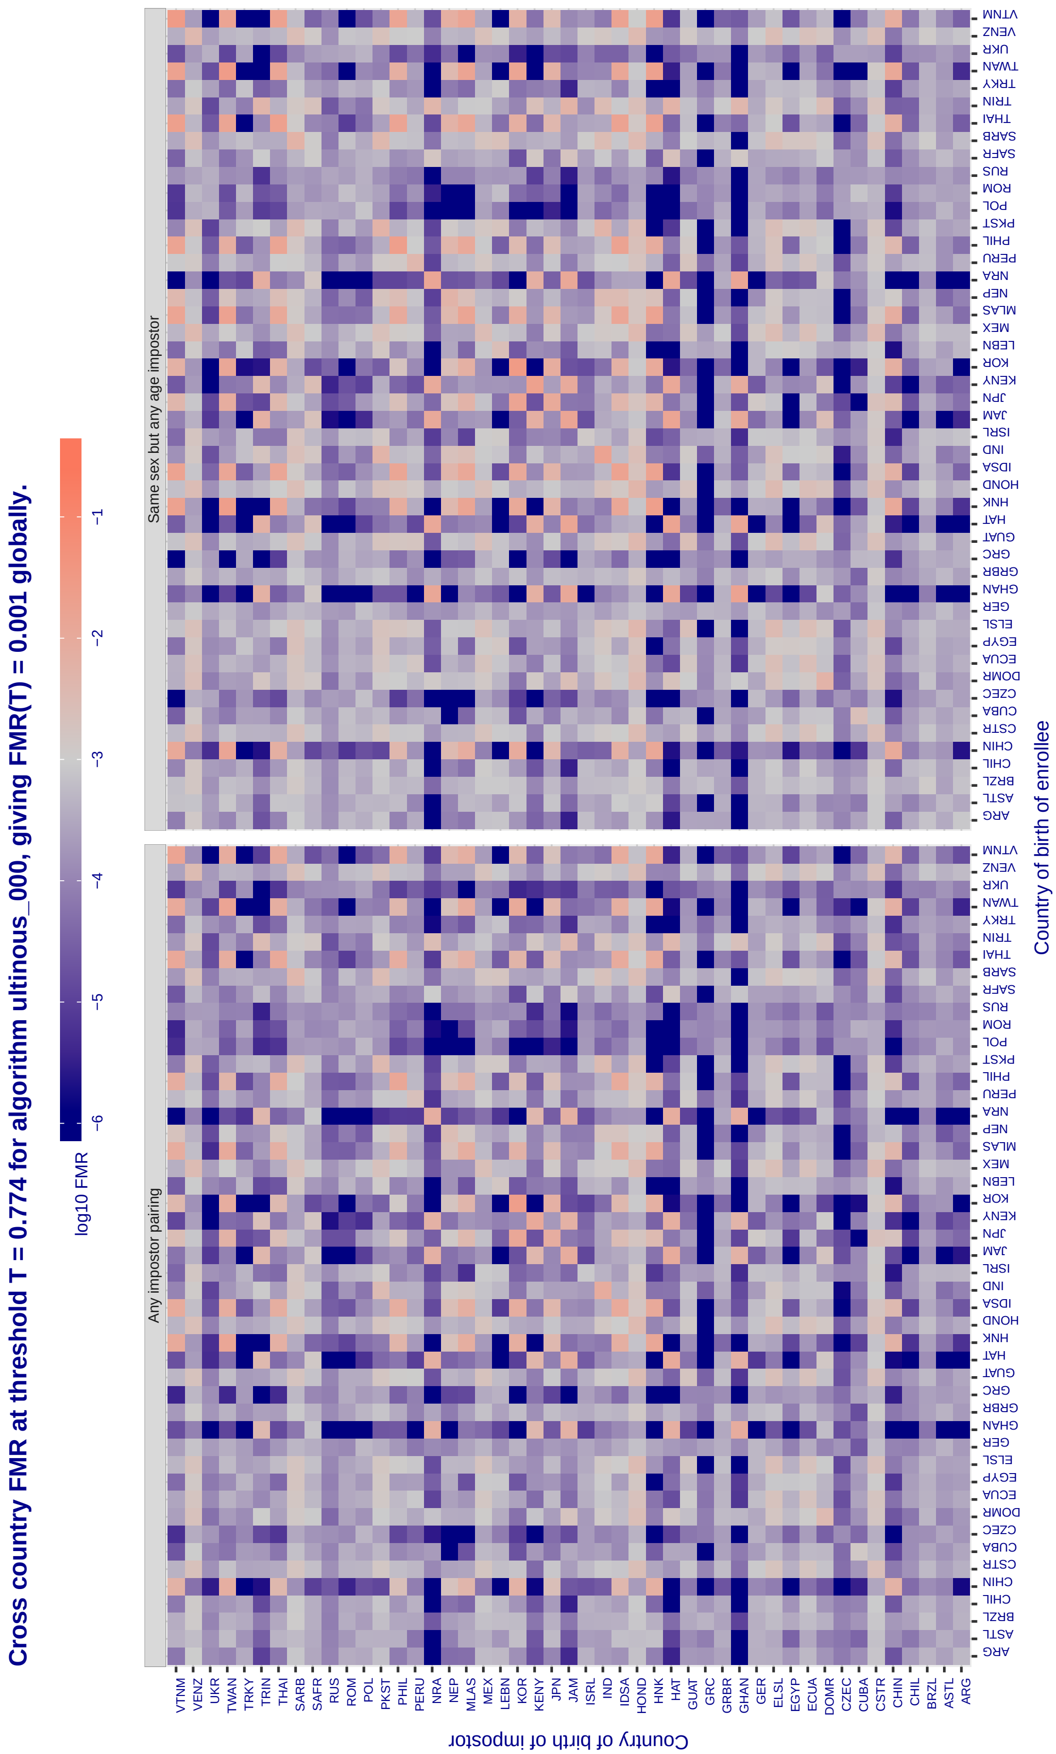

Figure 124: For algorithm ultinous-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

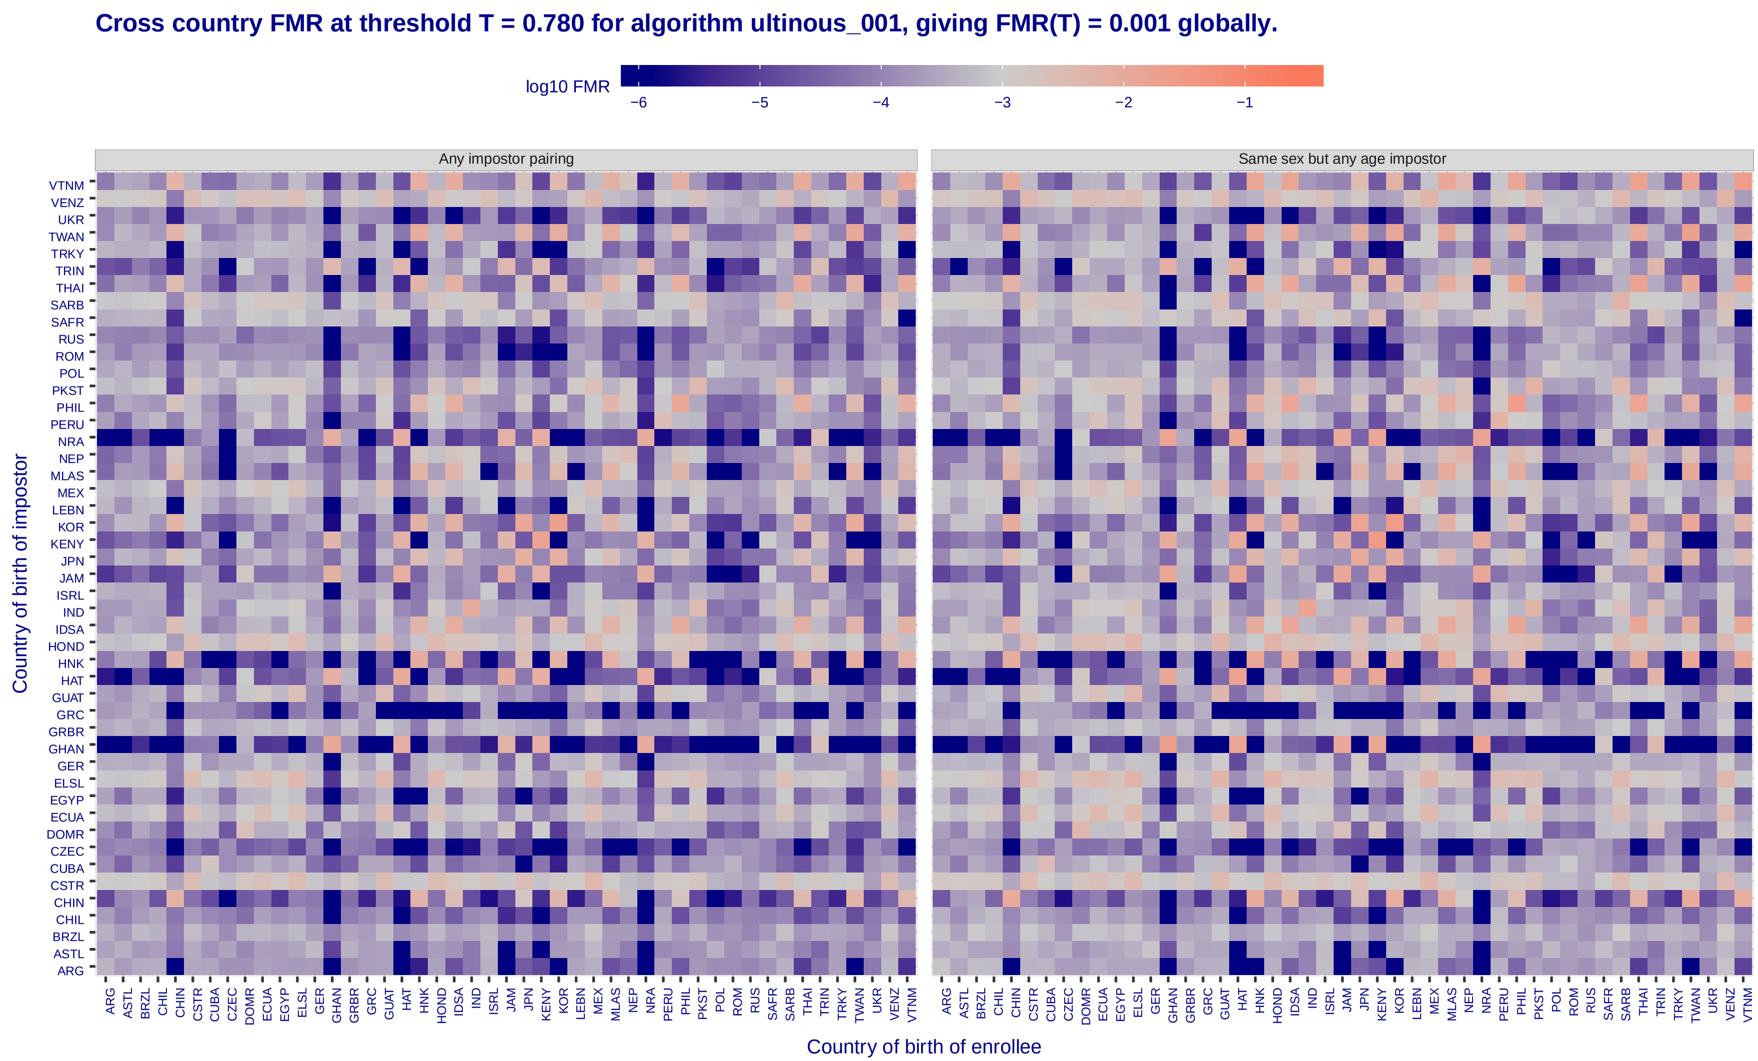

Figure 125: For algorithm ultinous-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

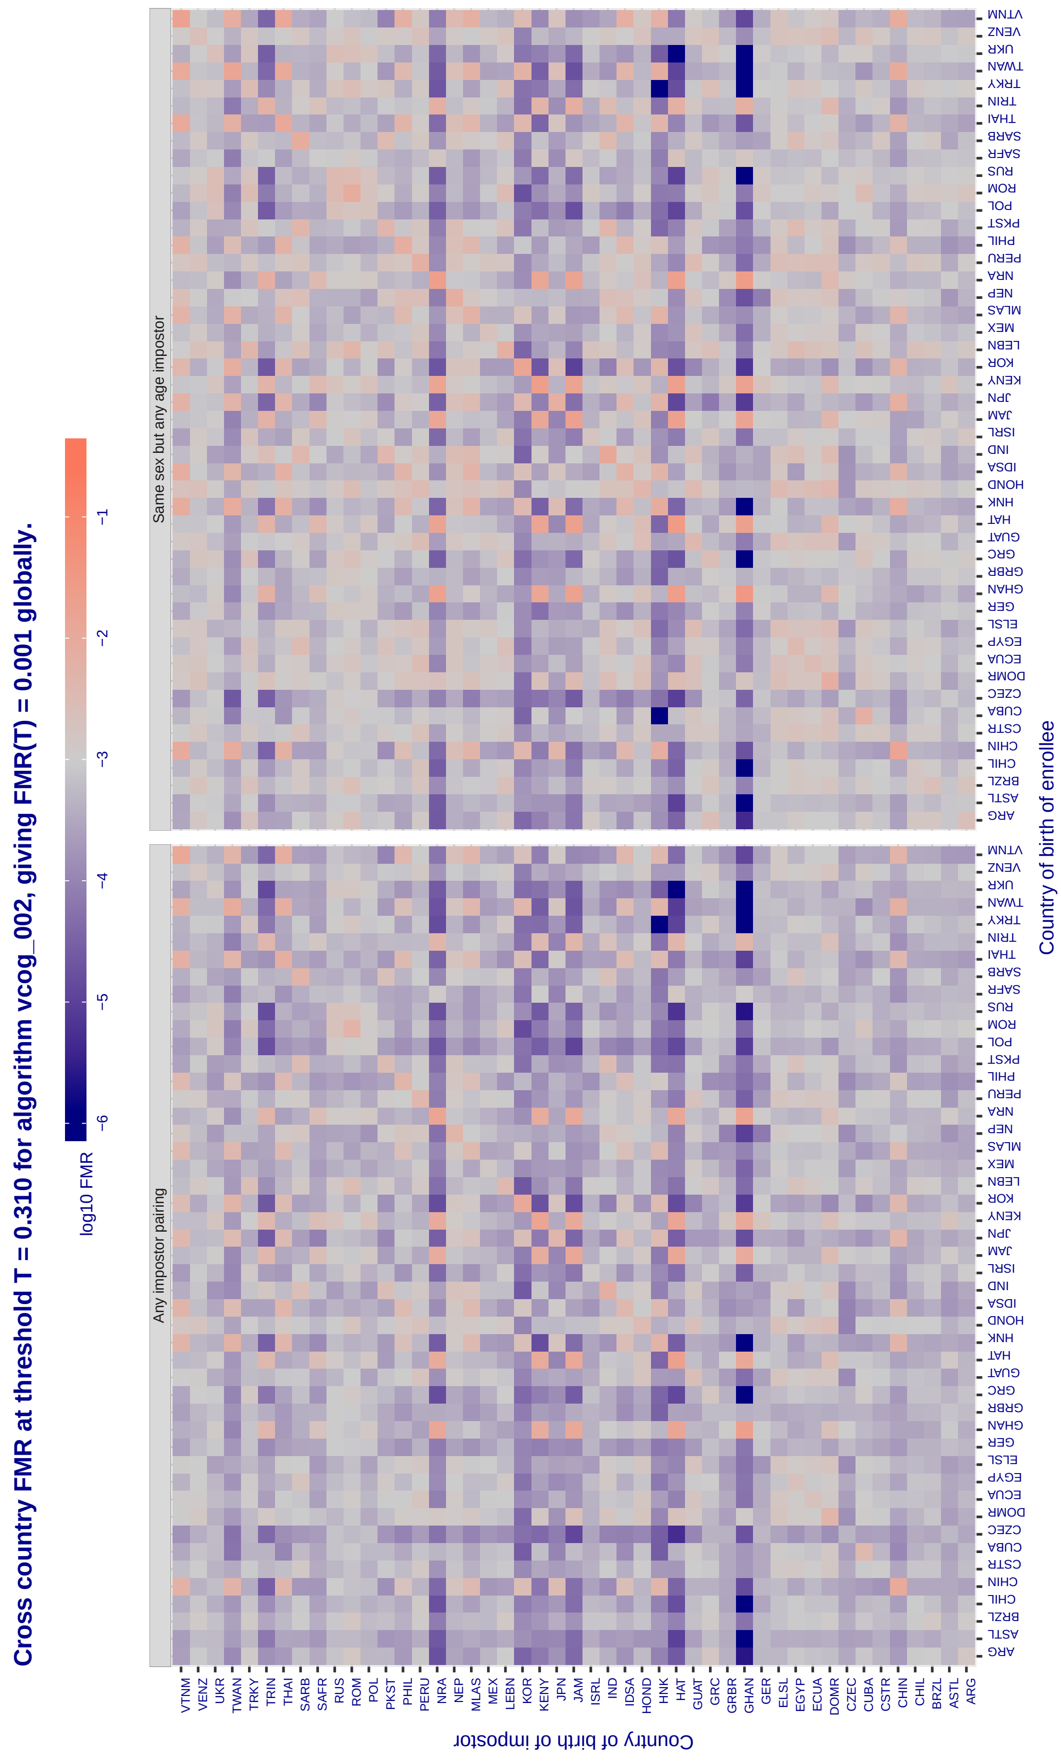

Figure 126: For algorithm vcog-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

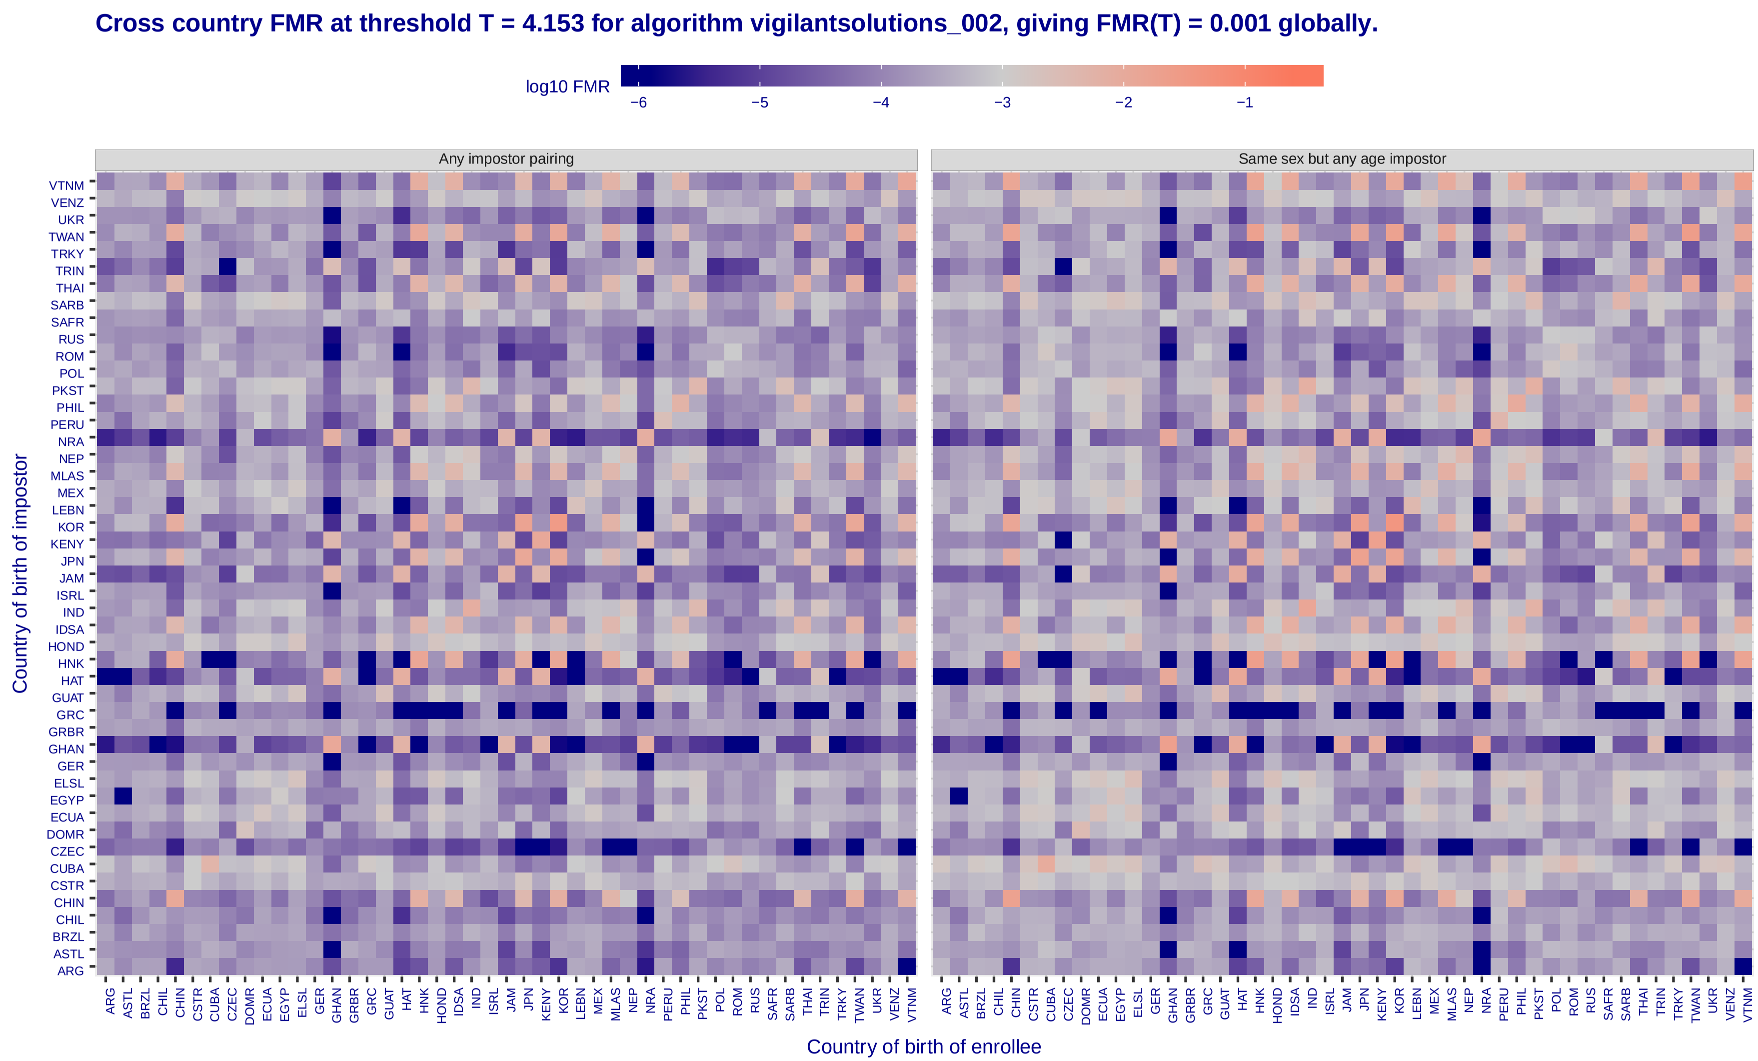

Figure 127: For algorithm vigilantsolutions-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

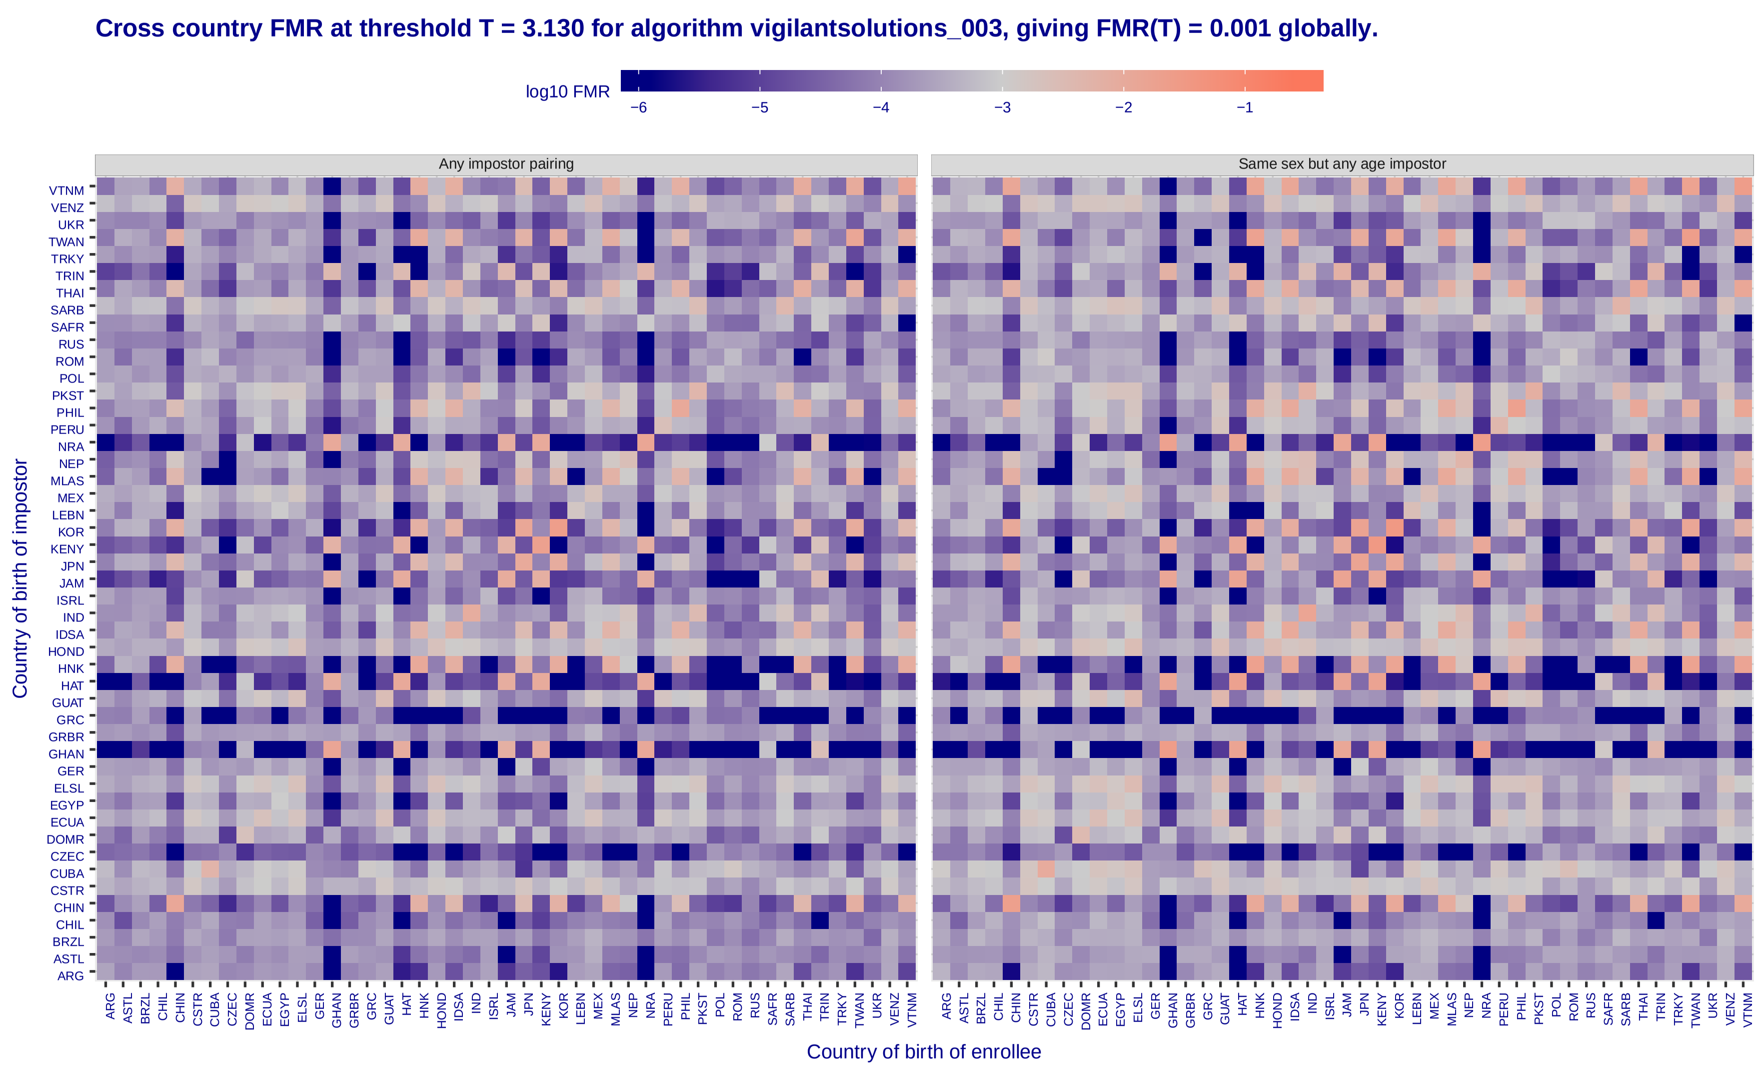

Figure 128: For algorithm vigilantsolutions-003 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

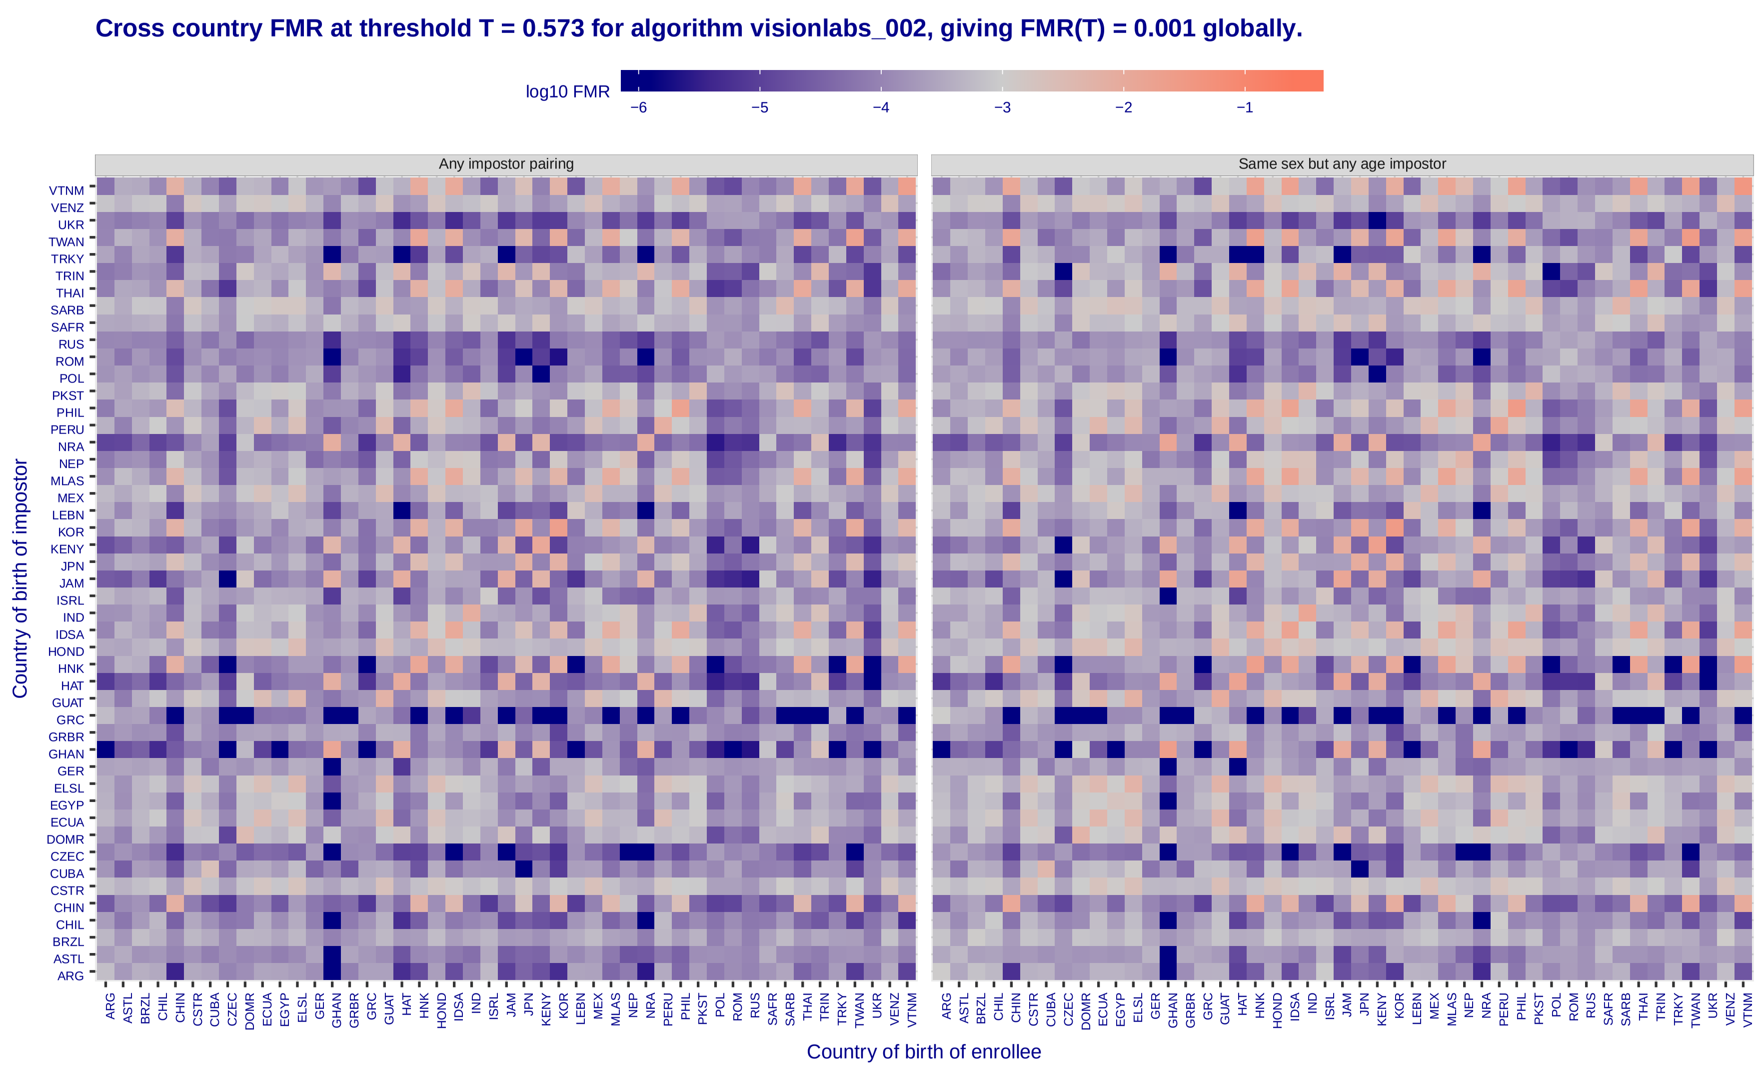

Figure 129: For algorithm visionlabs-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

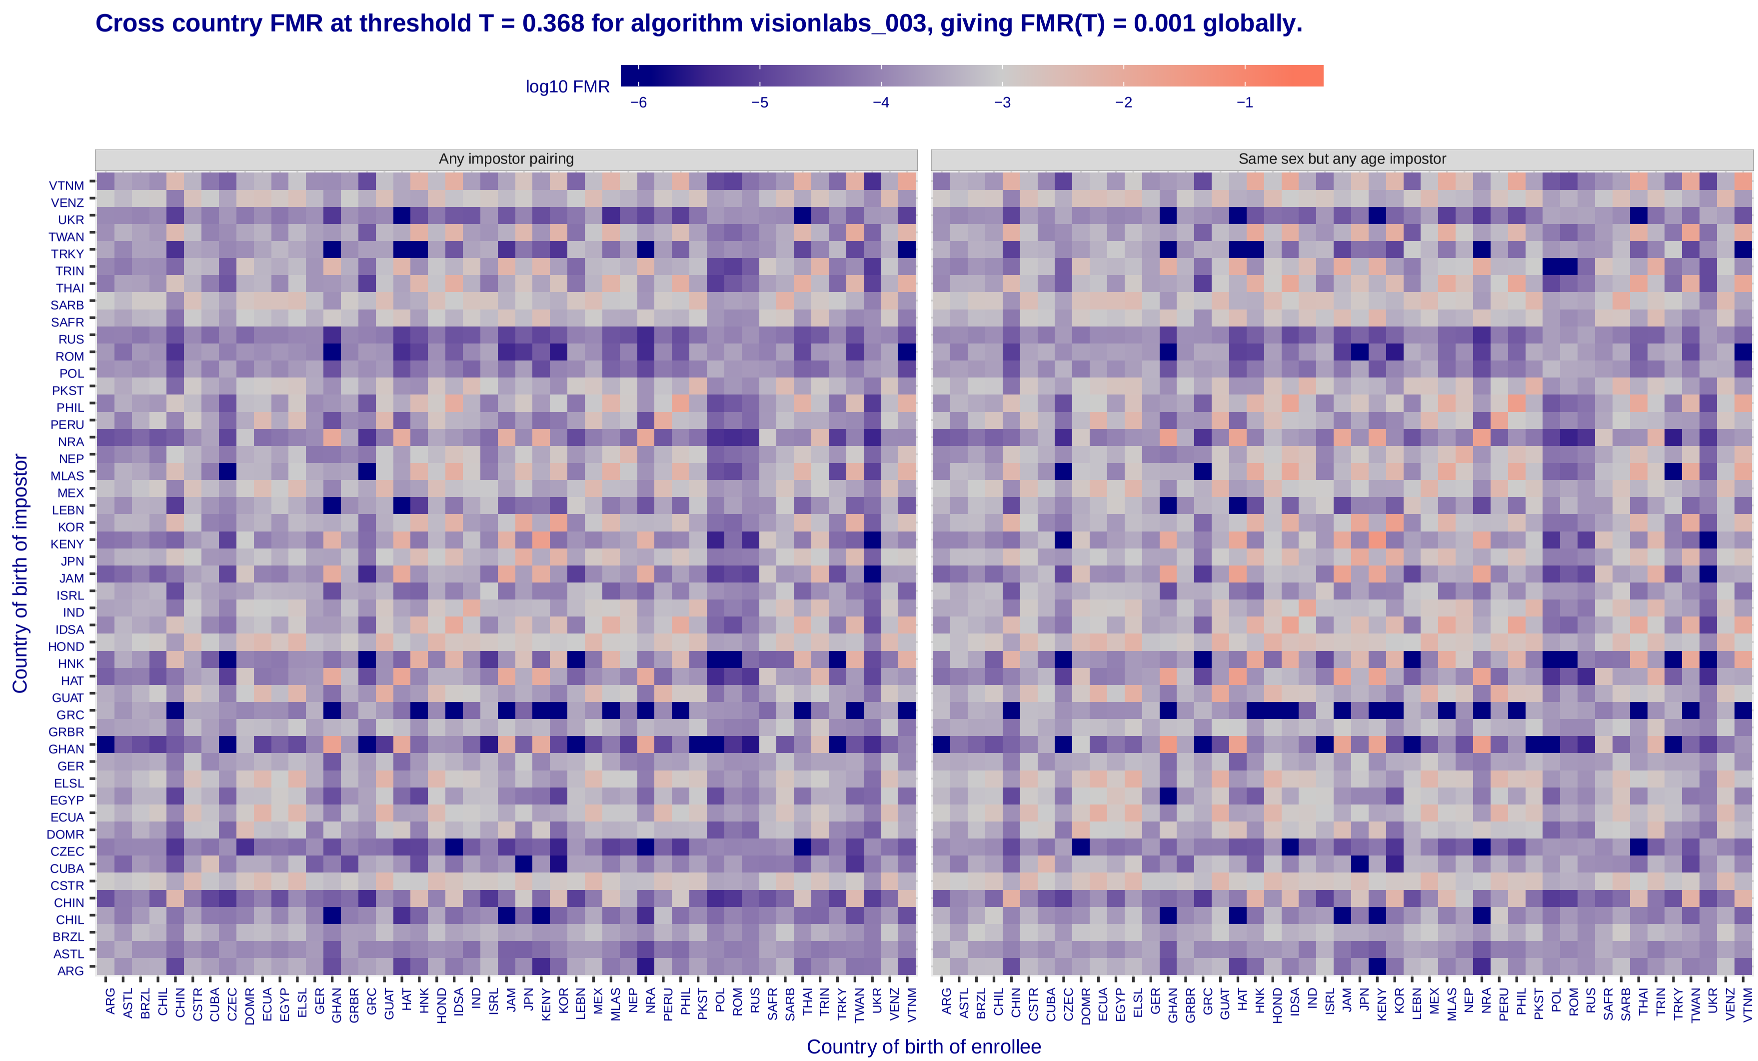

Figure 130: For algorithm visionlabs-003 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

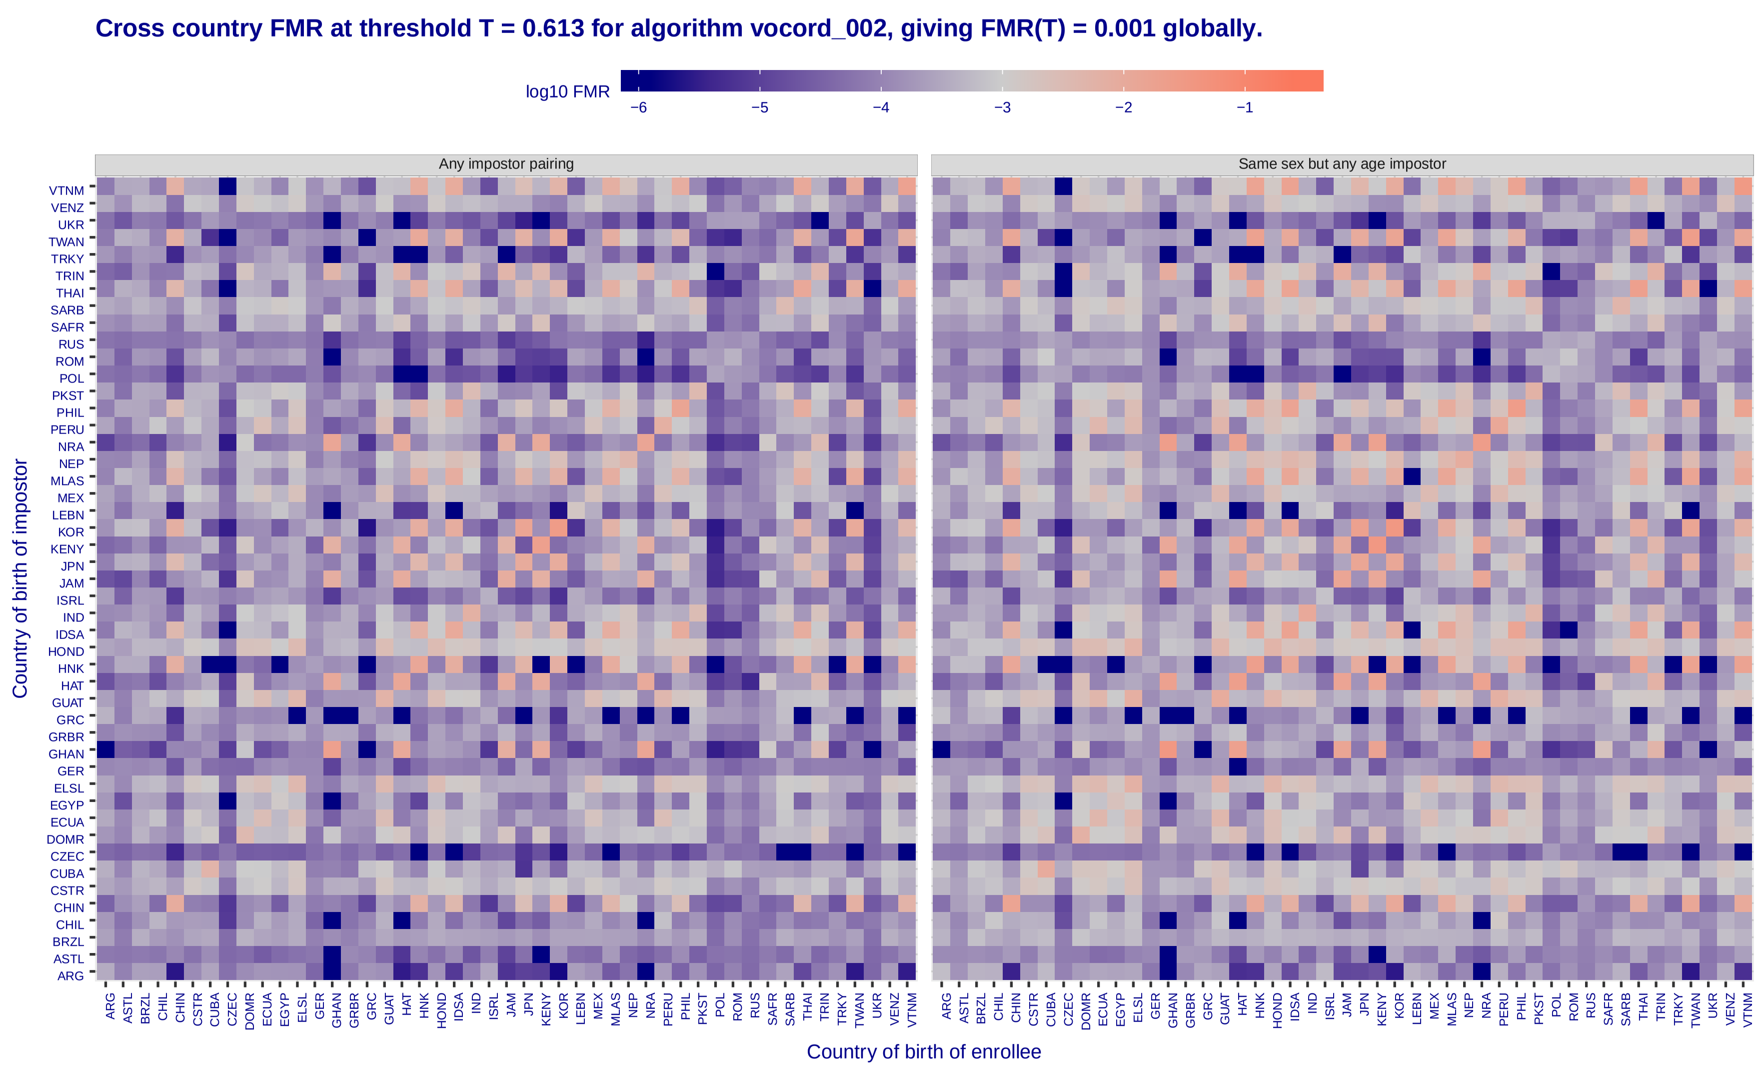

Figure 131: For algorithm vocord-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

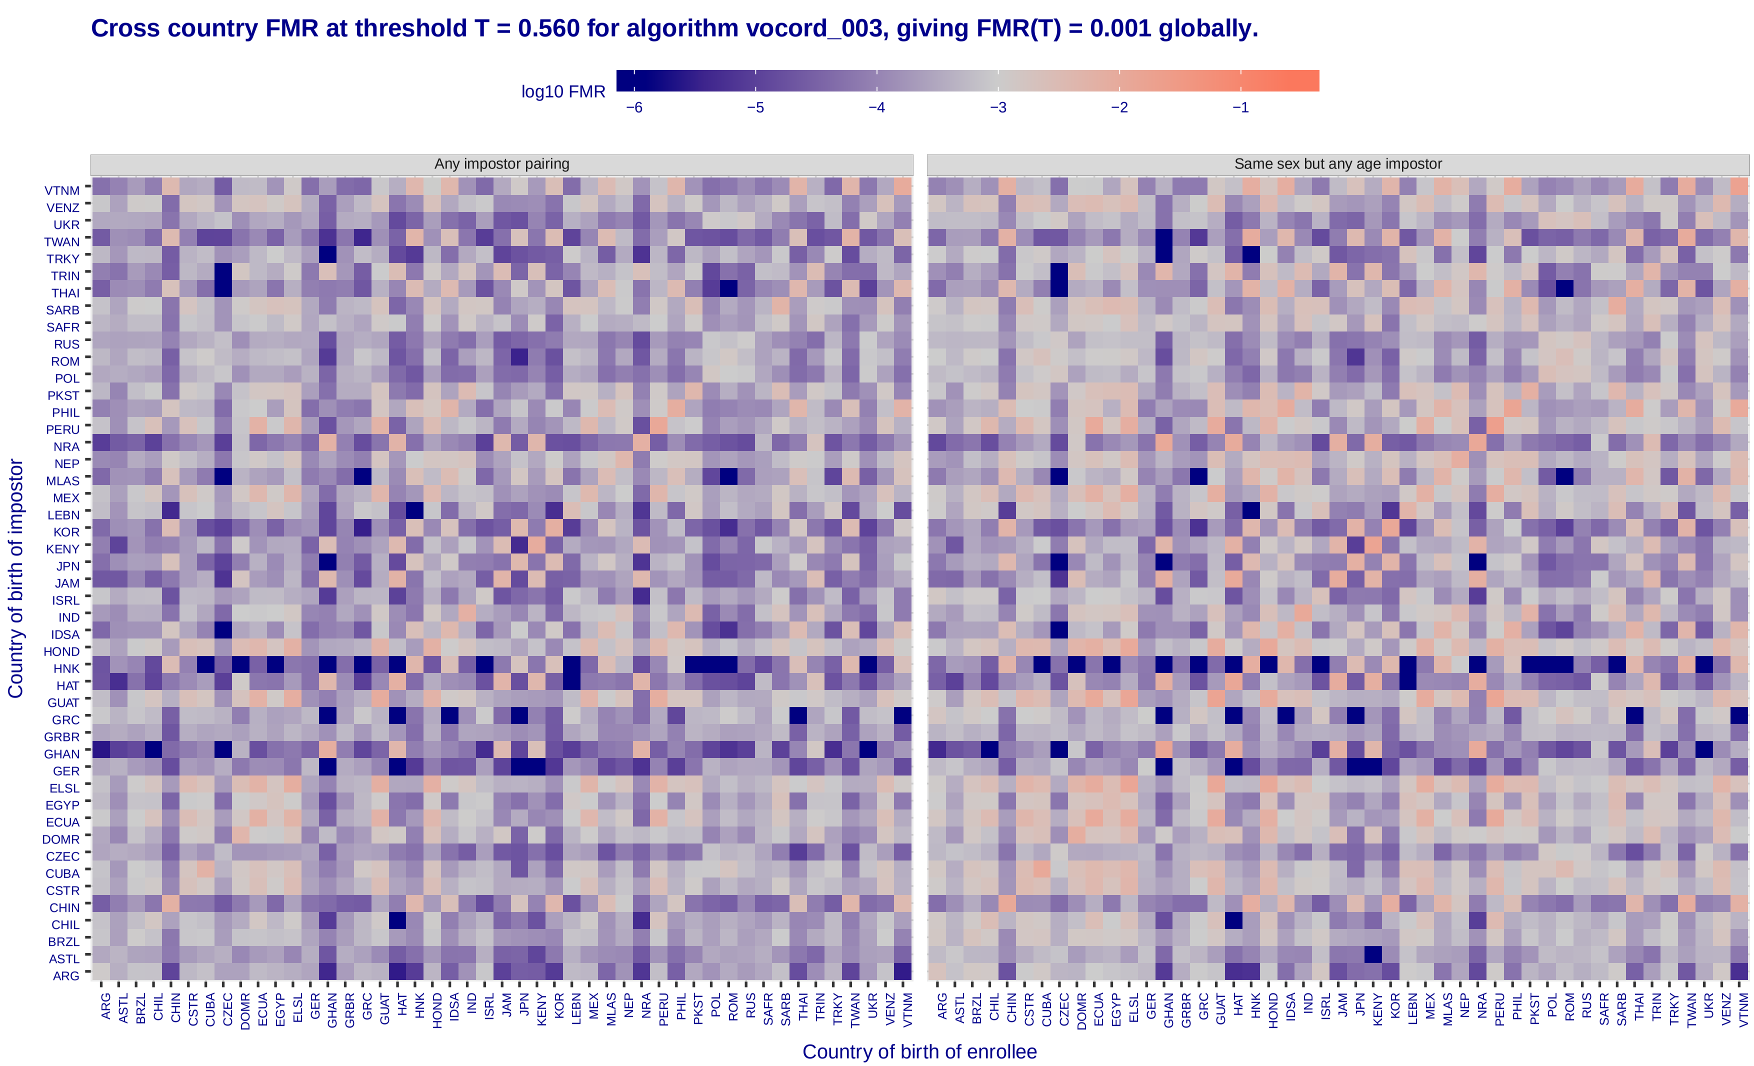

Figure 132: For algorithm vocord-003 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

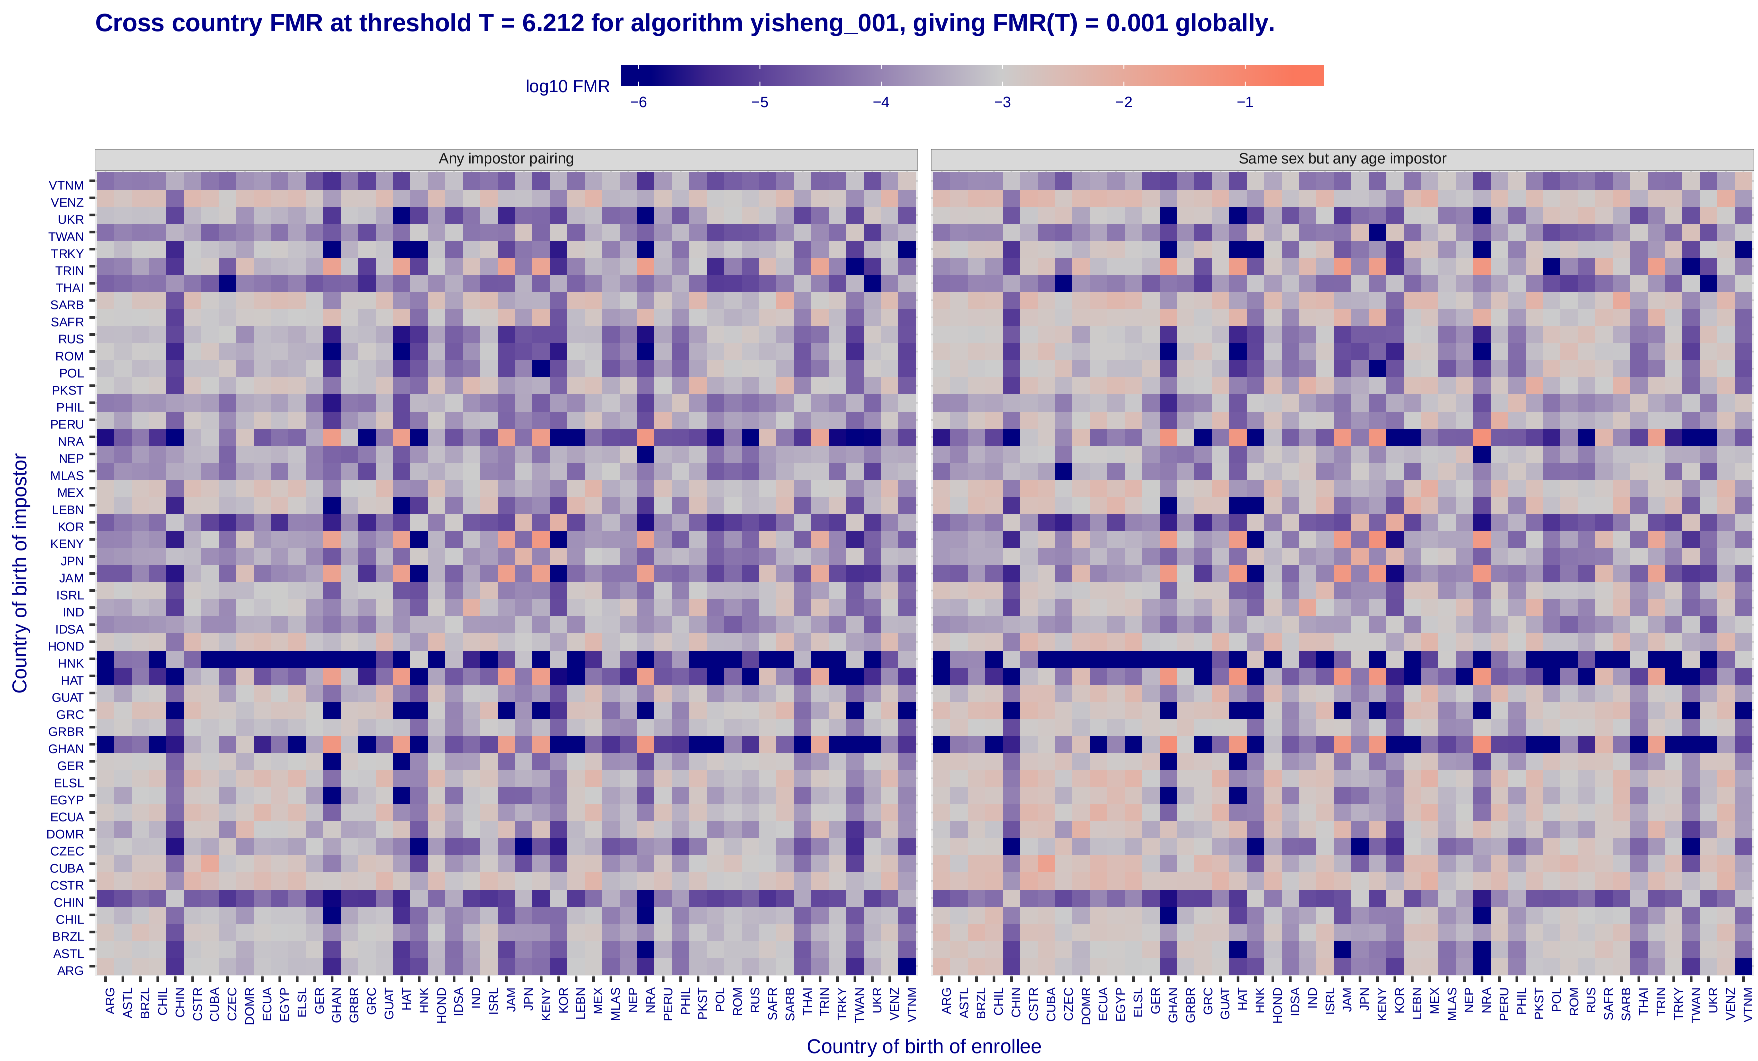

Figure 133: For algorithm yisheng-001 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

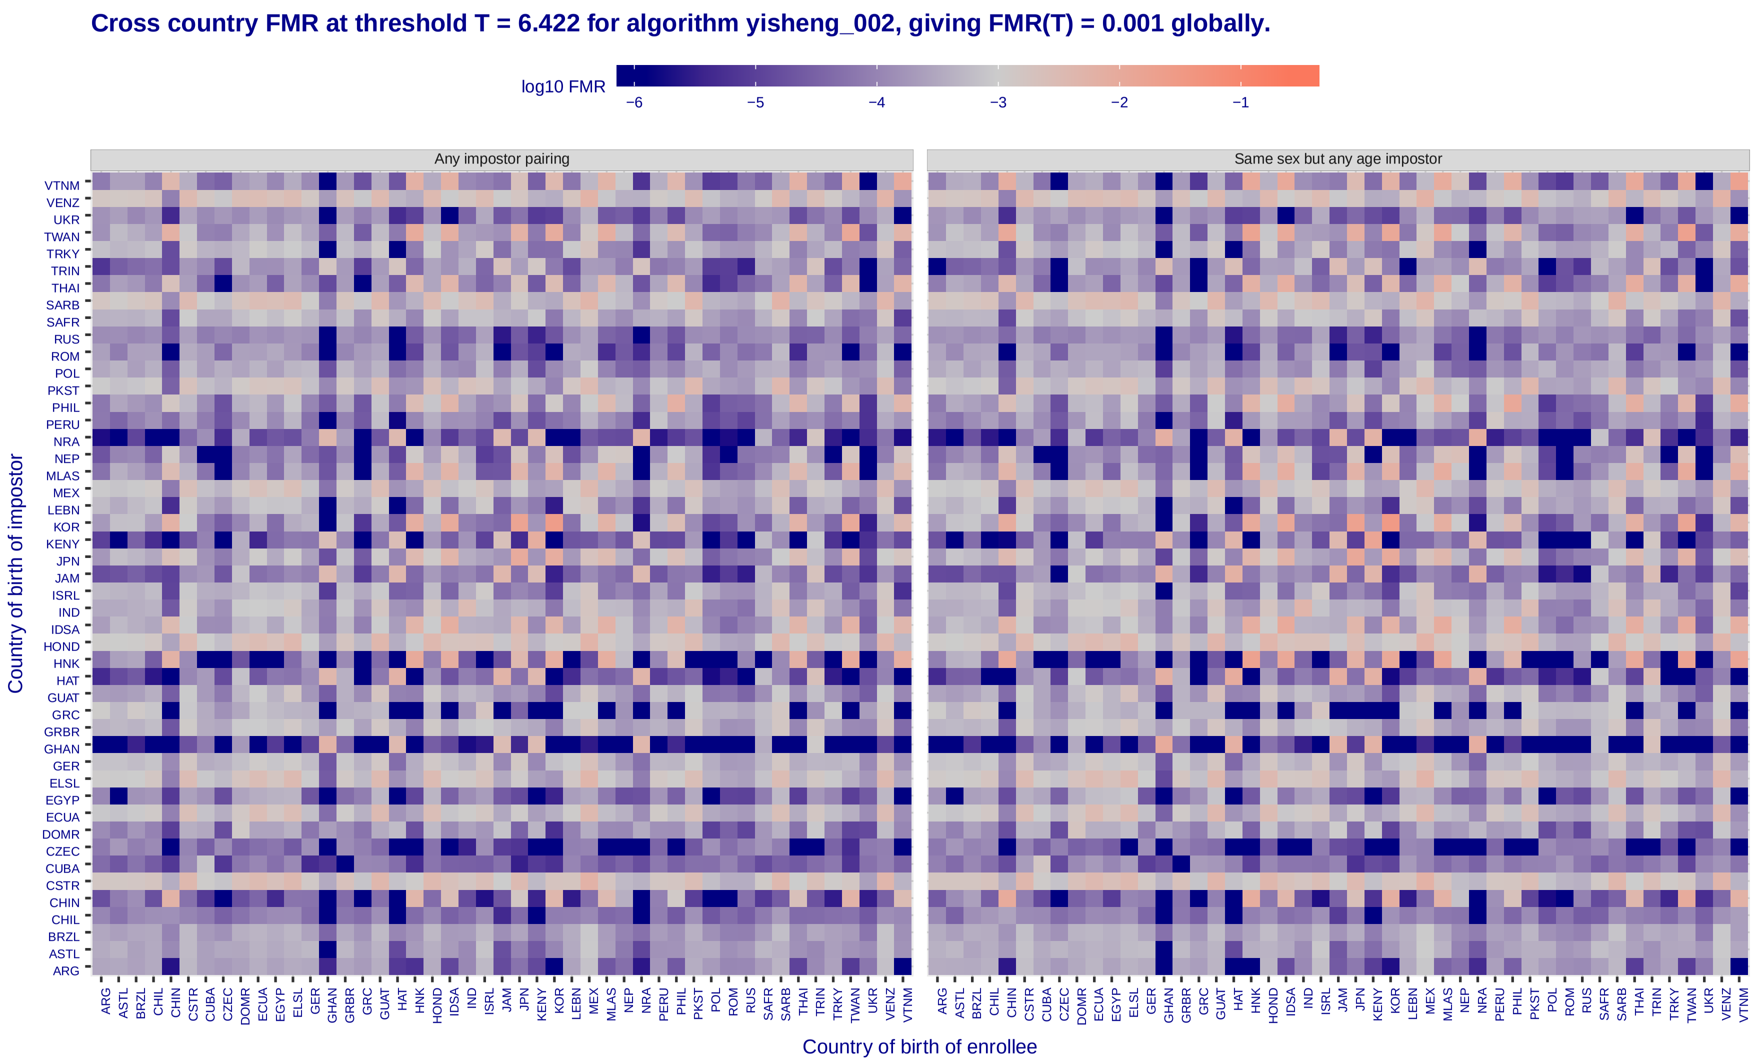

Figure 134: For algorithm yisheng-002 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all $O(10^{10})$ impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

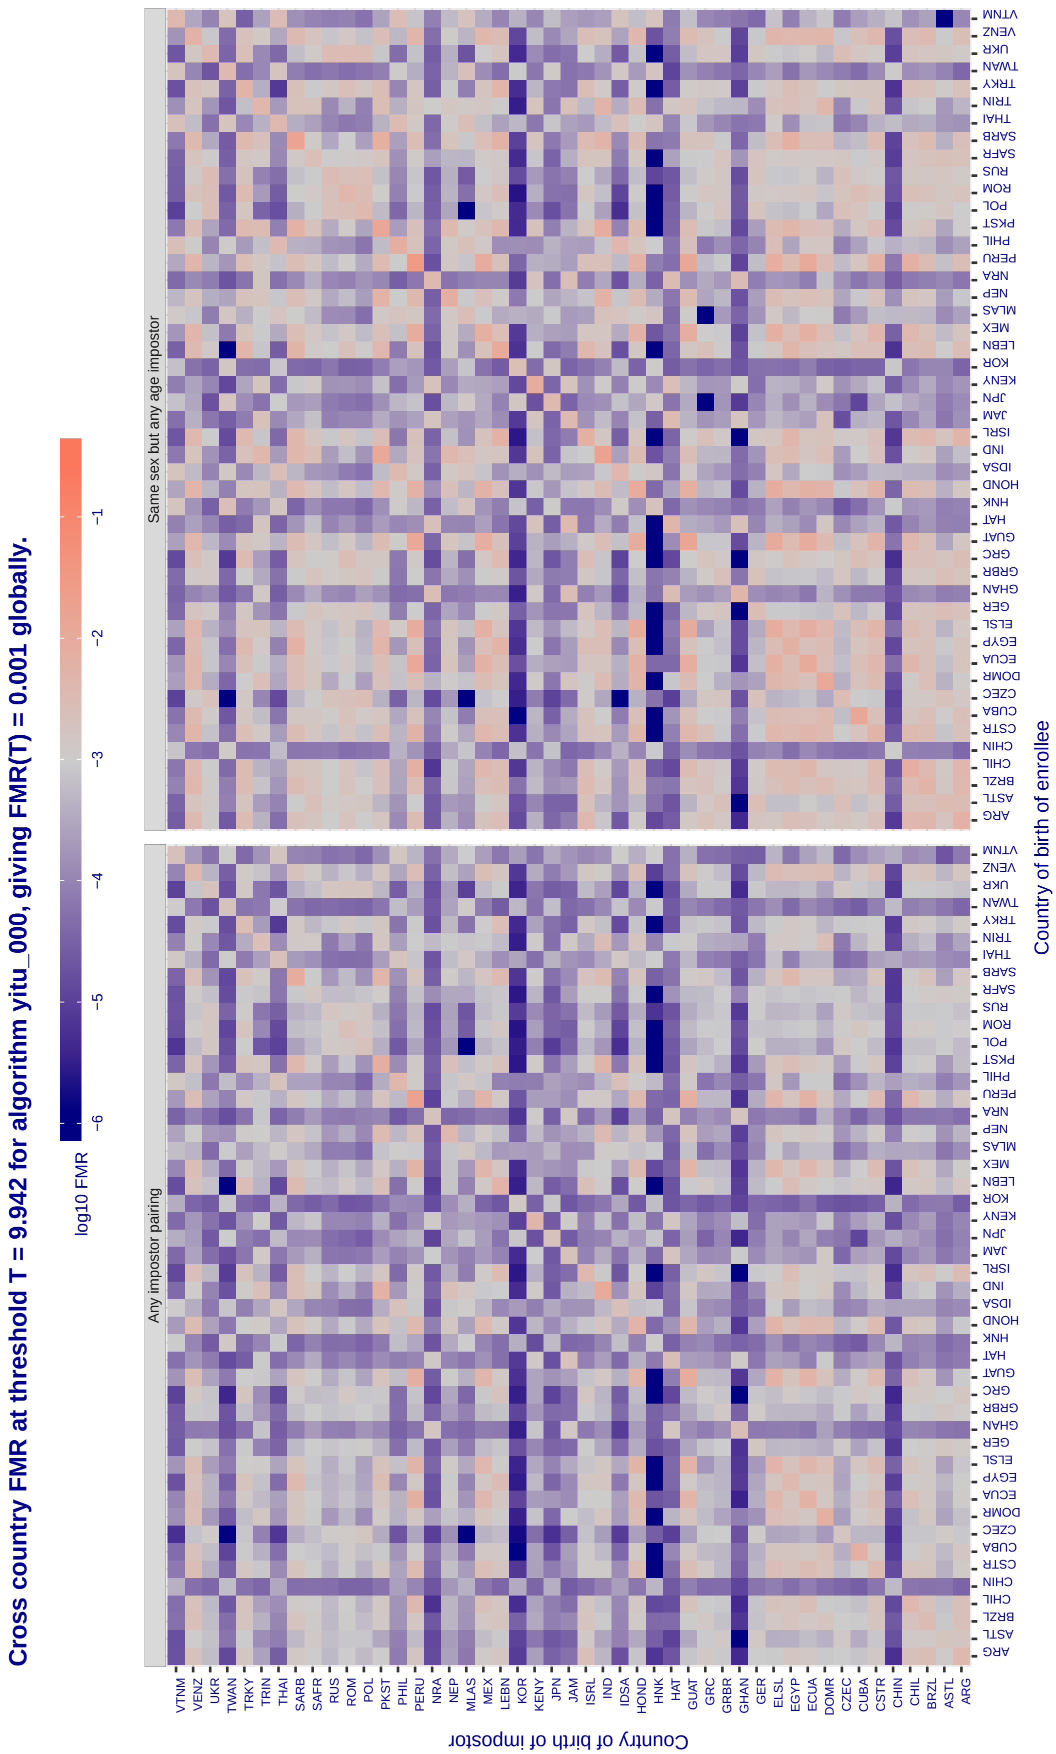

Figure 135: For algorithm yitu-000 operating on visa images, the heatmap shows false match rates observed over impostor comparisons of faces from different individuals who were born in the given country pair. False matches are counted against a recognition threshold fixed globally to give the target FMR in the plot title, computed over all O(10^{10}) impostor comparisons. If text appears in each box it give the same quantity as that coded by the color. Grey indicates FMR is at the intended FMR target level. Light red colors present a security vulnerability to, for example, a passport gate. Each +1 increase in log10 FMR corresponds to a factor of 10 increase in FMR. The matrix is not quite symmetric because images in the enrollment and verification sets are different.

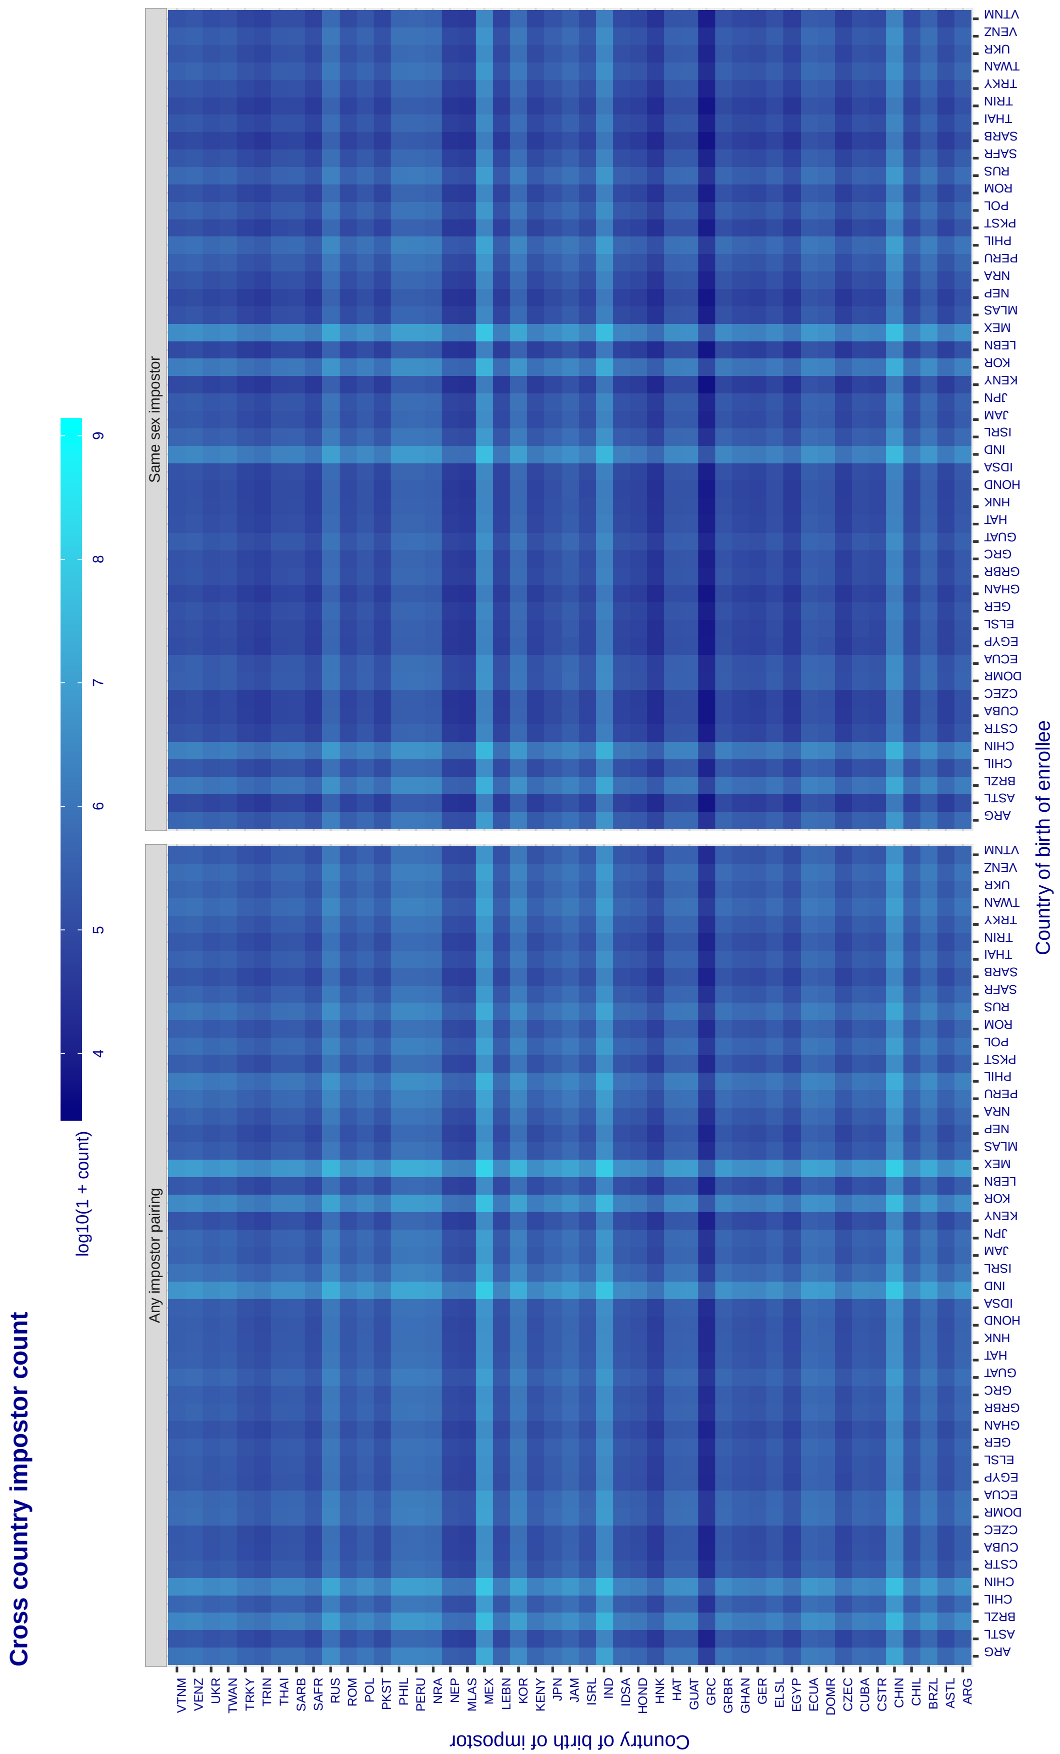

Figure 136: For visa images, the heatmap shows the count of impostor comparisons of faces from different individuals who were born in the given country pair.

5.6.2 Effect of age on impostors

Background: This section shows the effect of age on the impostor distribution. The ideal behaviour is that the age of the enrollee and the impostor would not affect impostor scores. This would support FMR stability over sub-populations.

Goals:

- ▷ To show the effect of relative ages of the impostor and enrollee on false match rates.
- ▷ To determine whether some algorithms have better impostor distribution stability.

Methods:

- ▷ Define 14 age group bins, spanning 0 to over 100 years old.
- ▷ Compute FMR over all impostor comparisons for which the subjects in the enrollee and impostor images have ages in two bins.
- ▷ Compute FMR over all impostor comparisons for which the subjects are additionally of the same sex, and born in the same geographic region.

Results:

The notable aspects are:

- ▷ Diagonal dominance: Impostors are more likely to be matched against their same age group.
- ▷ Same sex and same region impostors are more successful. On the diagonal, an impostor is more likely to succeed by posing as someone of the same sex. If $\Delta \log_{10} \text{FMR} = 0.2$, then same-sex same-region FMR exceeds the all-pairs FMR by factor of $10^{0.2} = 1.6$.
- ▷ Young children impostors give elevated FMR against young children. Older adult impostor give elevated FMR against older adults. These effects are quite large, for example if $\Delta \log_{10} \text{FMR} = 1.0$ larger than a 32 year old, then these groups have higher FMR by a factor of $10^1 = 10$. This would imply an FMR above 0.01 for a nominal (global) FMR = 0.001.
- ▷ Algorithms vary.
- ▷ We computed the same quantities for a global FMR = 0.0001. The effects are similar.

Note the calculations in this section include impostors paired across all countries of birth.

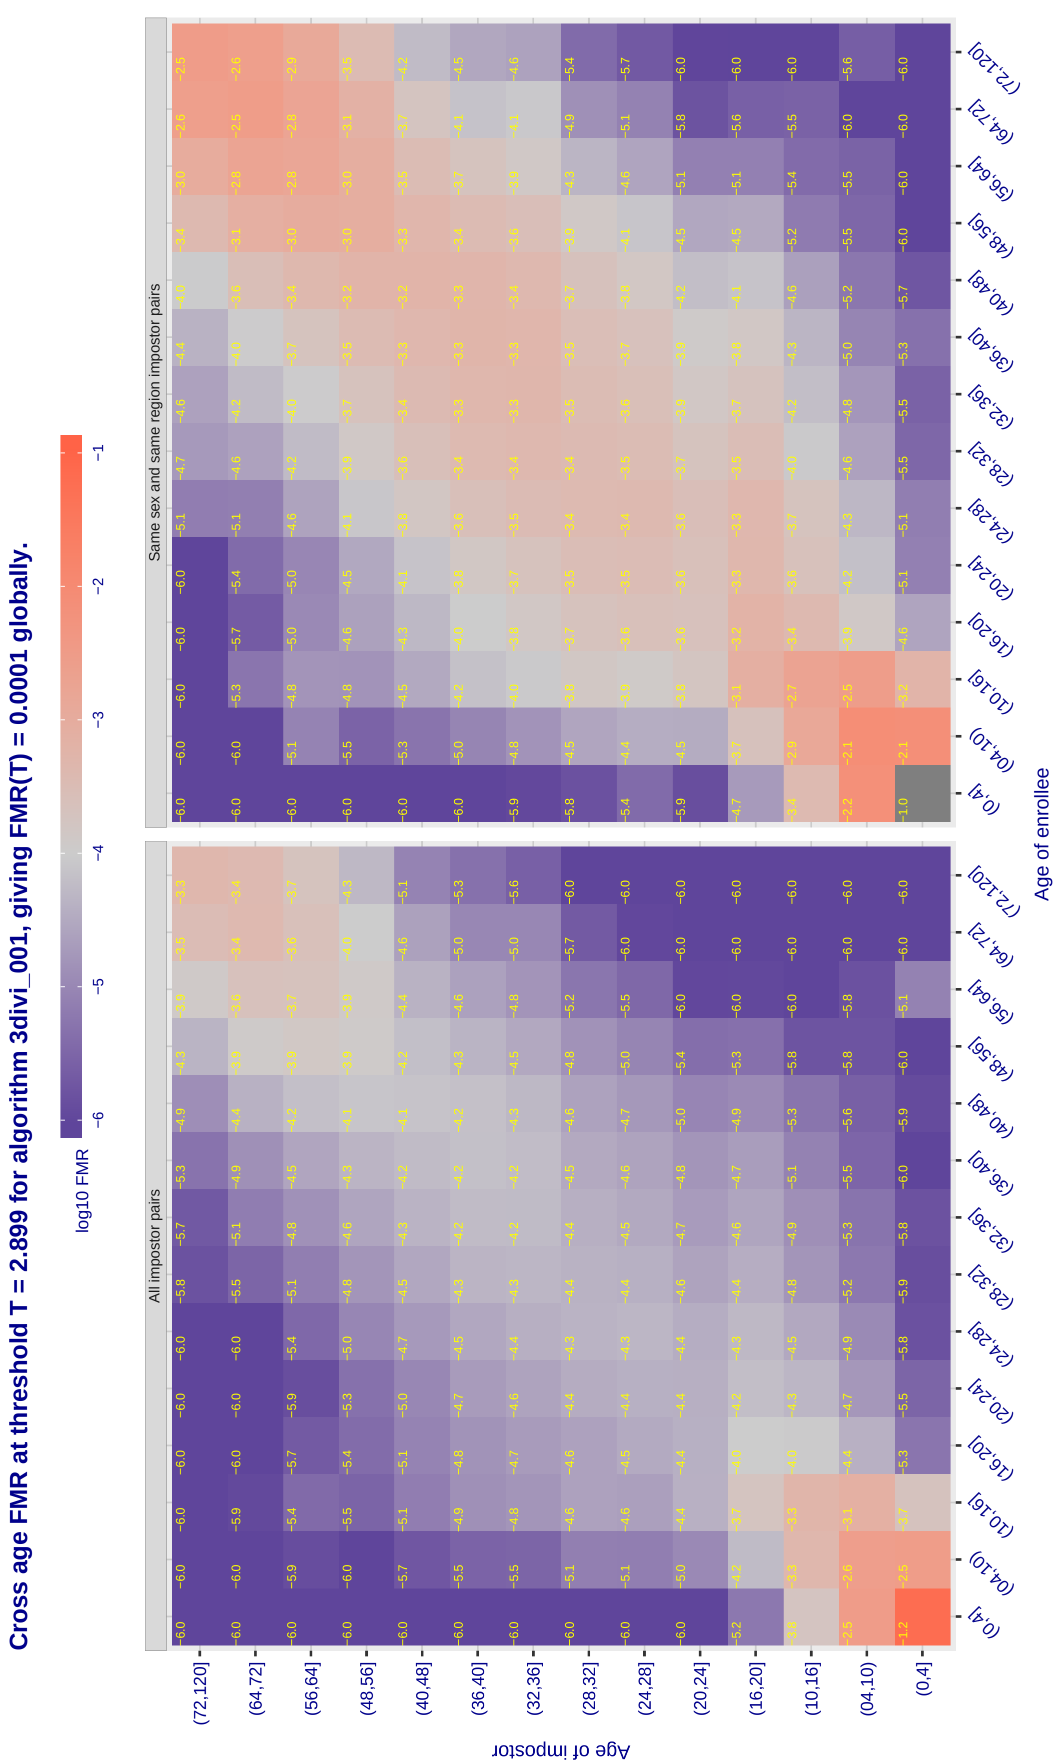

Figure 137: For algorithm 3divi-001 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

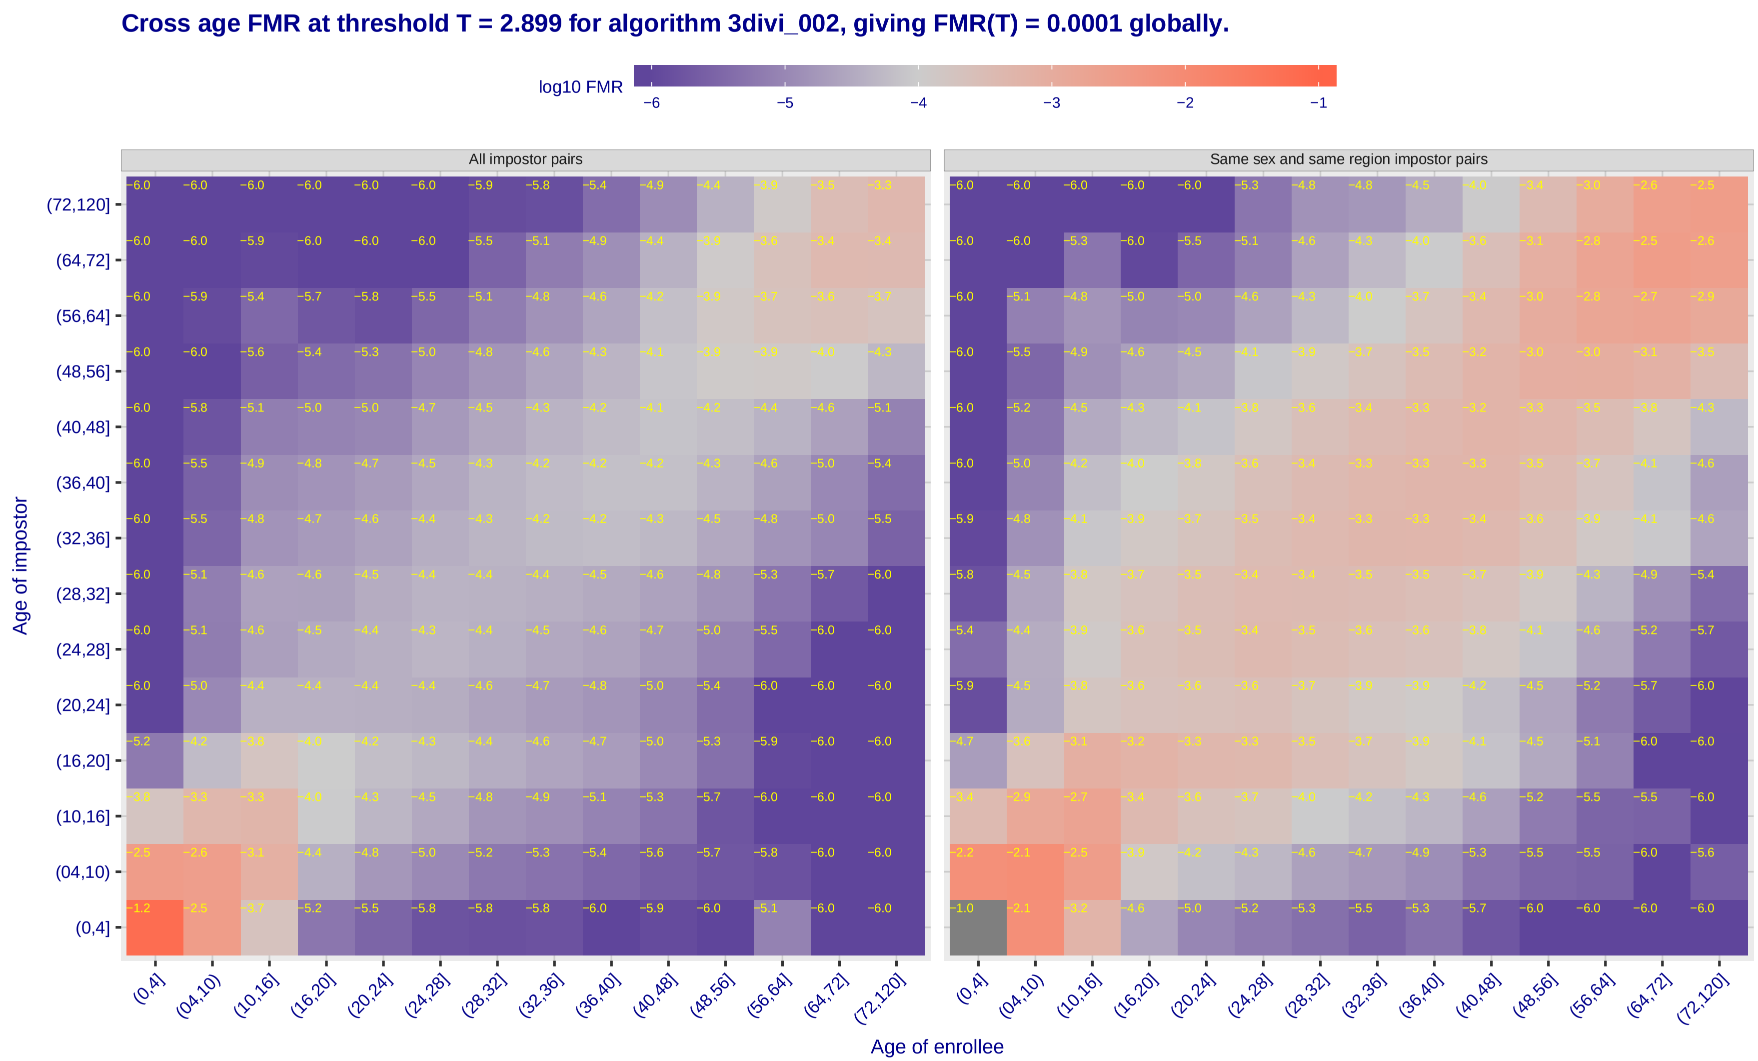

Figure 138: For algorithm 3divi-002 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

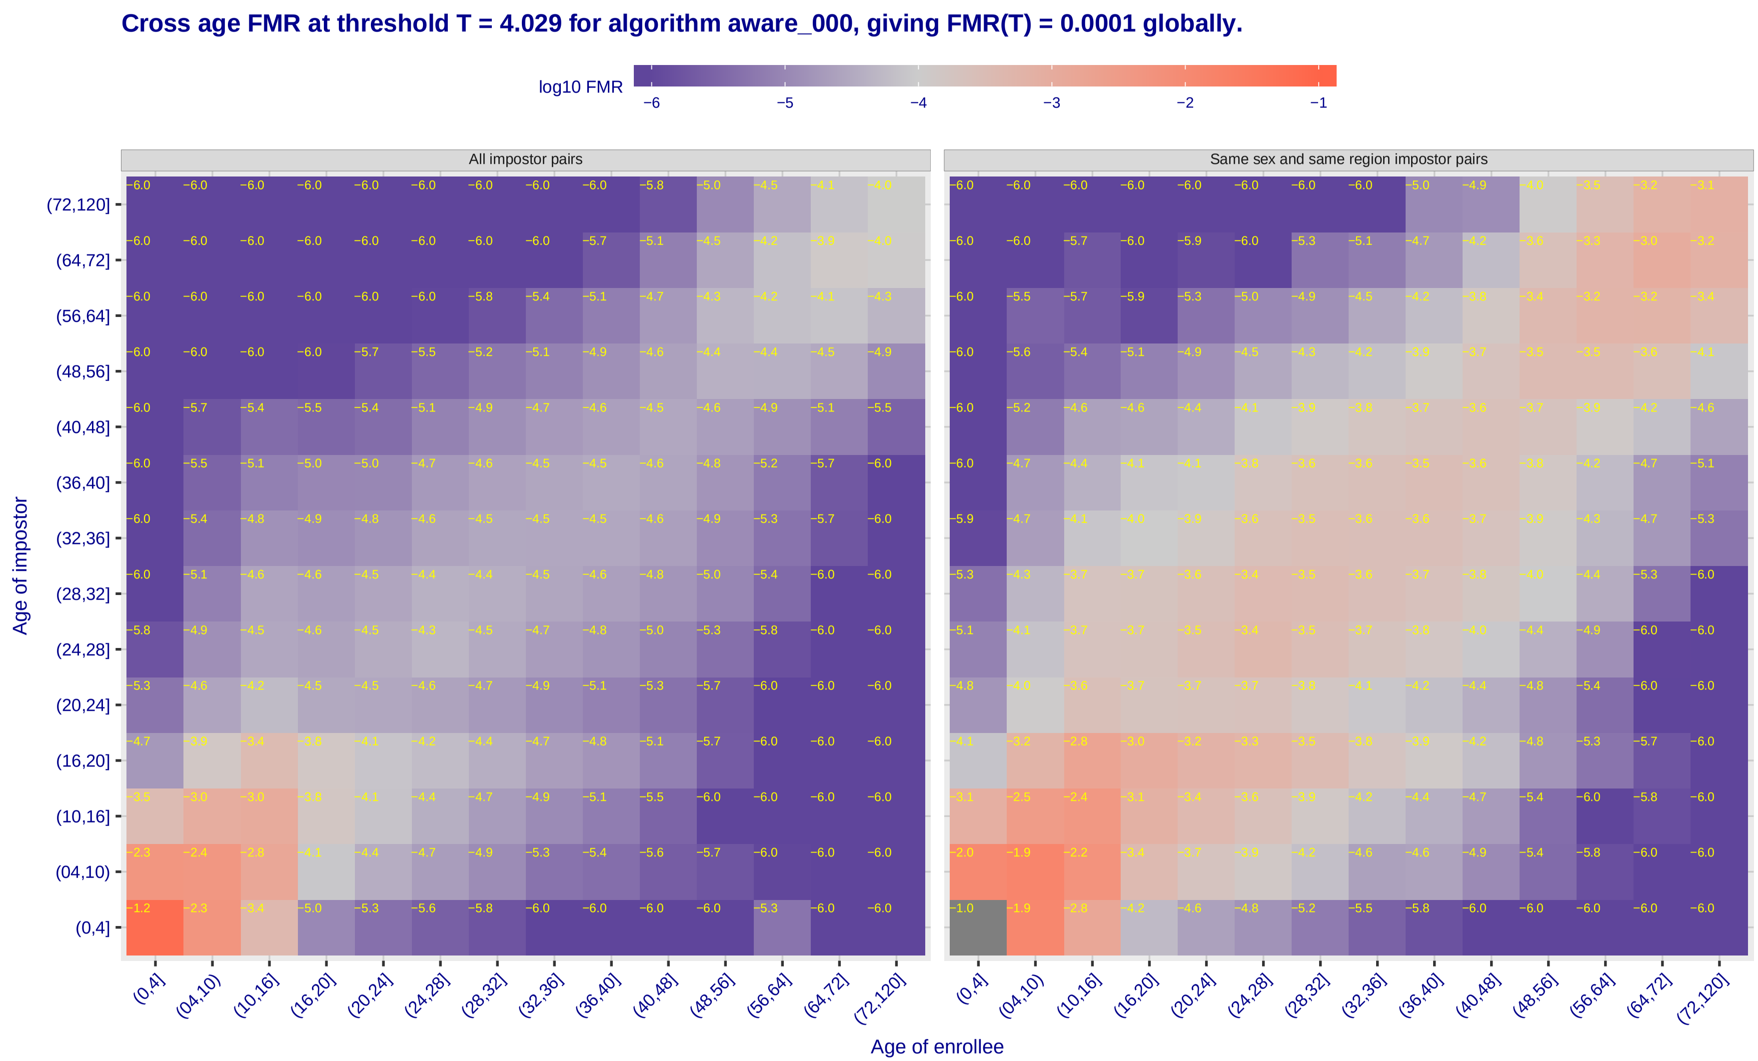

Figure 139: For algorithm aware-000 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

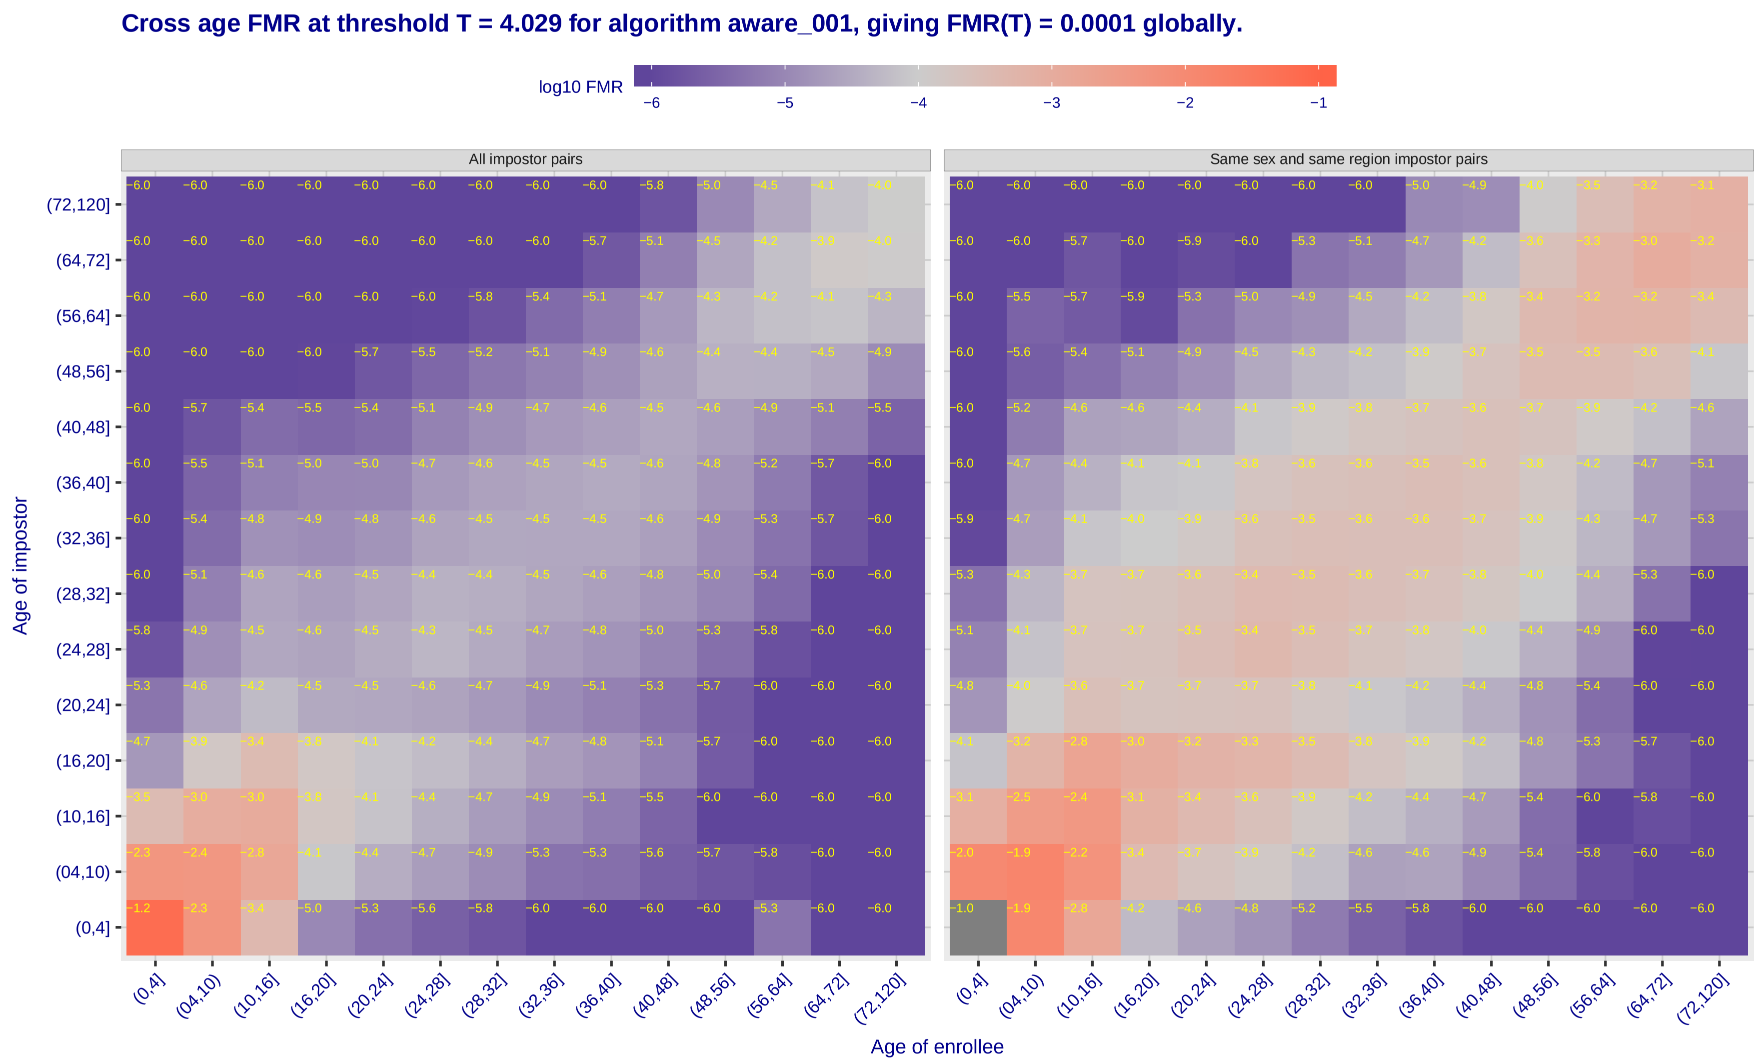

Figure 140: *For algorithm aware-001 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all $O(10^{10})$ impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.*

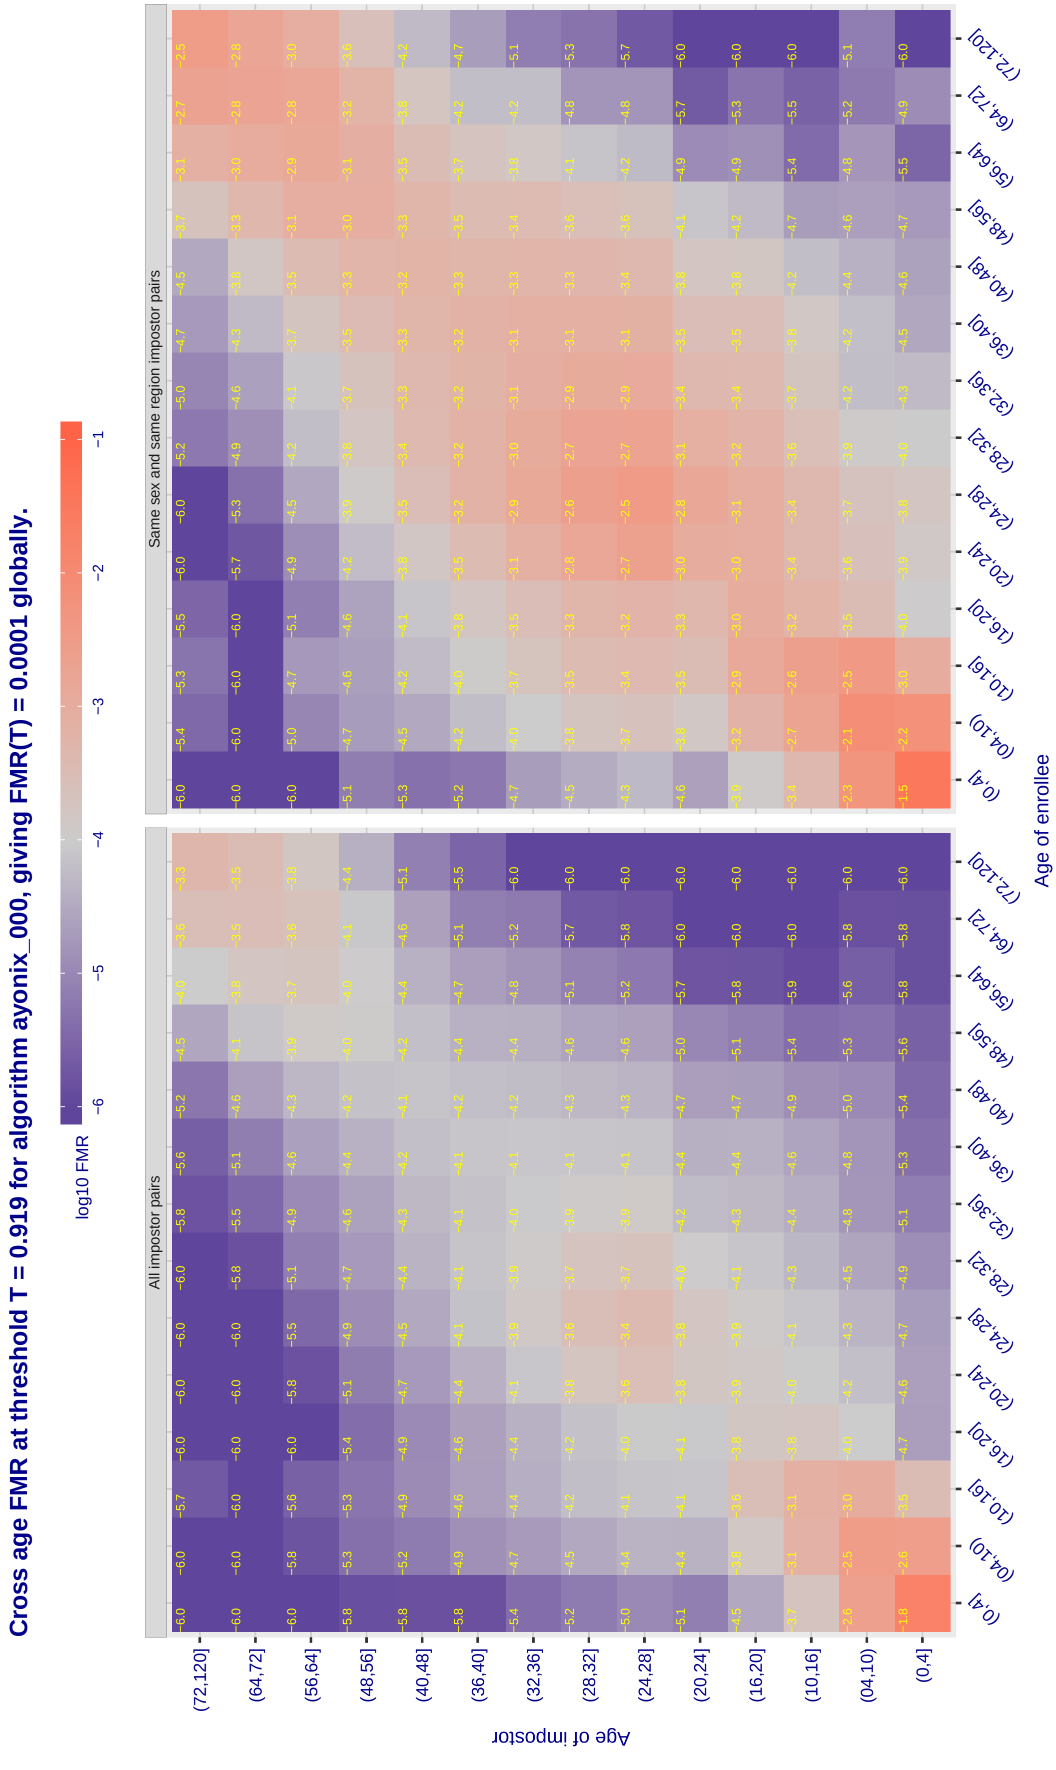

Figure 141: For algorithm ayonix-000 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all $O(10^{10})$ impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

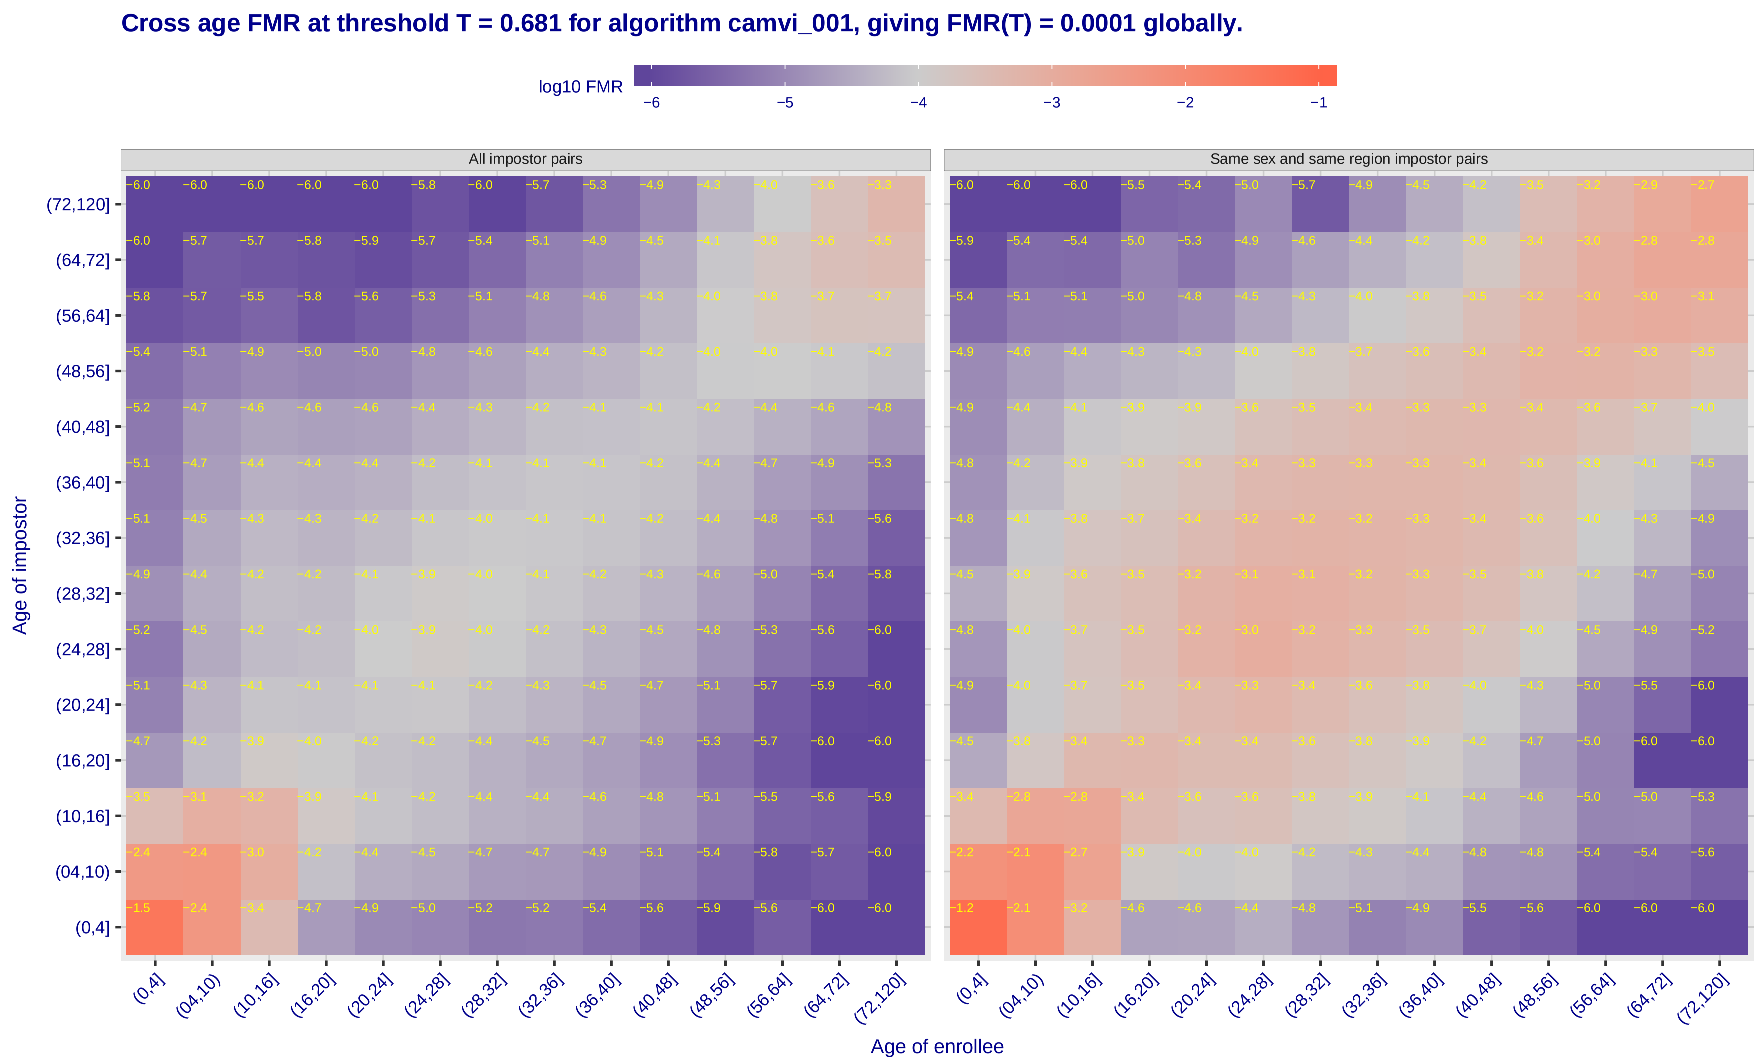

Figure 142: For algorithm camvi-001 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all $O(10^{10})$ impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

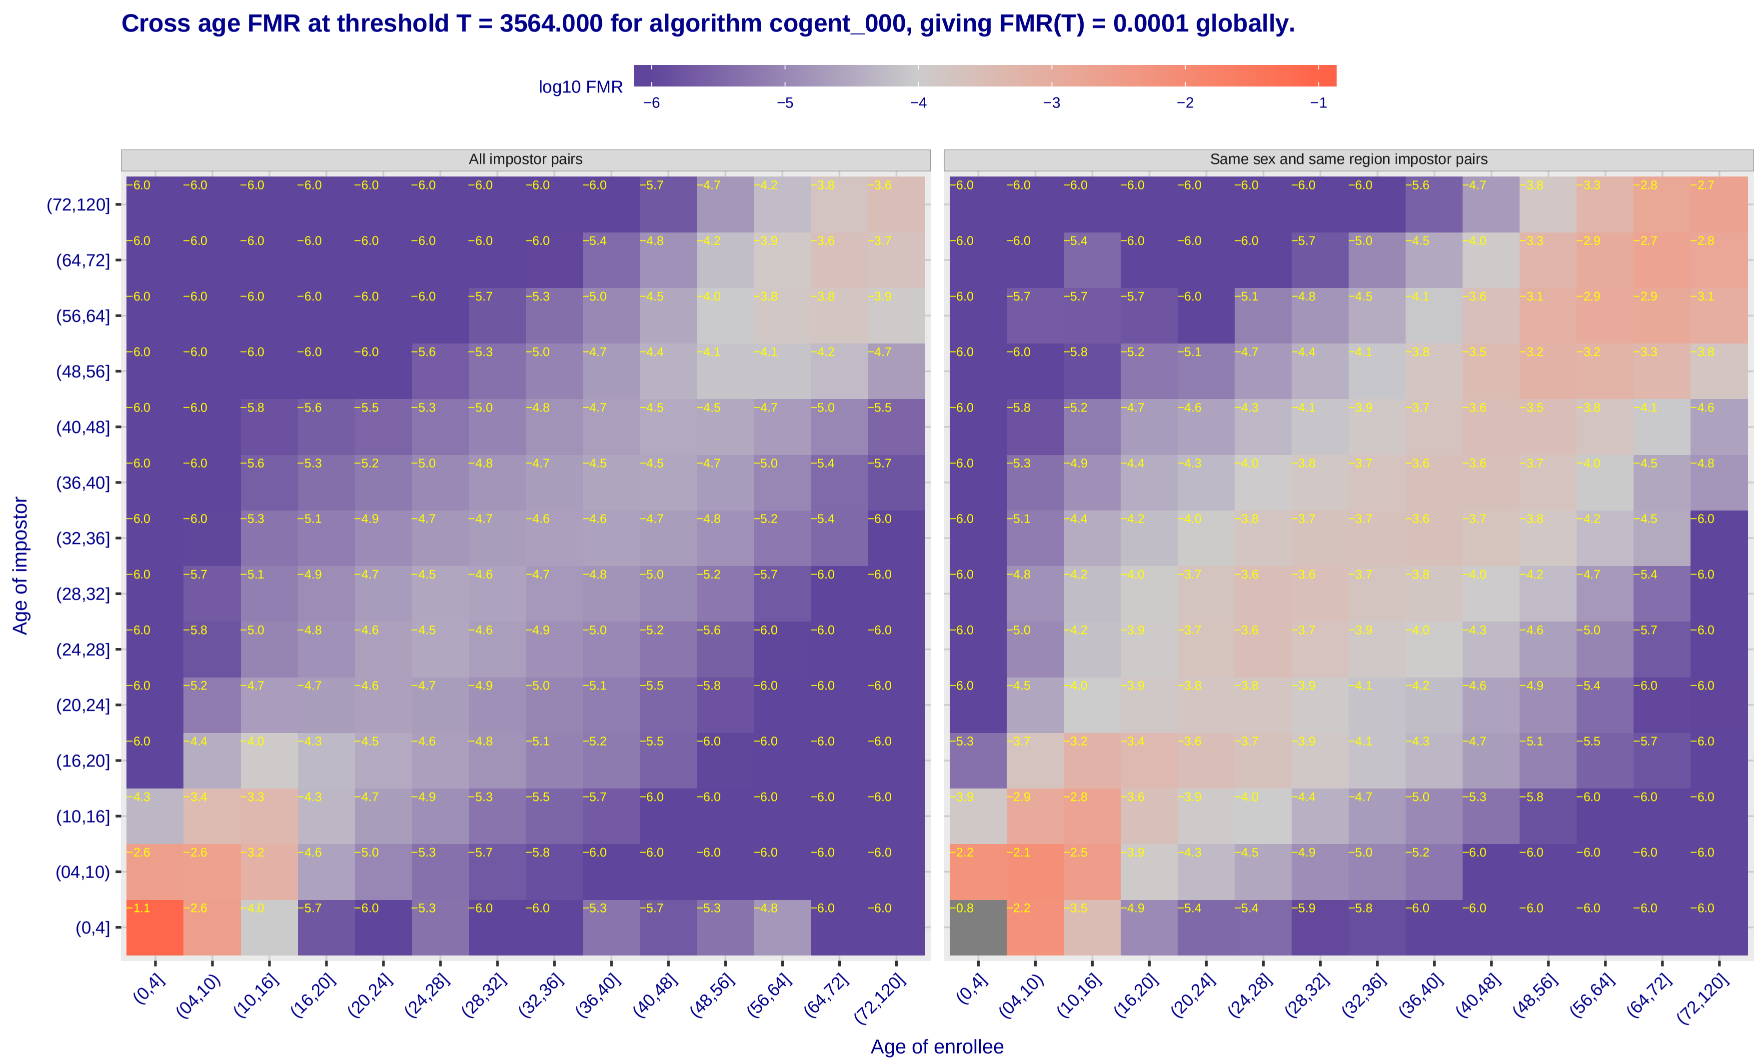

Figure 143: For algorithm cogent-000 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

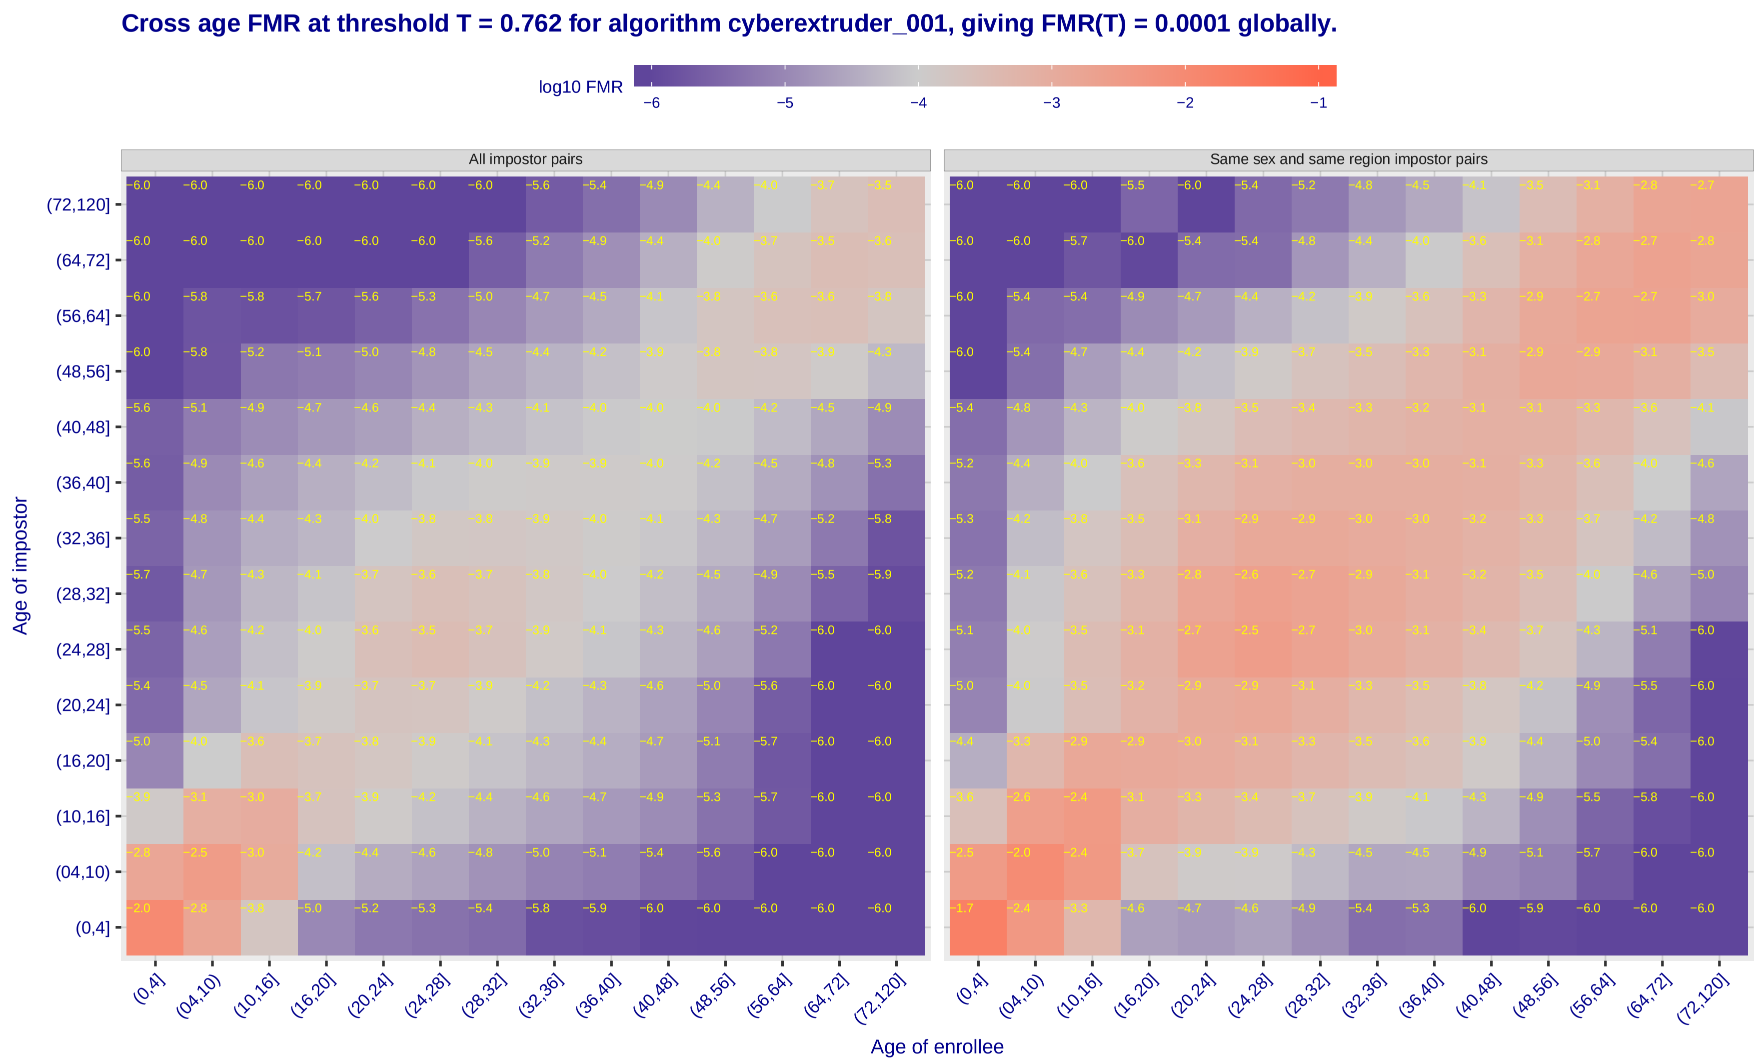

Figure 144: For algorithm cyberextruder-001 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

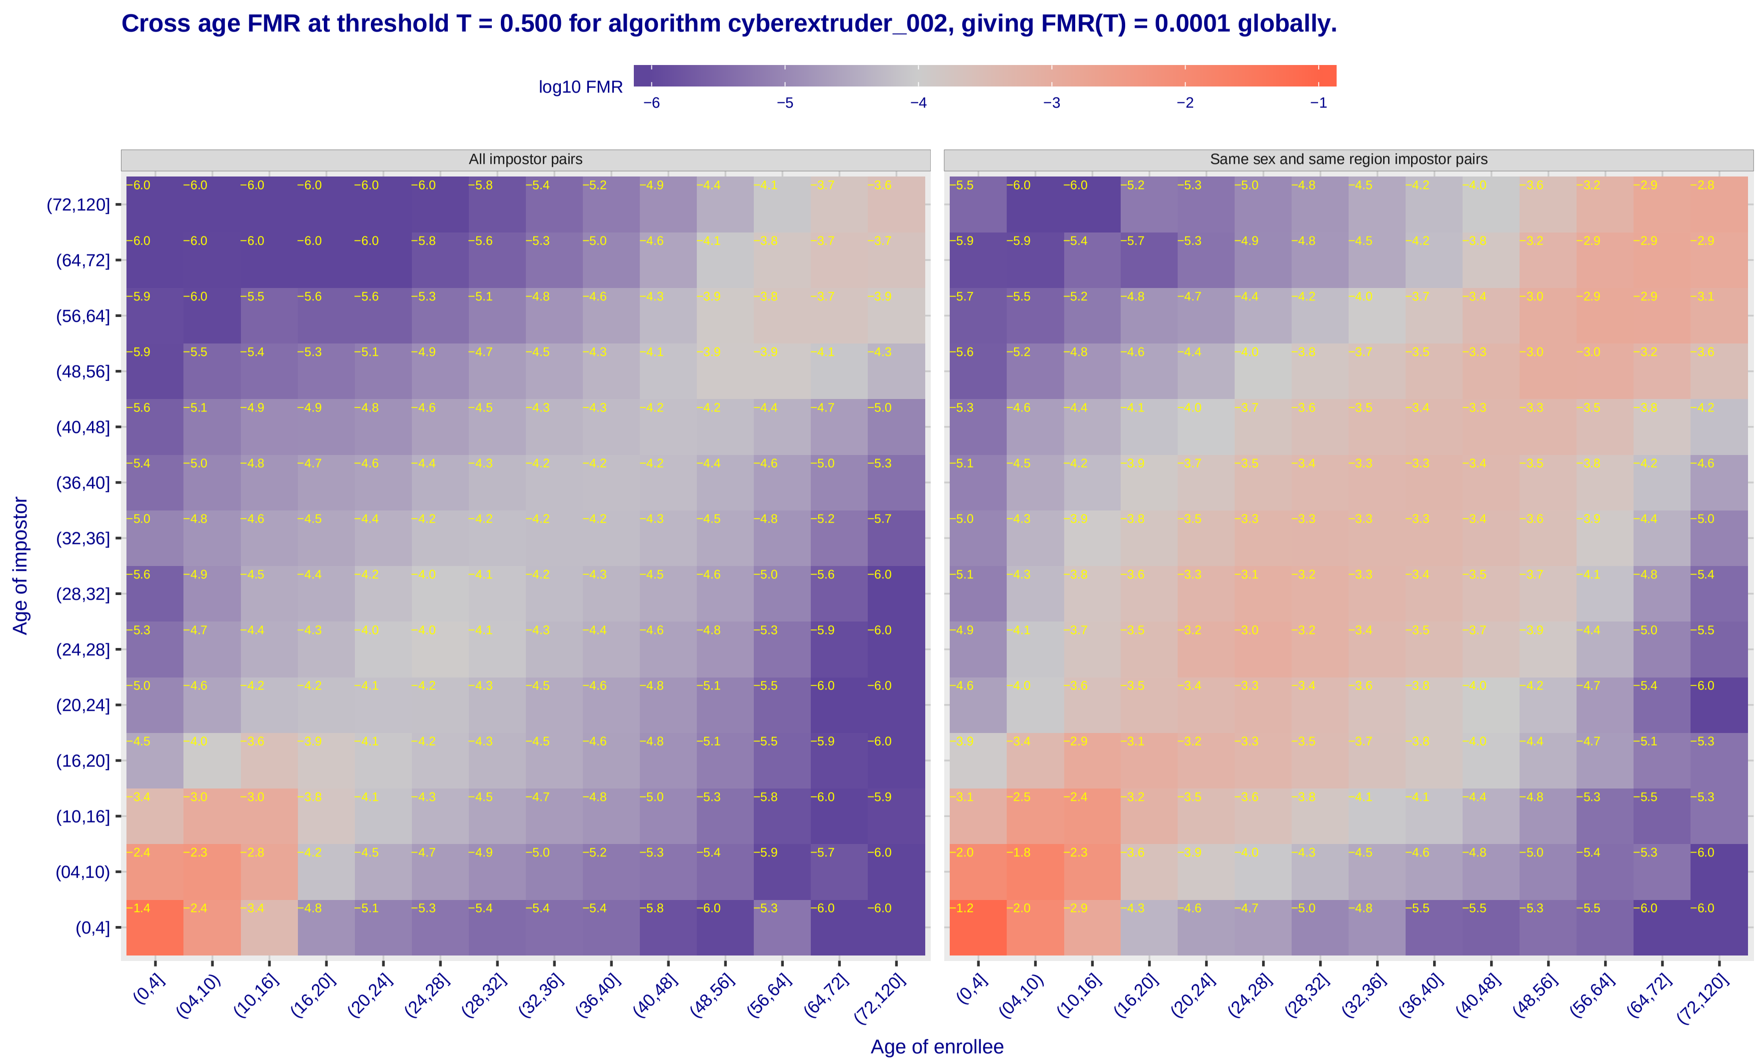

Figure 145: For algorithm cyberextruder-002 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

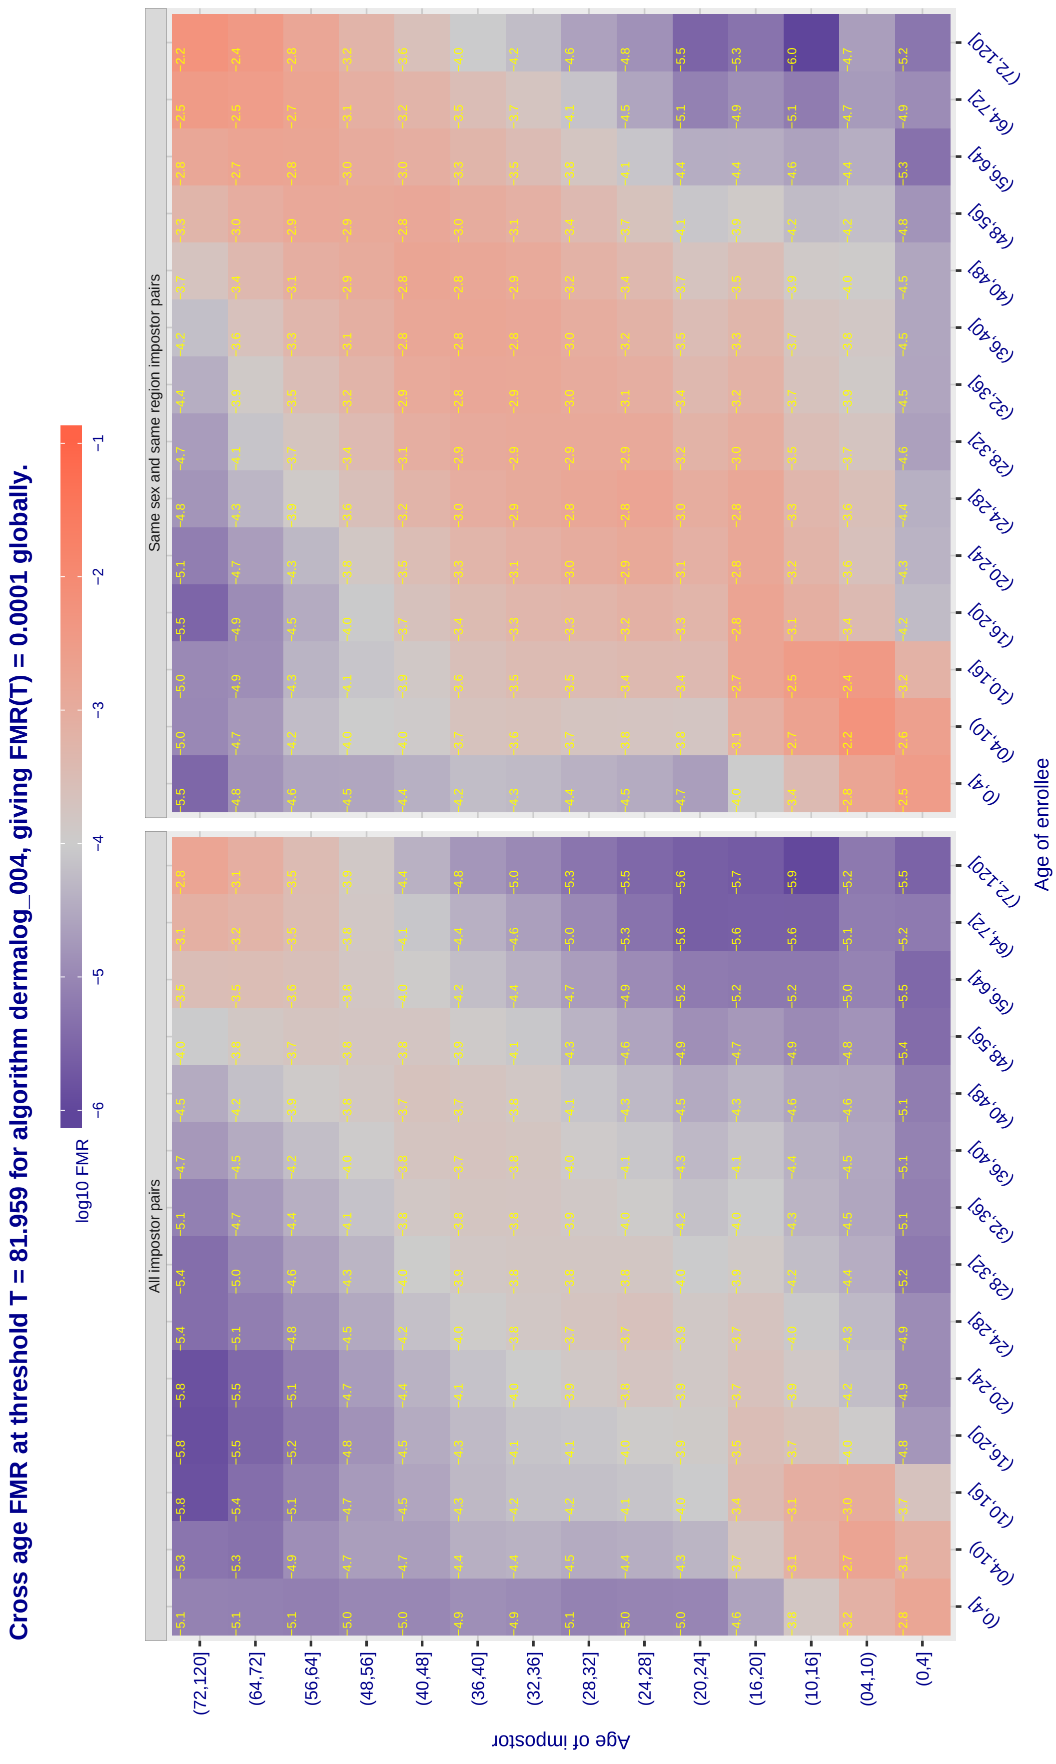

Figure 146: For algorithm dermalog-004 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

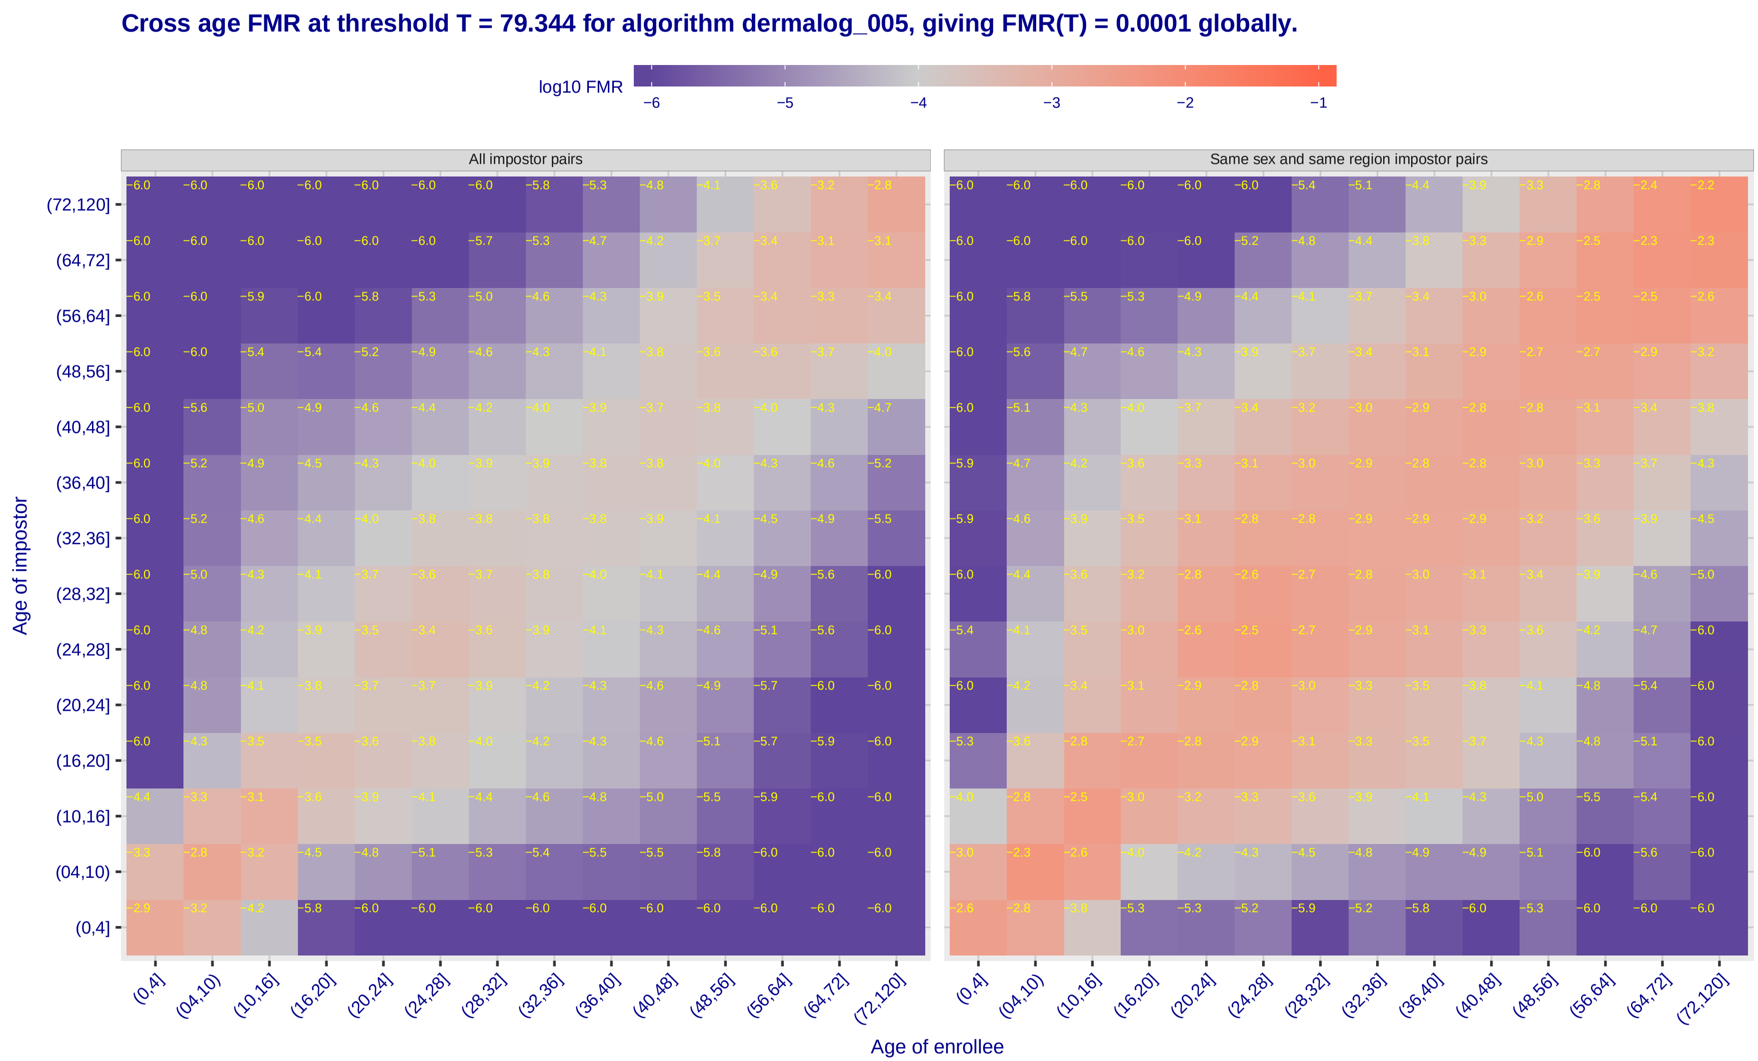

Figure 147: For algorithm dermalog-005 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

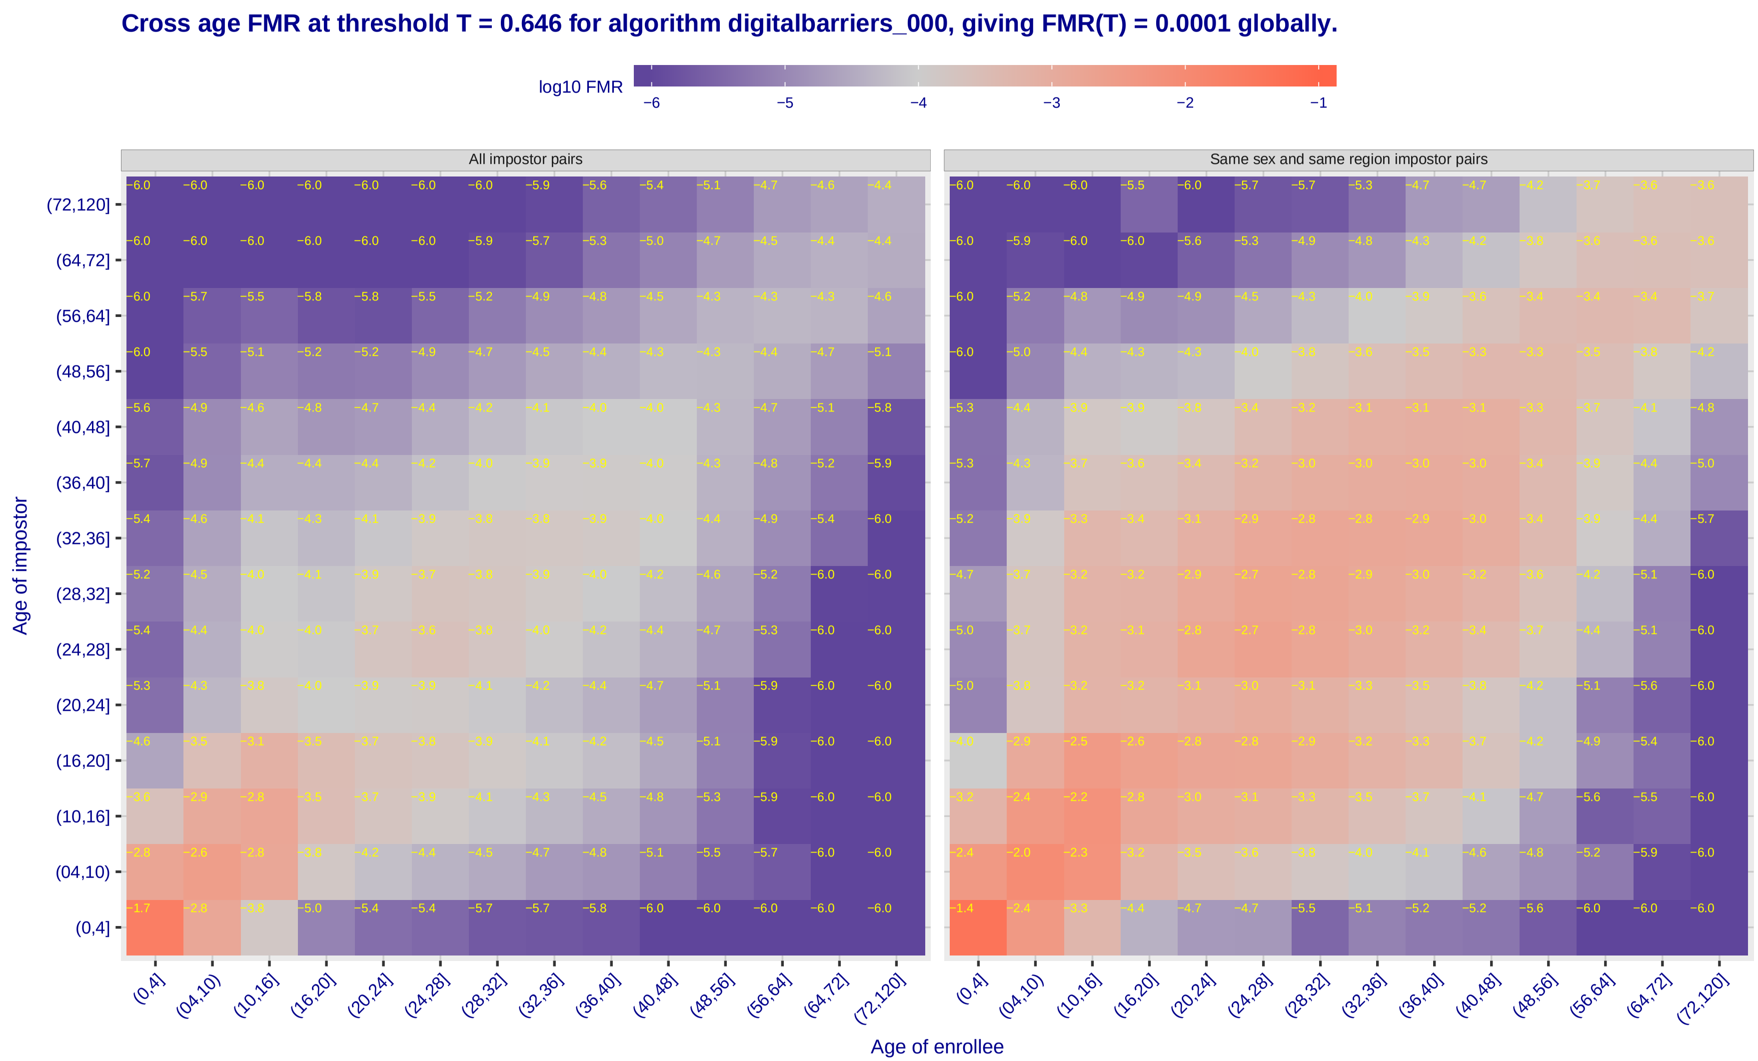

Figure 148: For algorithm digitalbarriers-000 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all $O(10^{10})$ impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

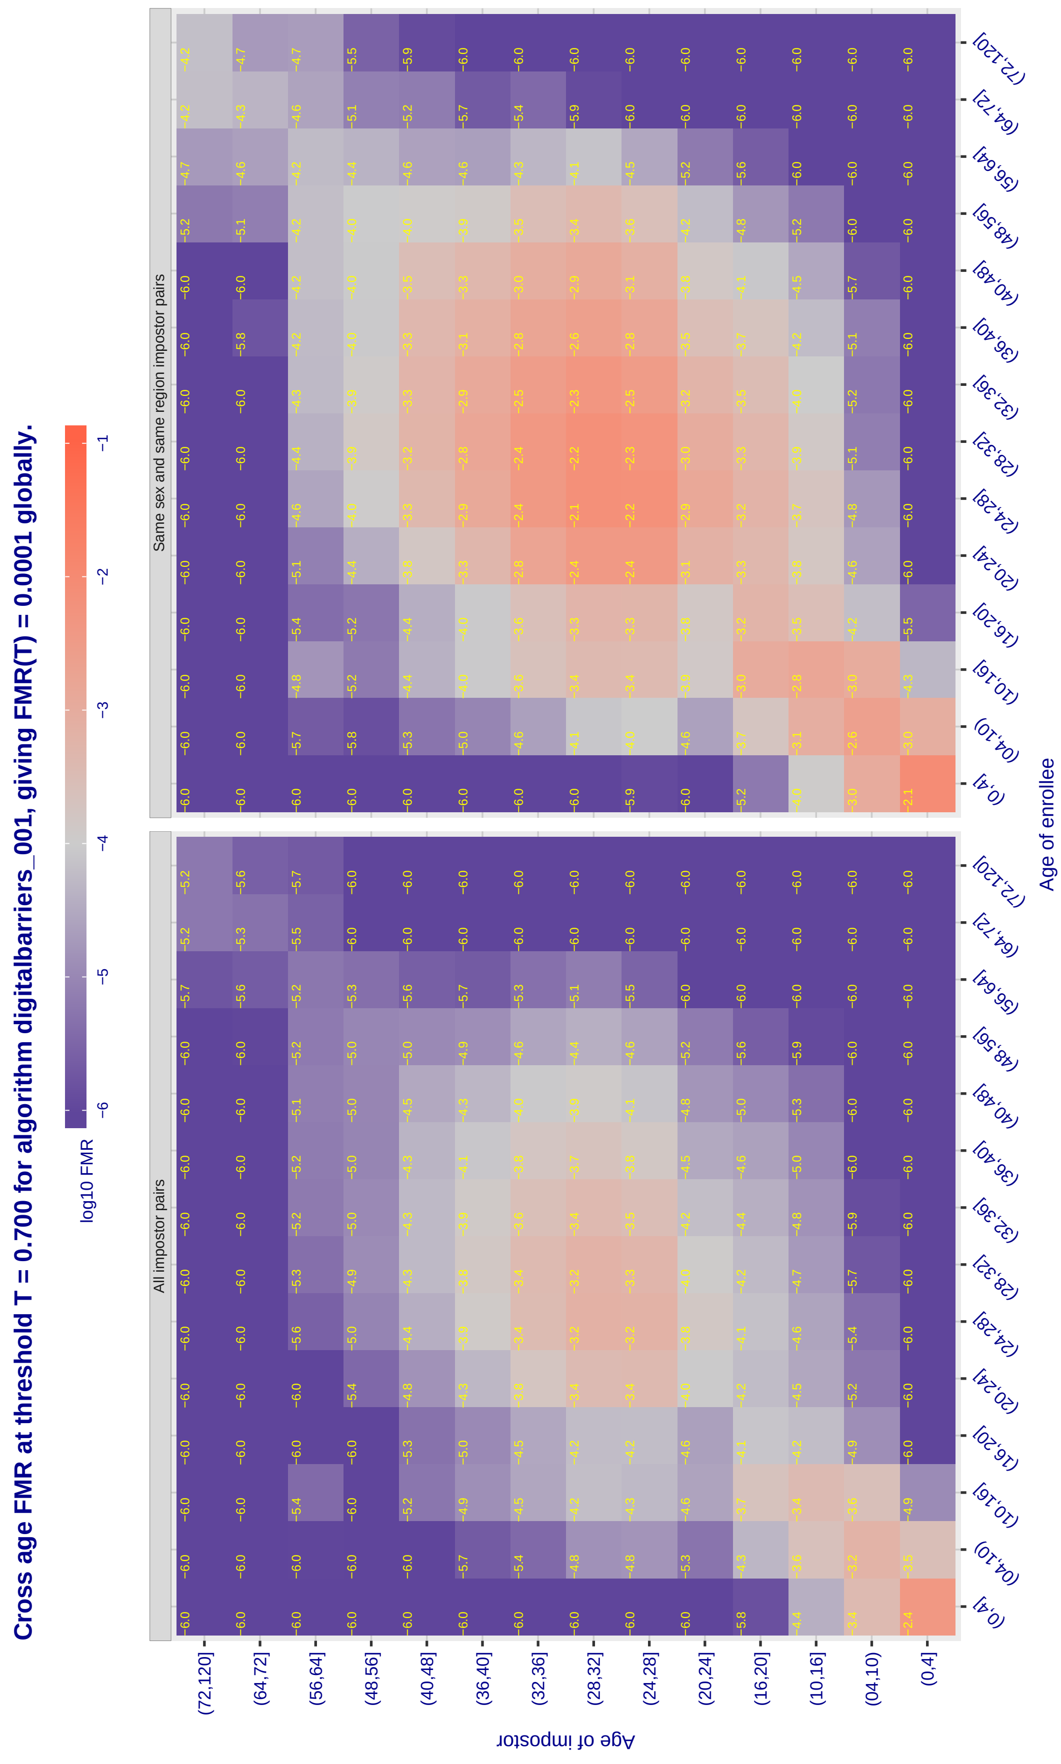

Figure 149: For algorithm digitalbarriers-001 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all $O(10^{10})$ impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

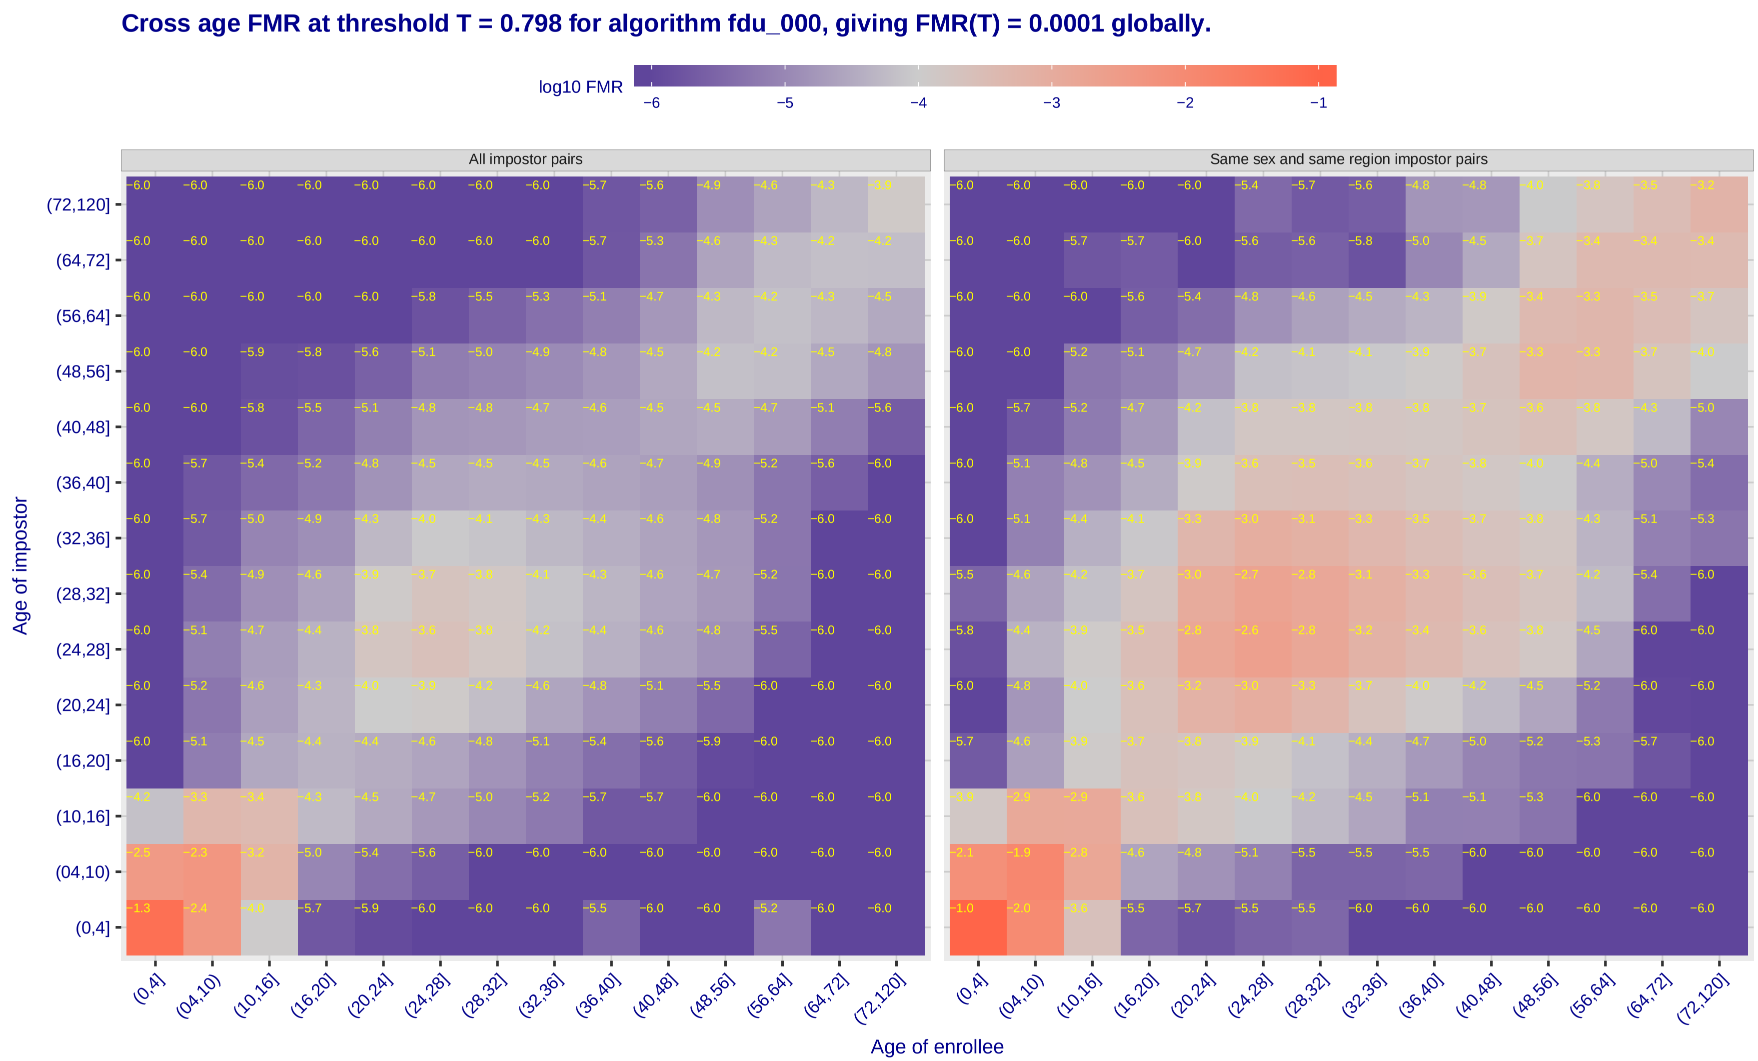

Figure 150: For algorithm fdu-000 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

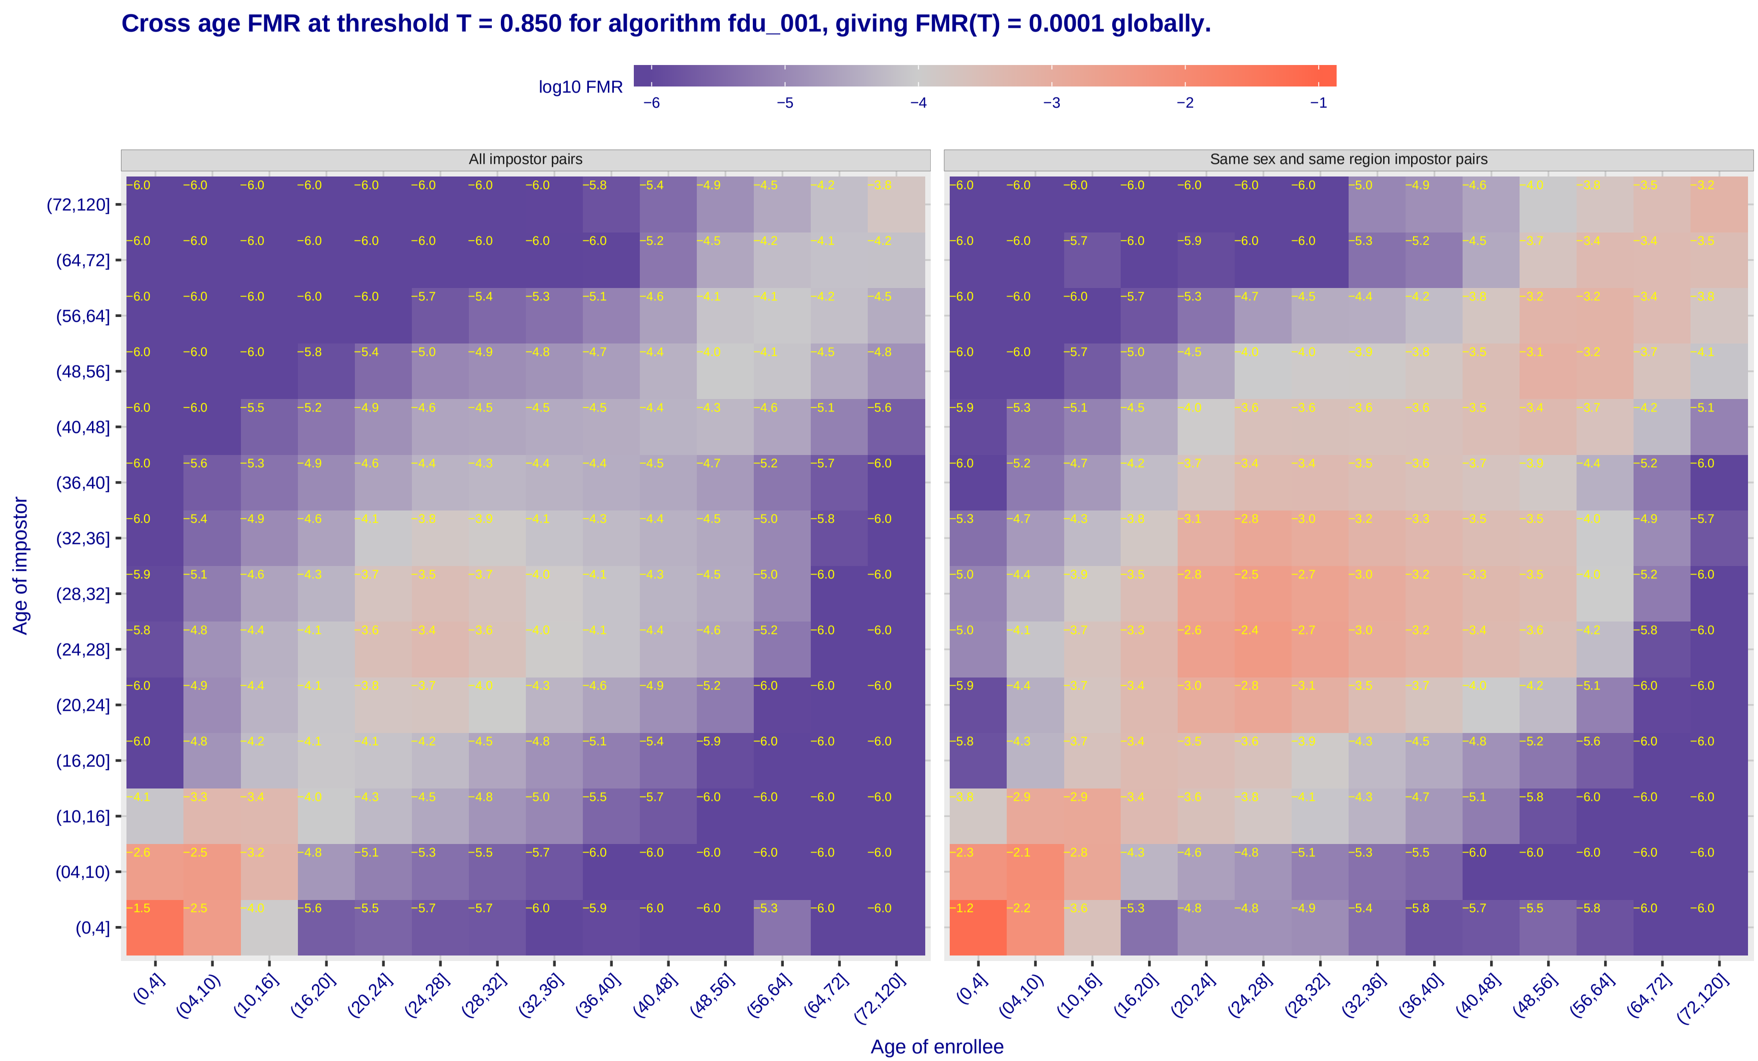

Figure 151: For algorithm fdu-001 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

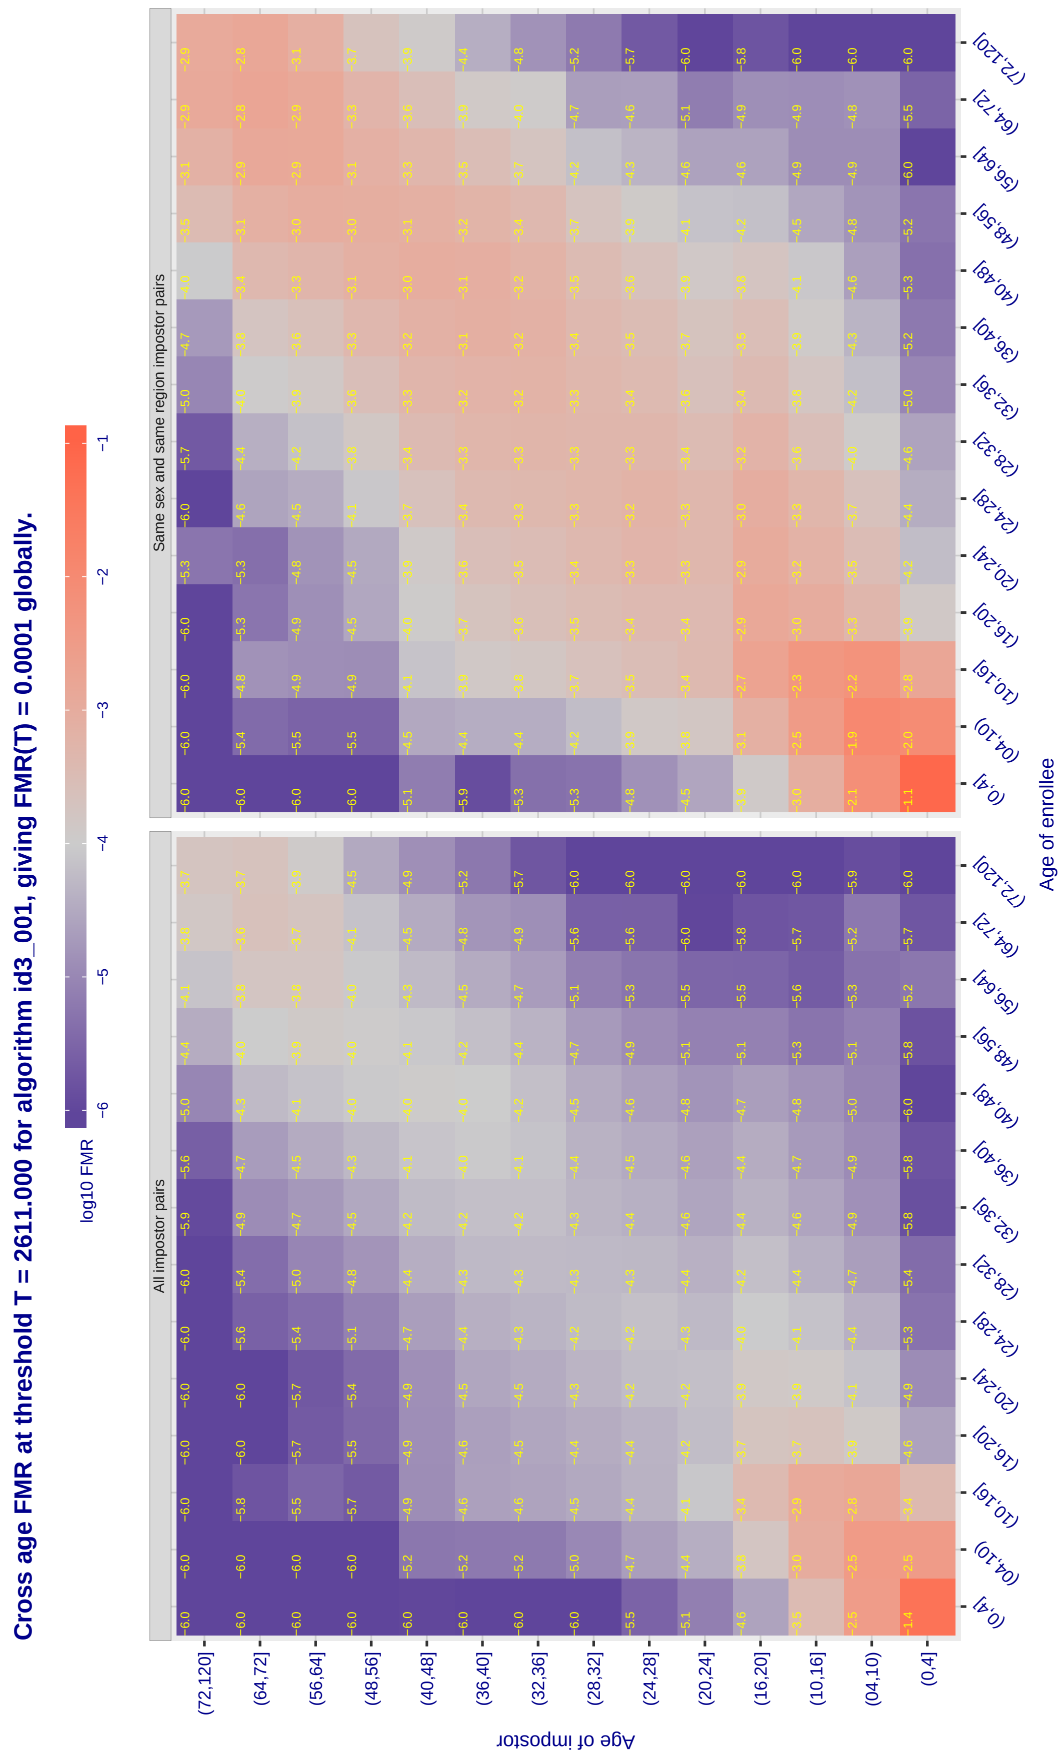

Figure 152: For algorithm id3-001 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

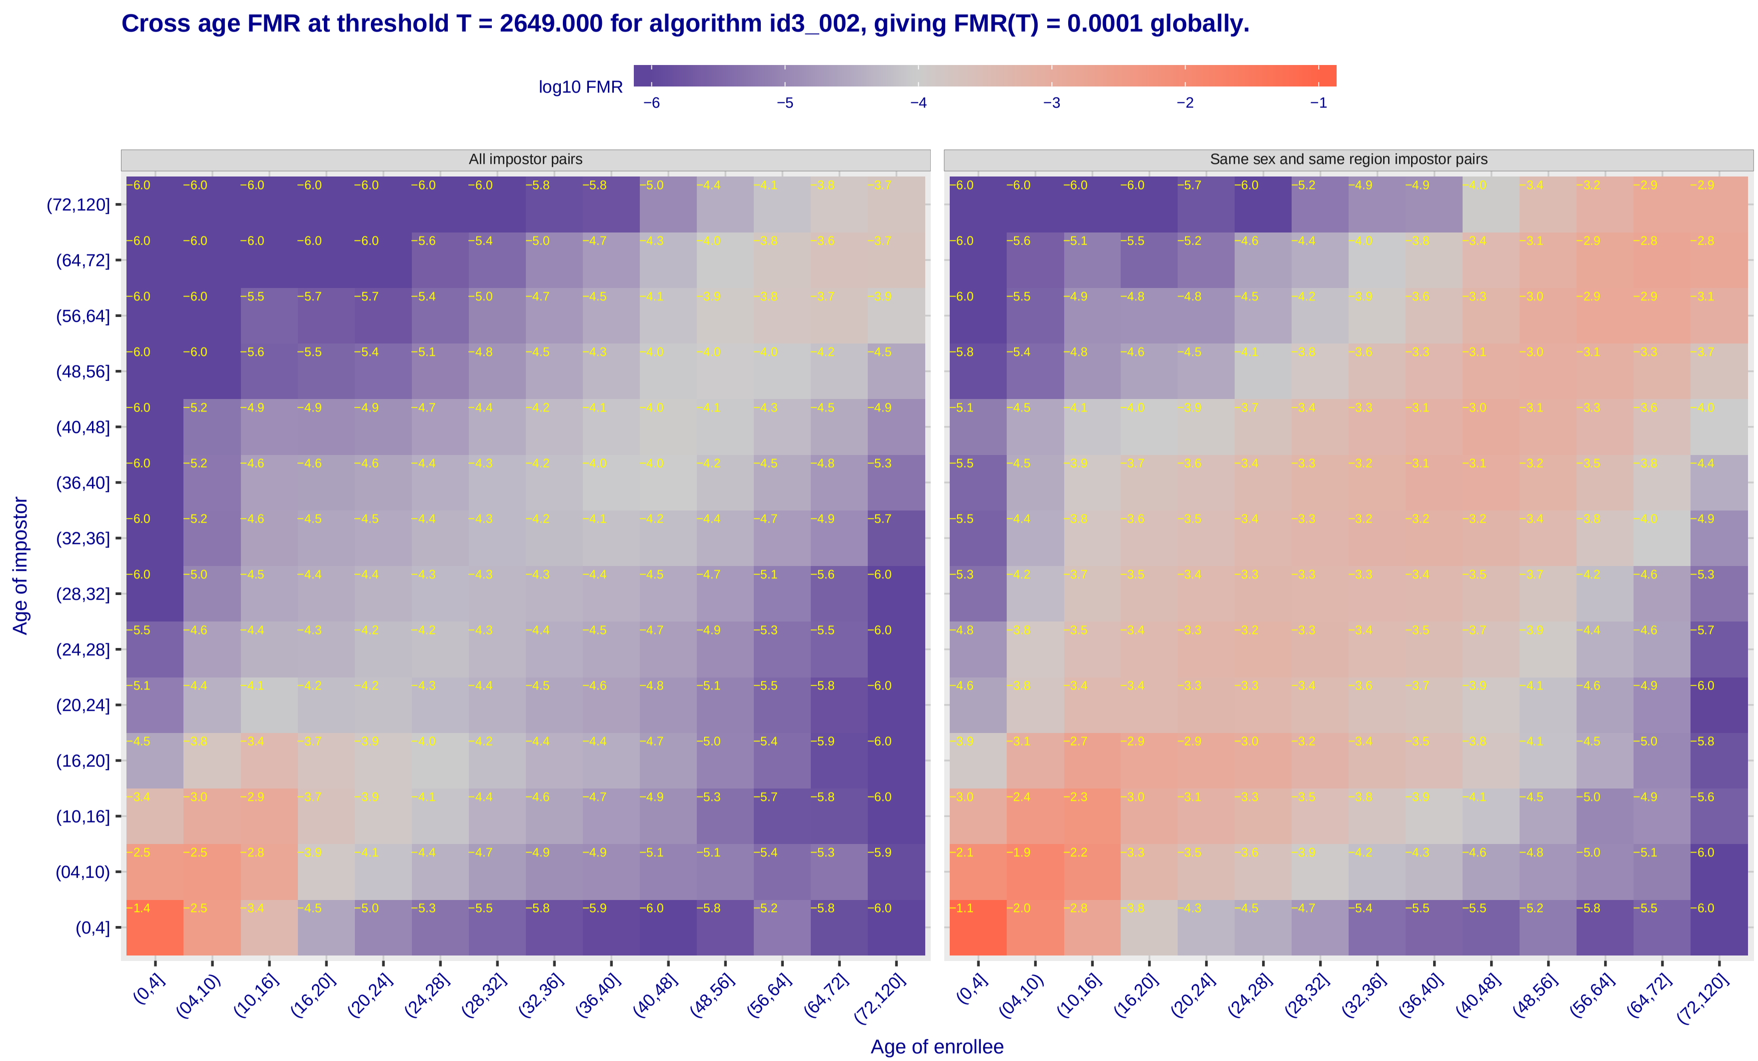

Figure 153: For algorithm id3-002 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all $O(10^{10})$ impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

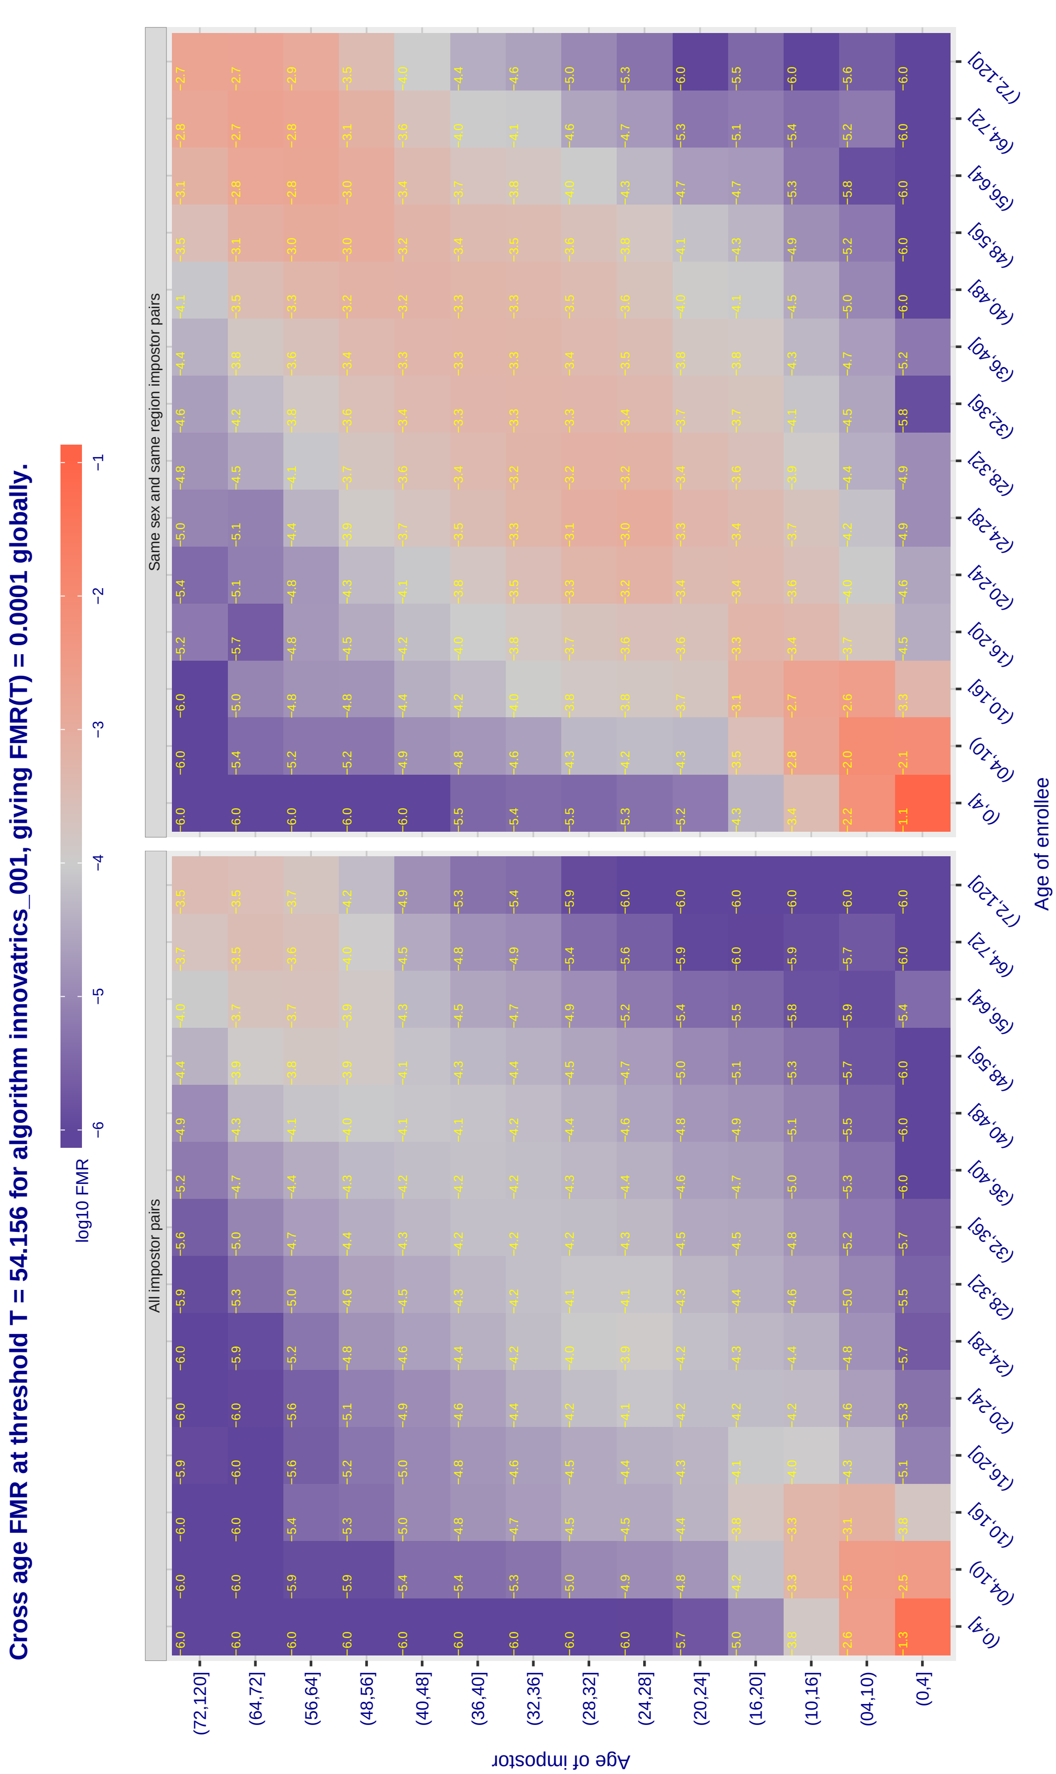

Figure 154: For algorithm innovatrics-001 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

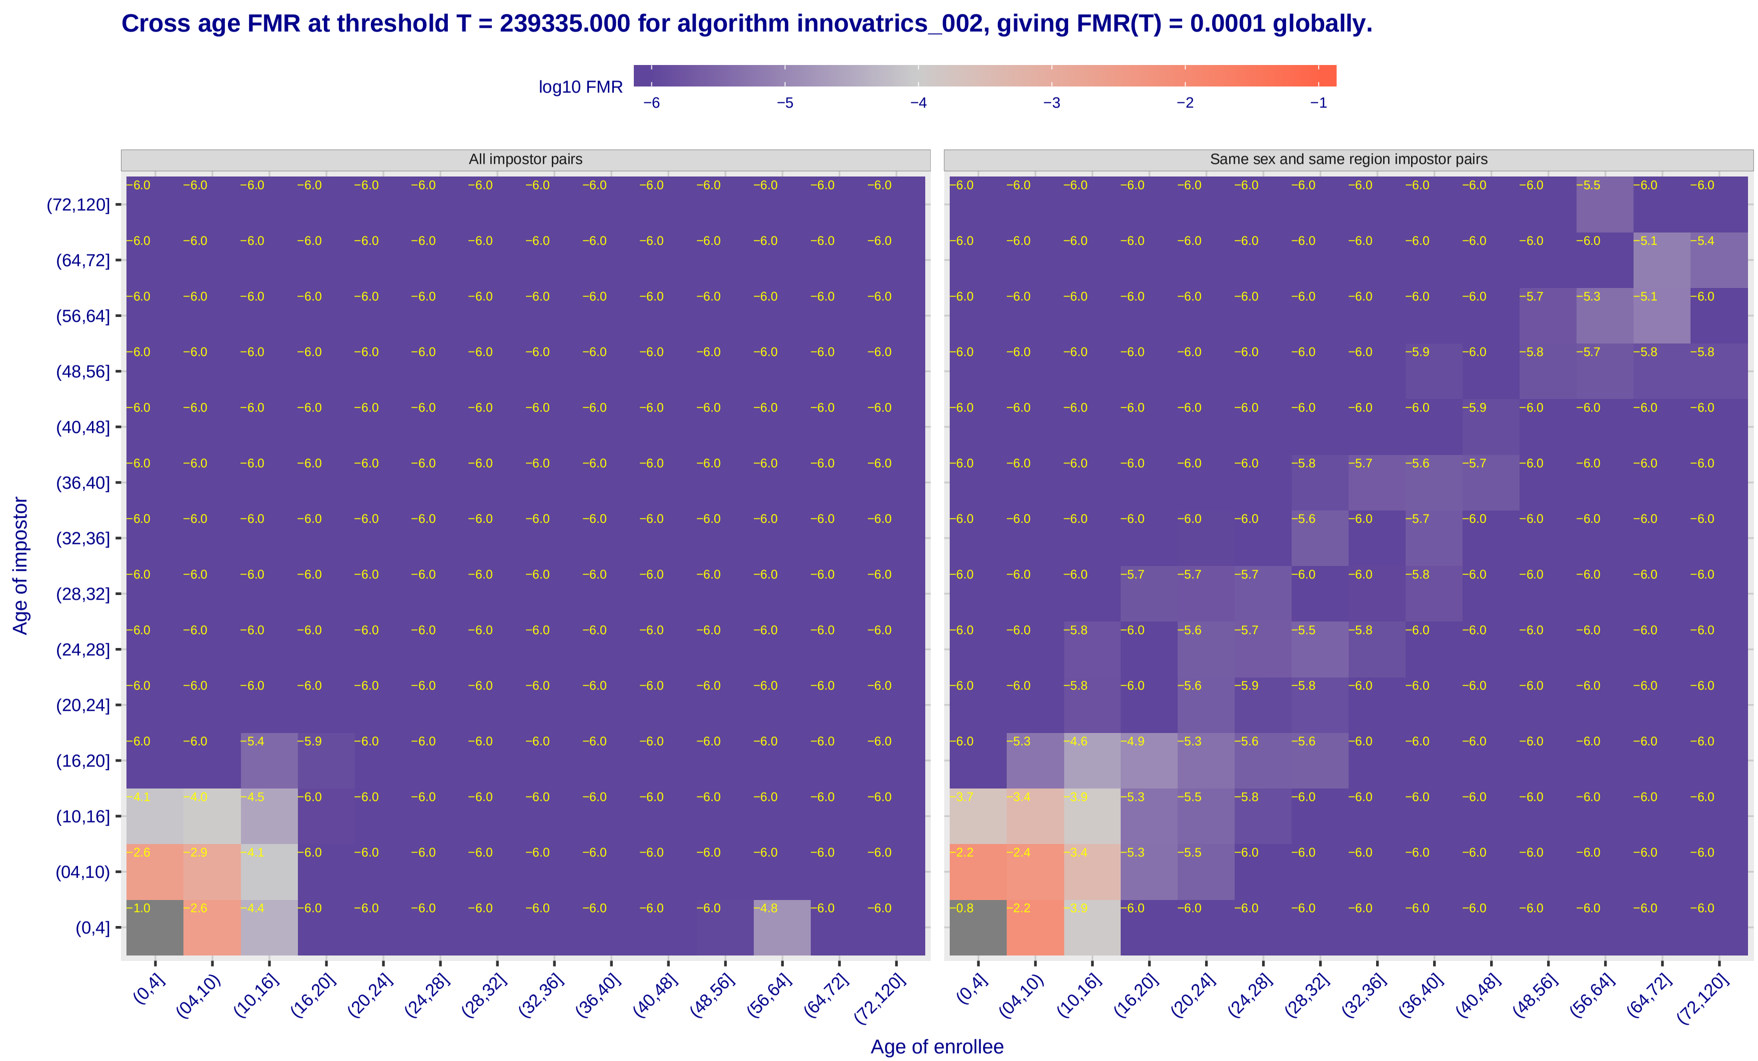

Figure 155: For algorithm innovatrics-002 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

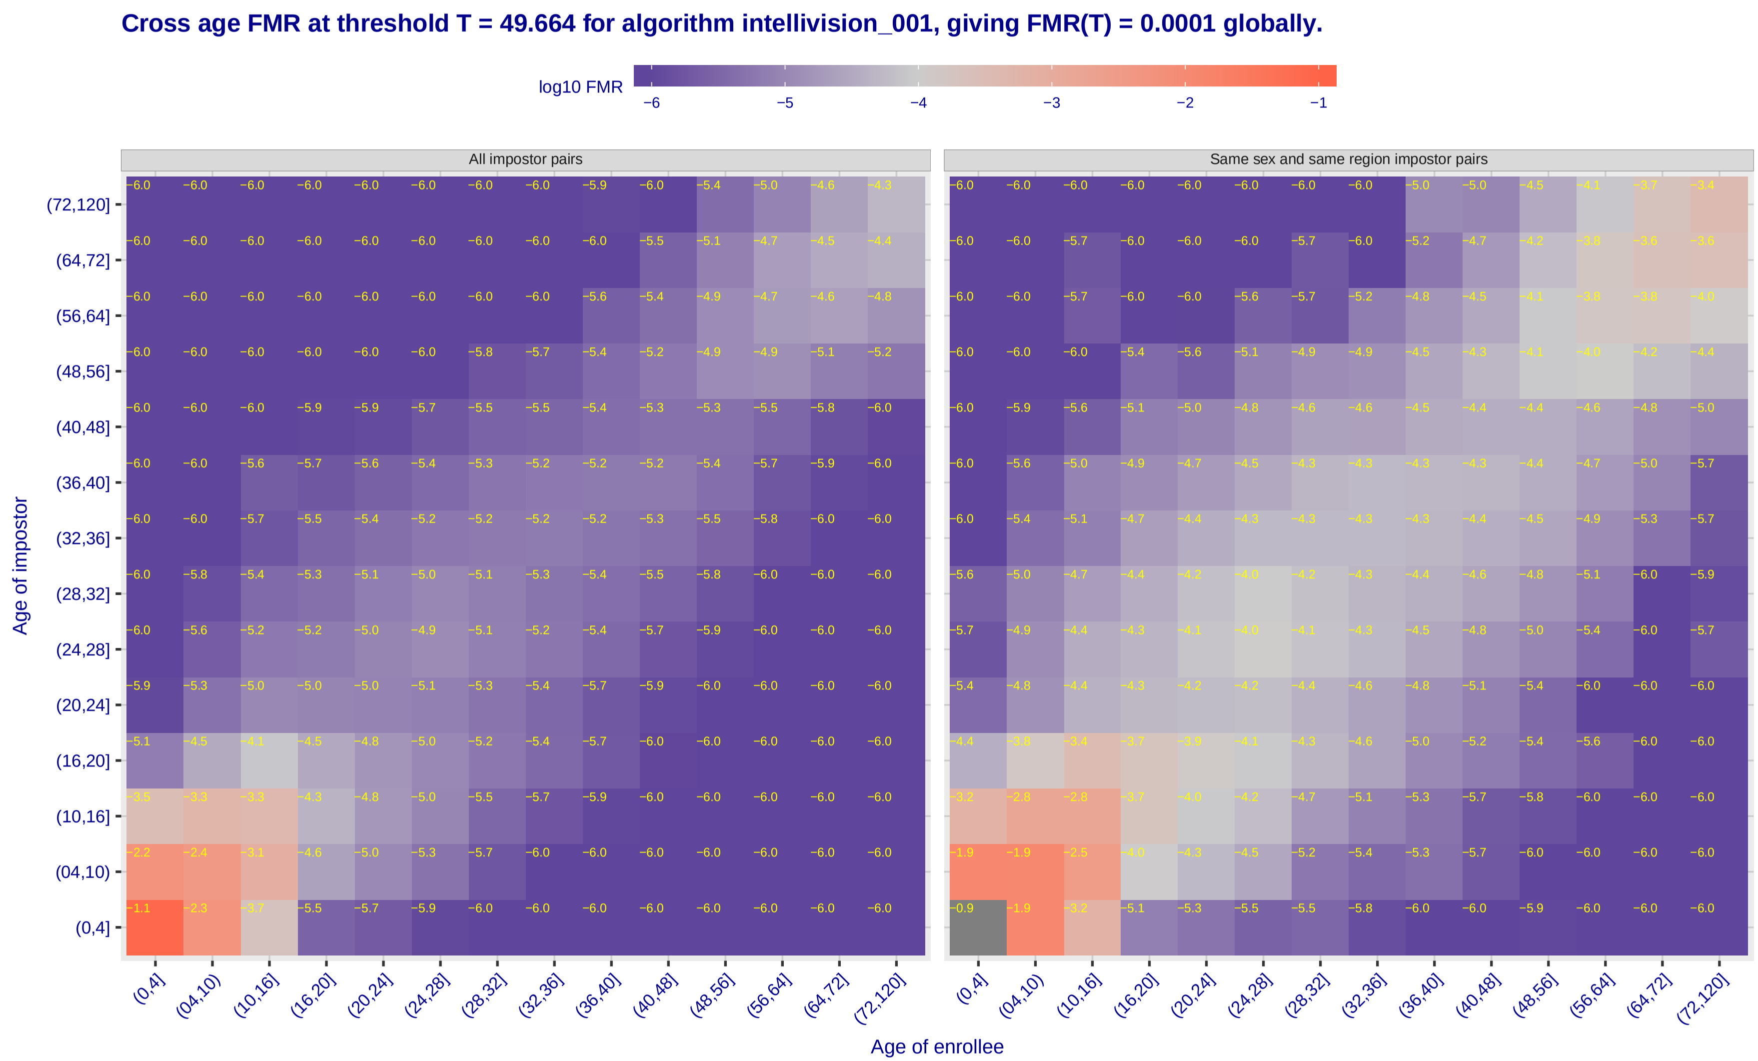

Figure 156: For algorithm intellivision-001 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

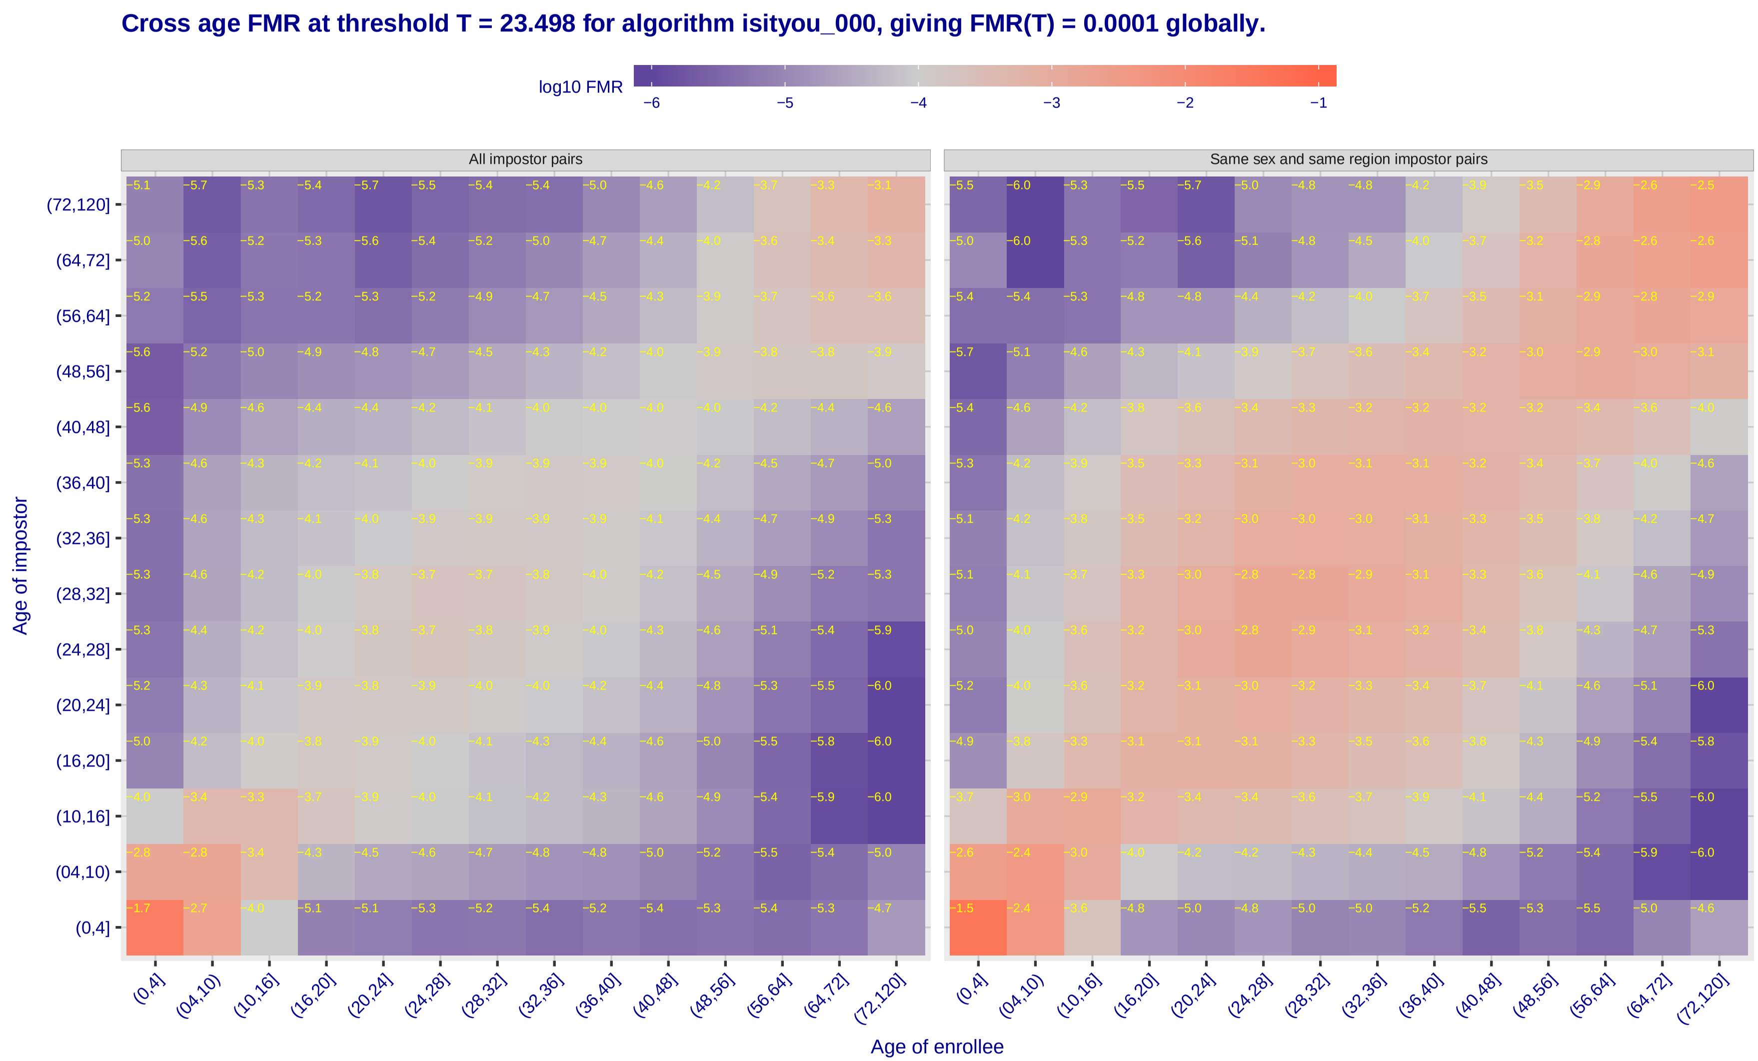

Figure 157: For algorithm isityou-000 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

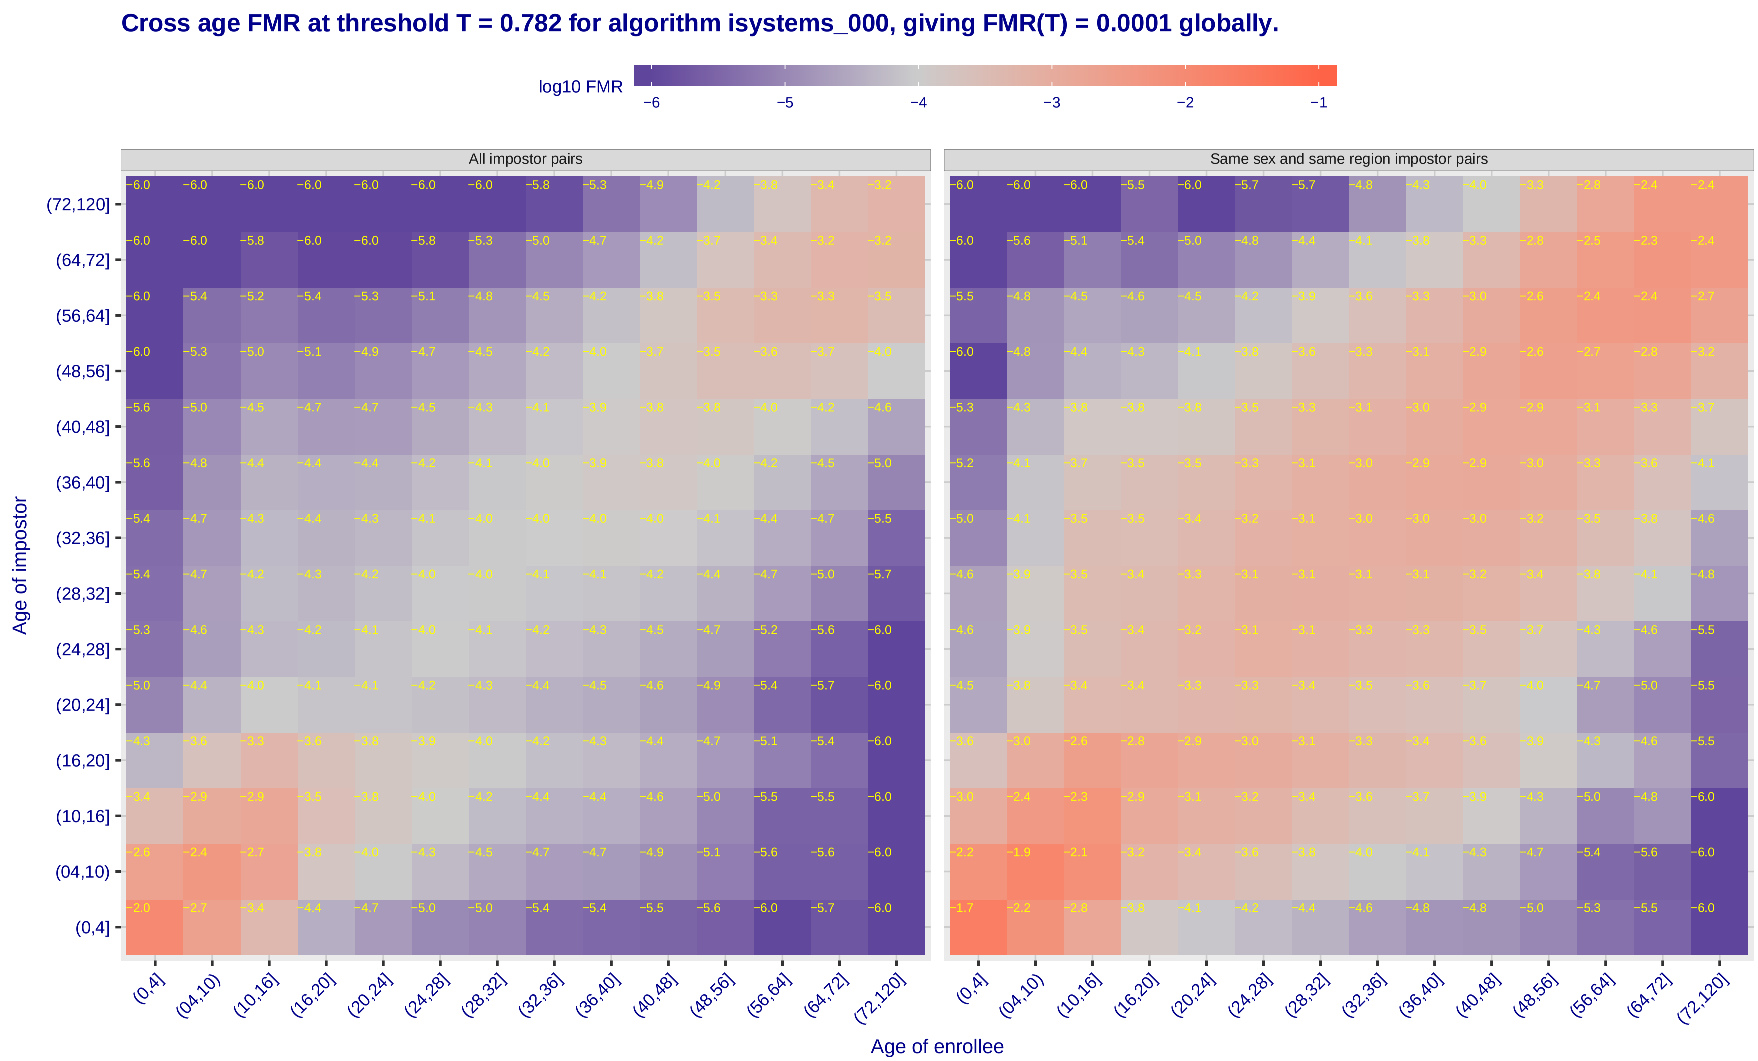

Figure 158: For algorithm isystems-000 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all $O(10^{10})$ impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

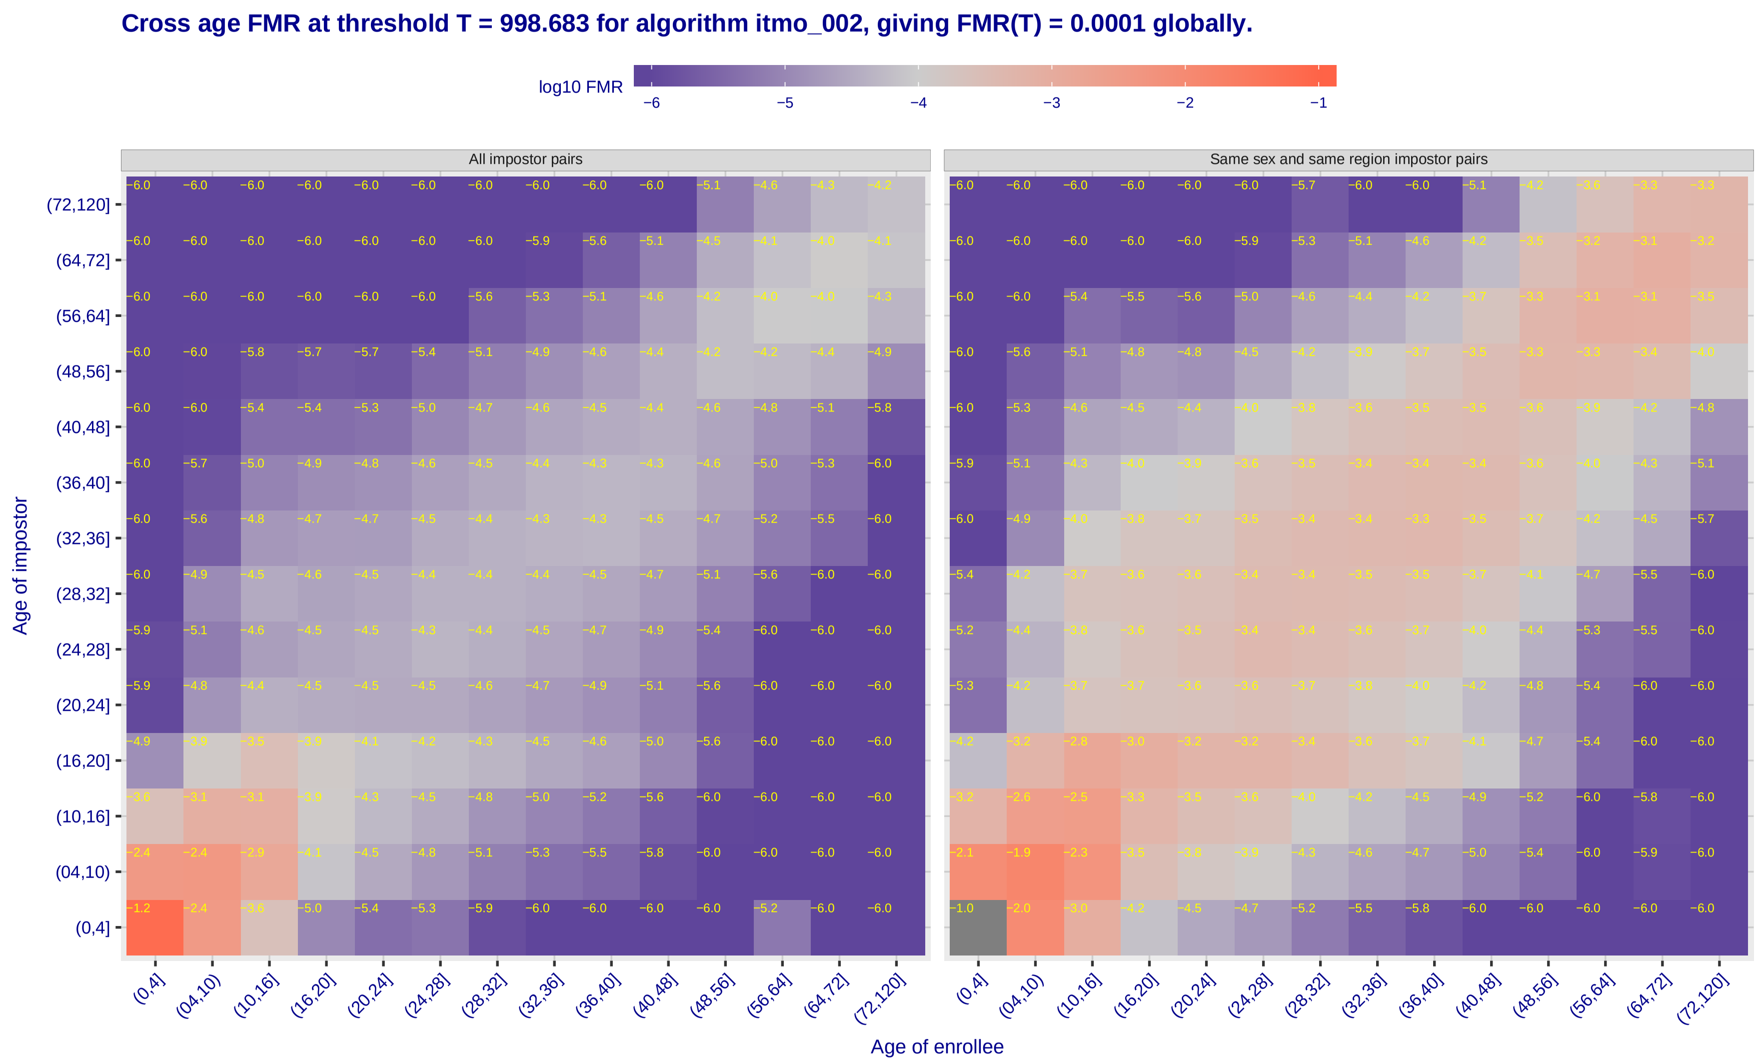

Figure 159: For algorithm itmo-002 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

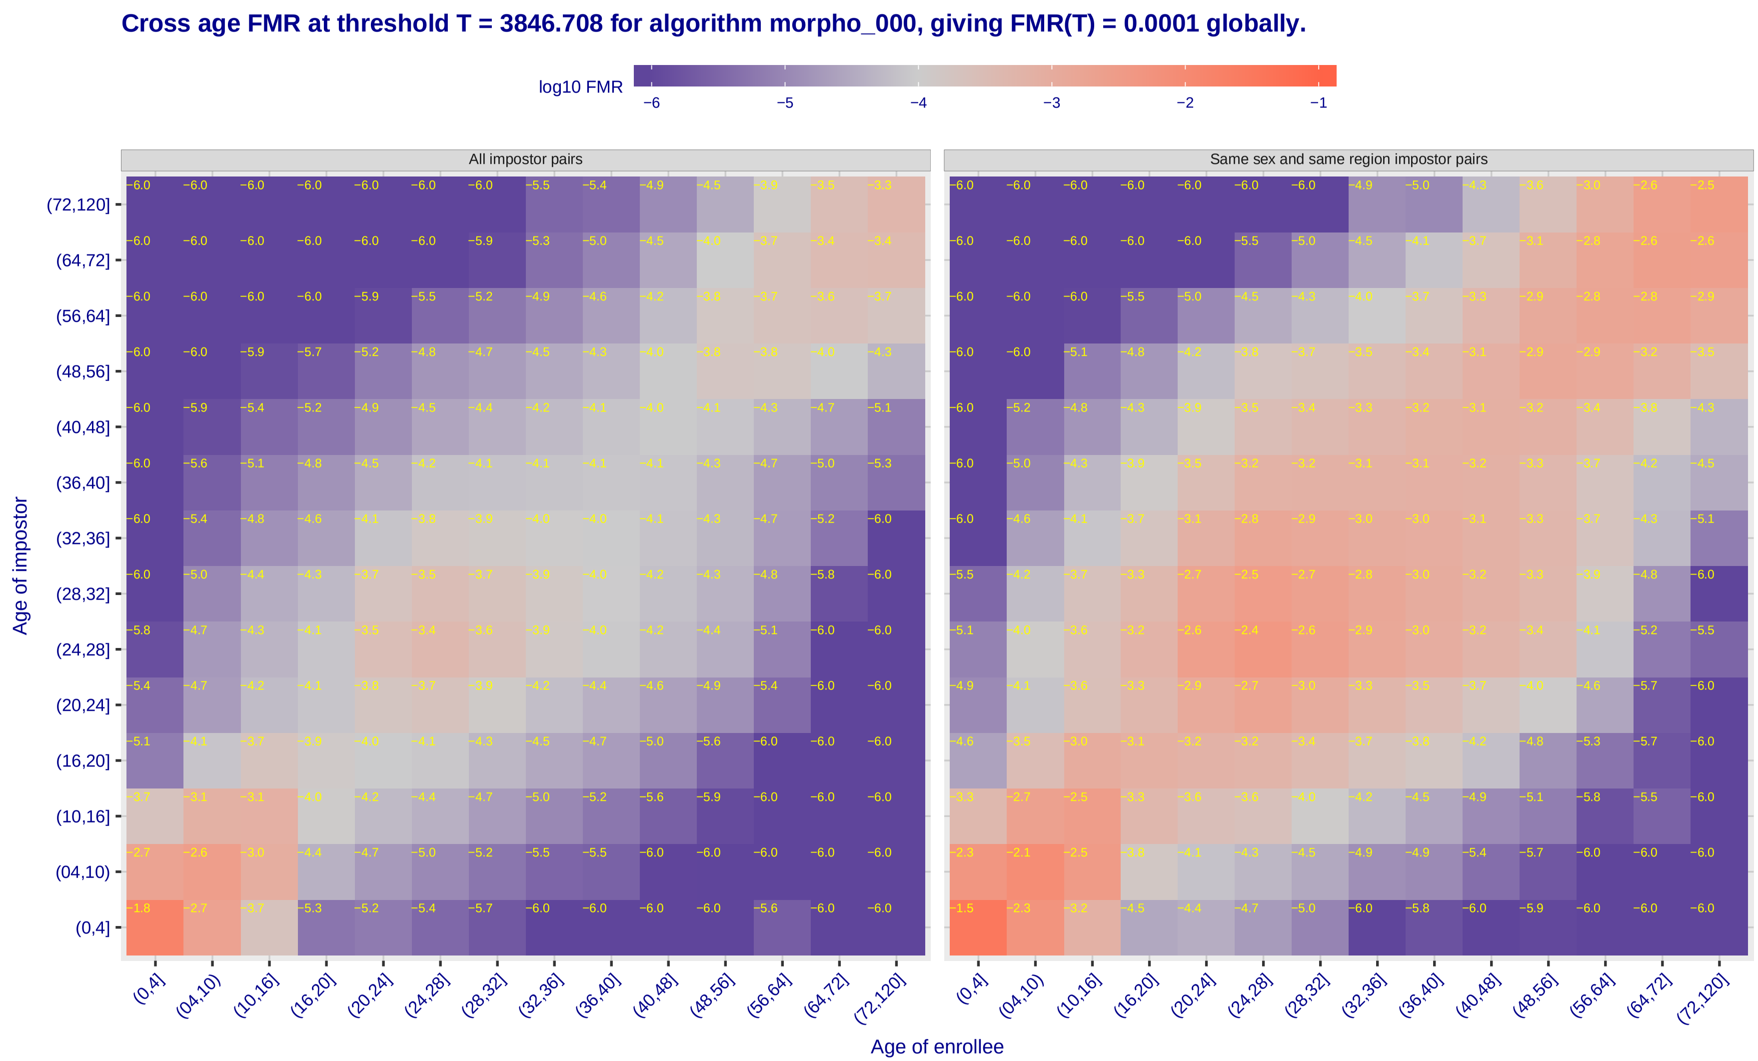

Figure 160: For algorithm morpho-000 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all $O(10^{10})$ impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

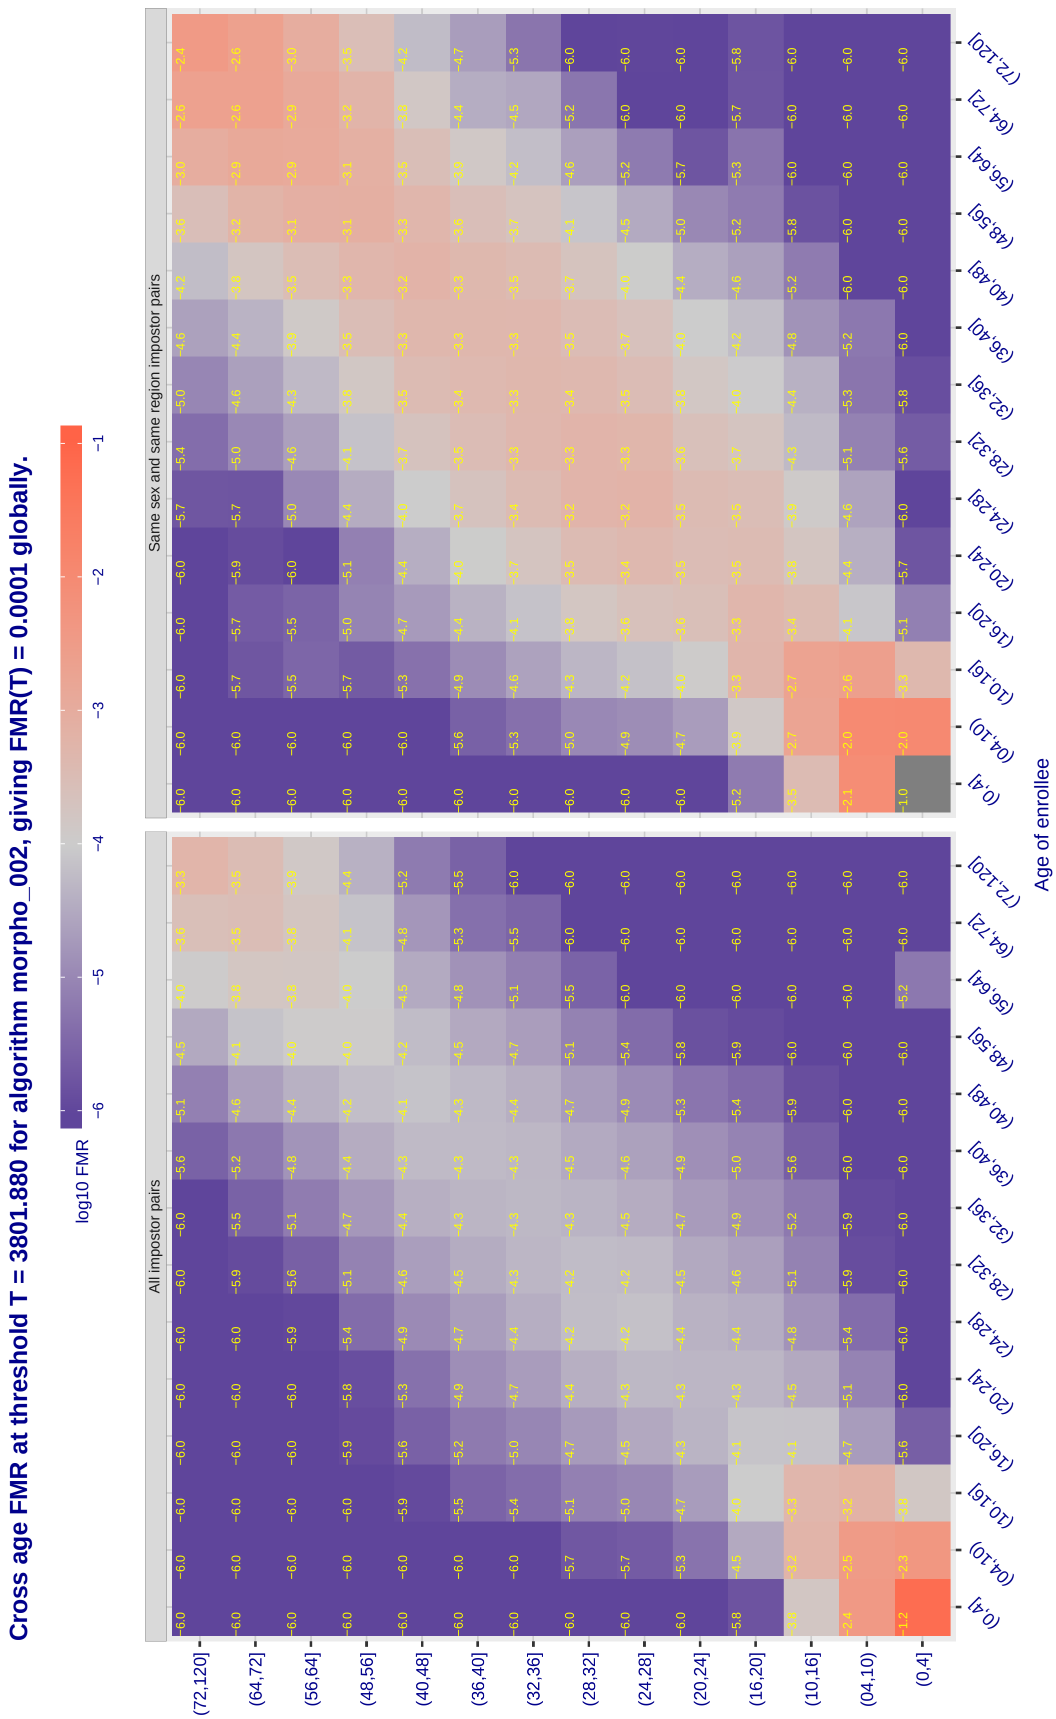

Figure 161: For algorithm morpho-002 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

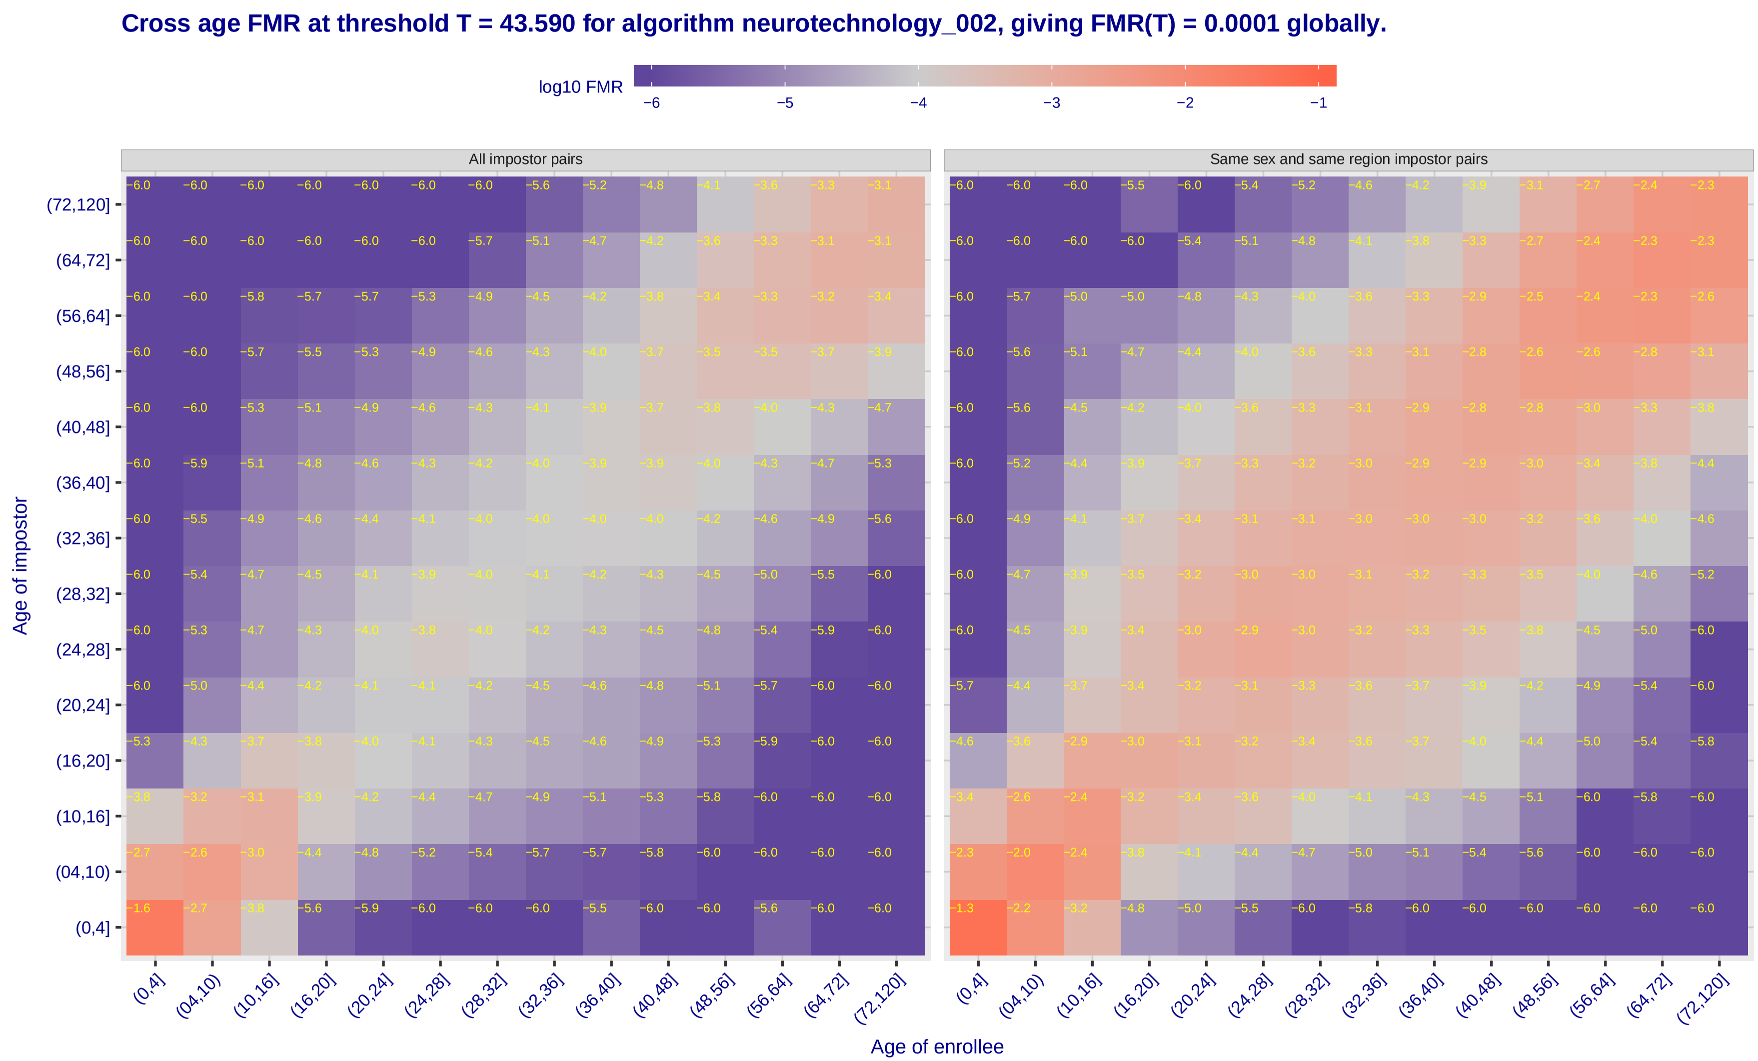

Figure 162: For algorithm neurotechnology-002 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

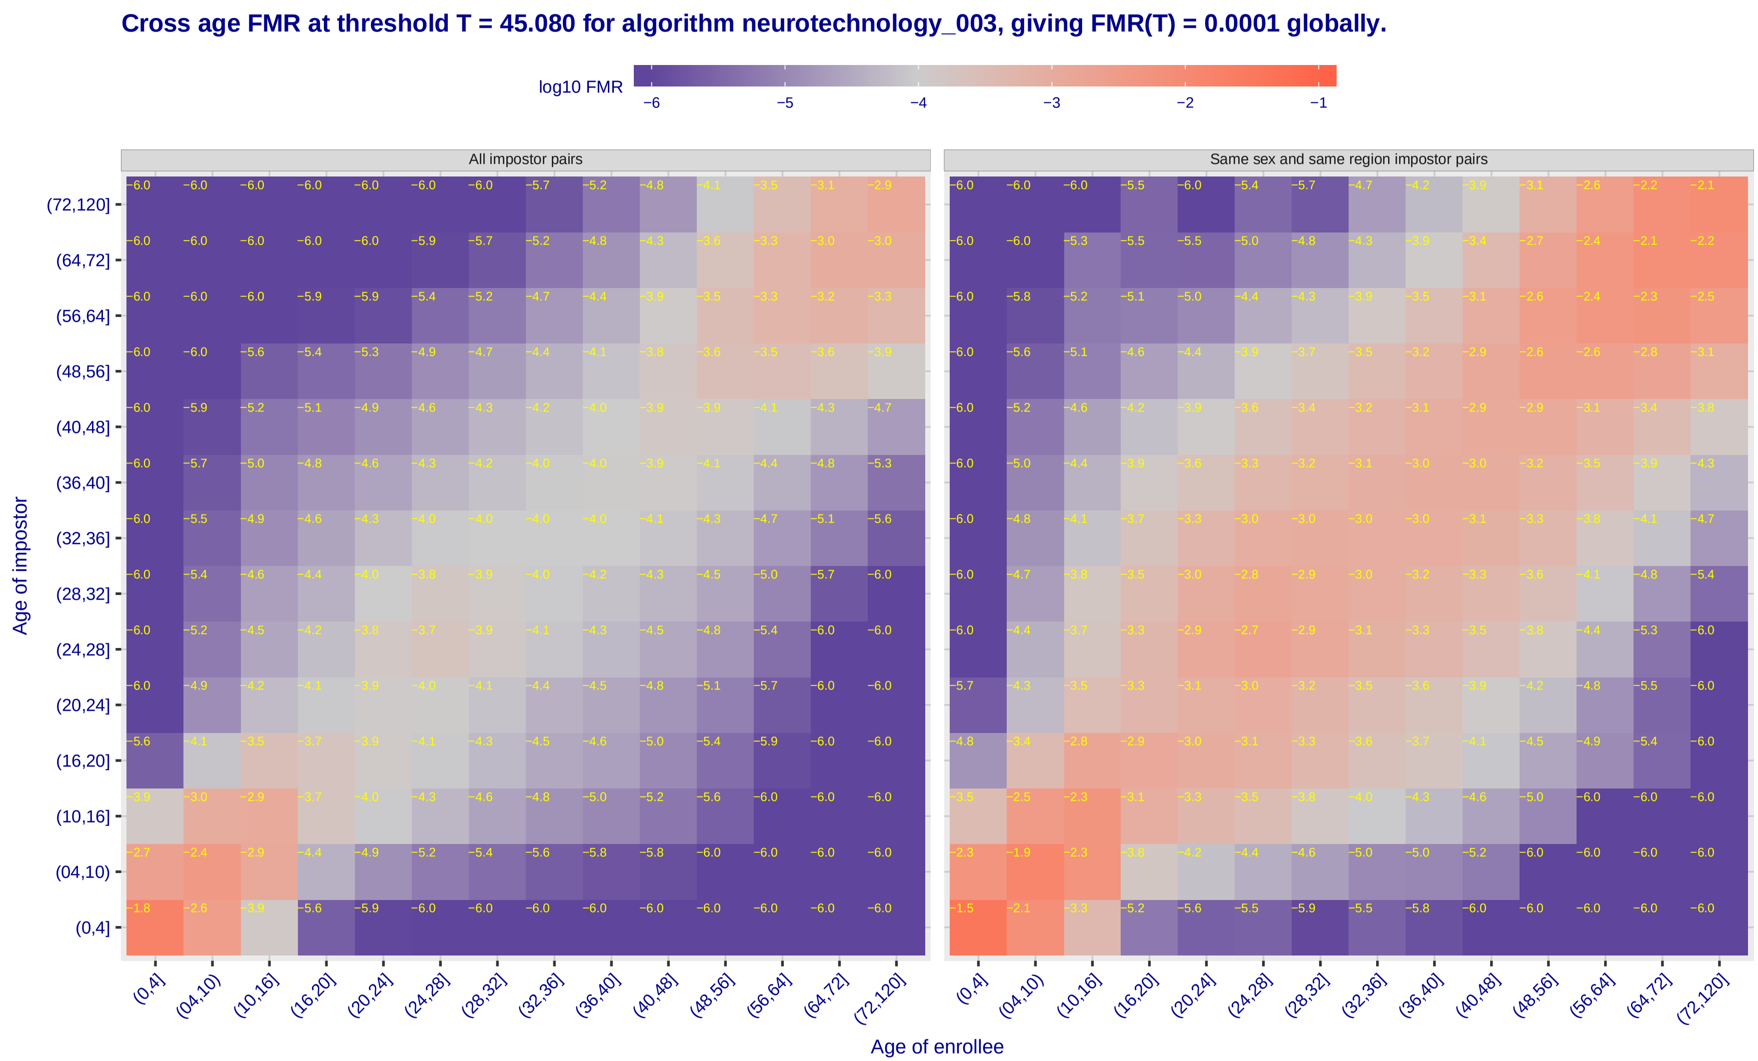

Figure 163: For algorithm neurotechnology-003 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

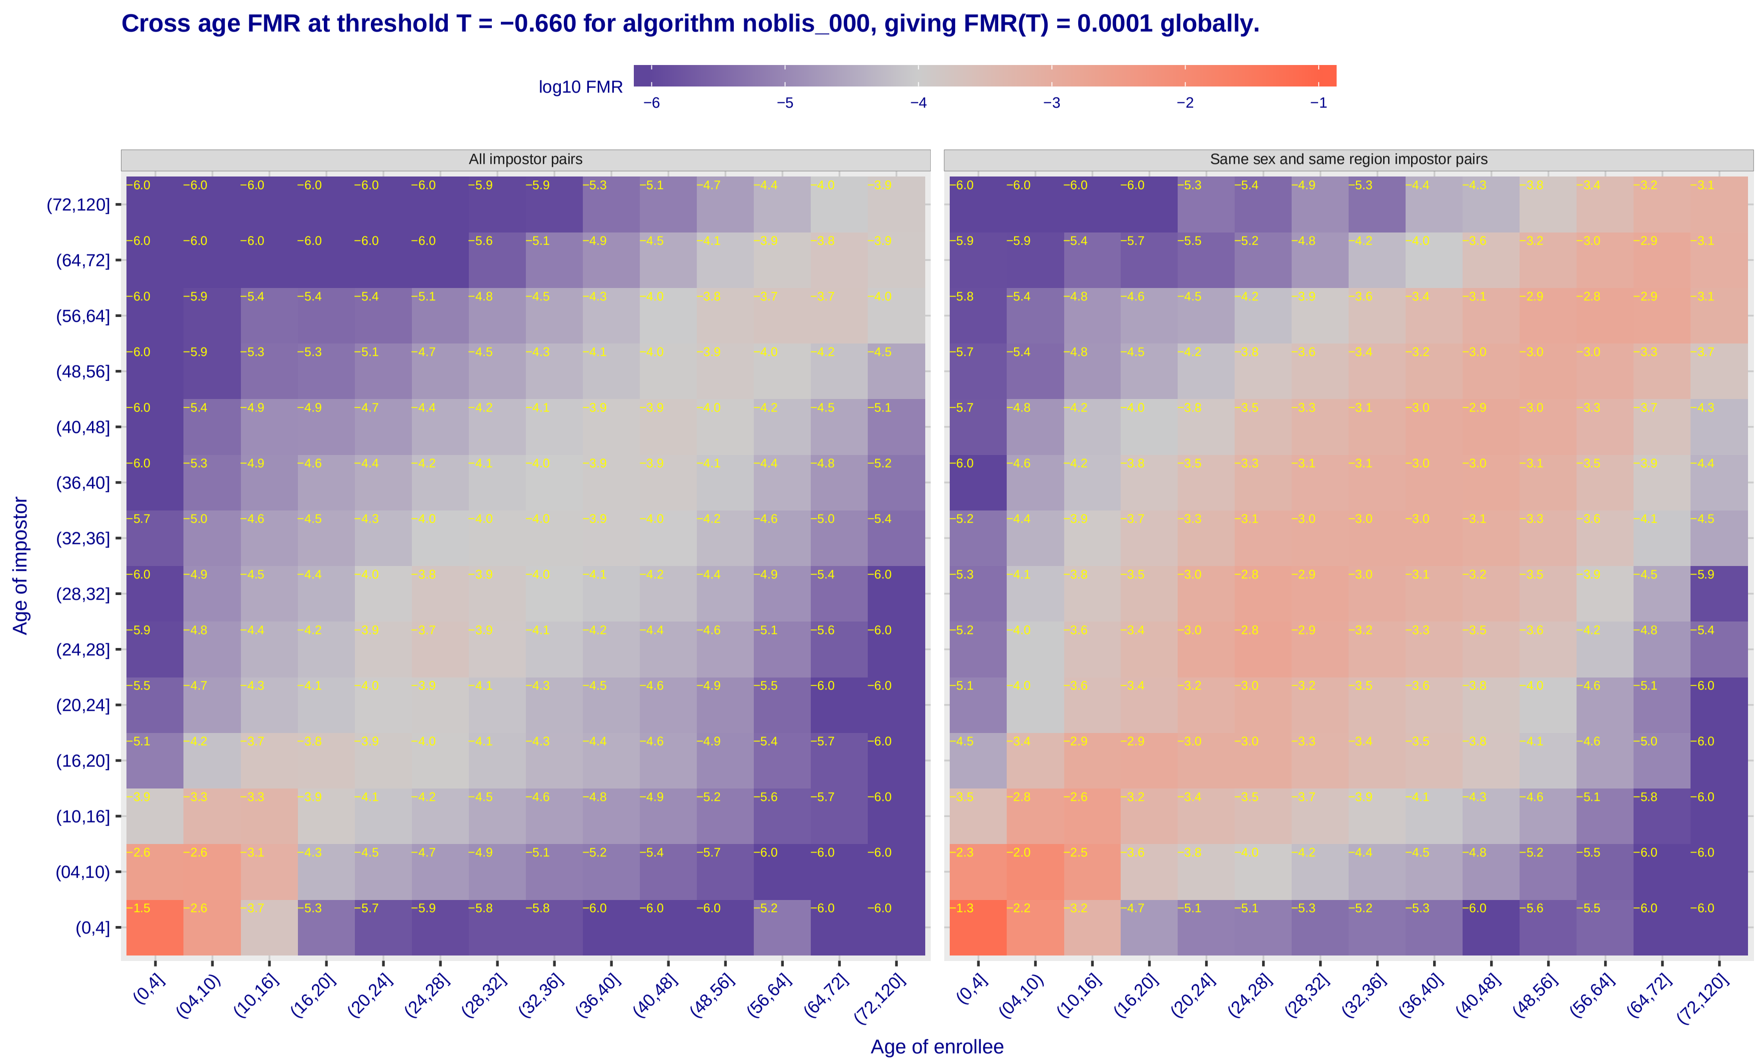

Figure 164: For algorithm noblis-000 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all $O(10^{10})$ impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

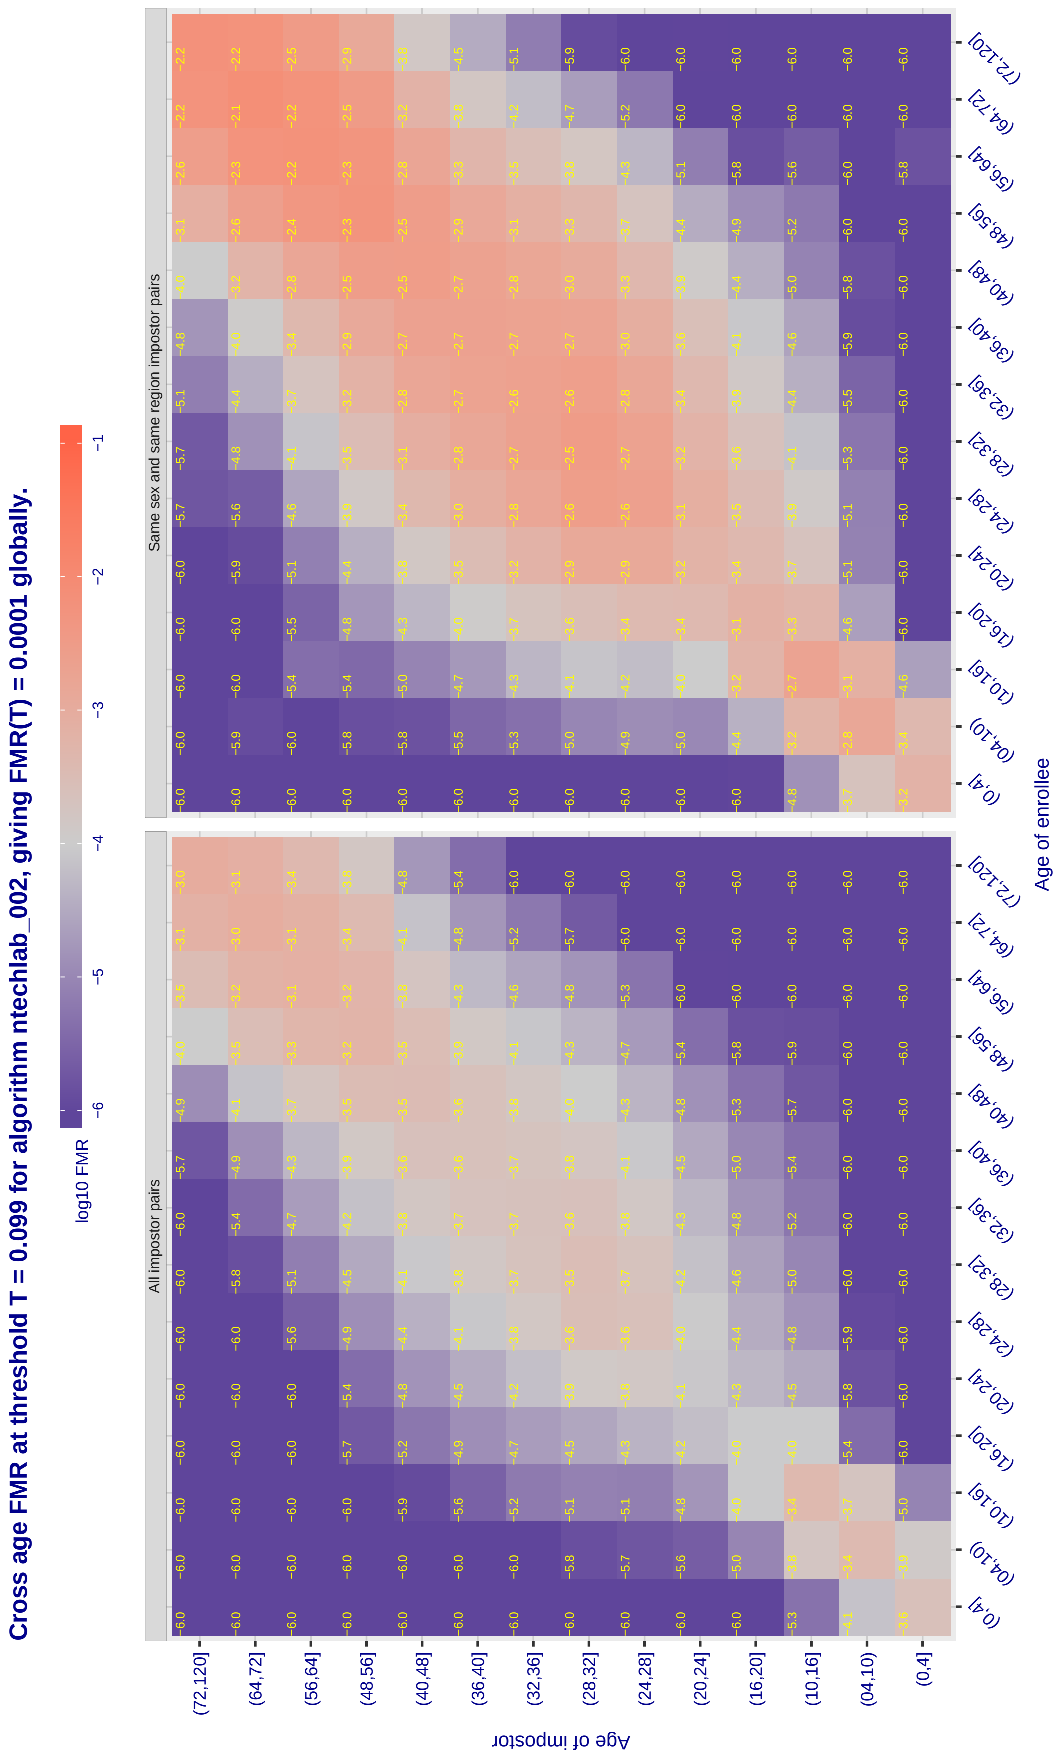

Figure 165: For algorithm ntechlab-002 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

FNMR(T) "False non-match rate"
FMR(T) "False match rate"

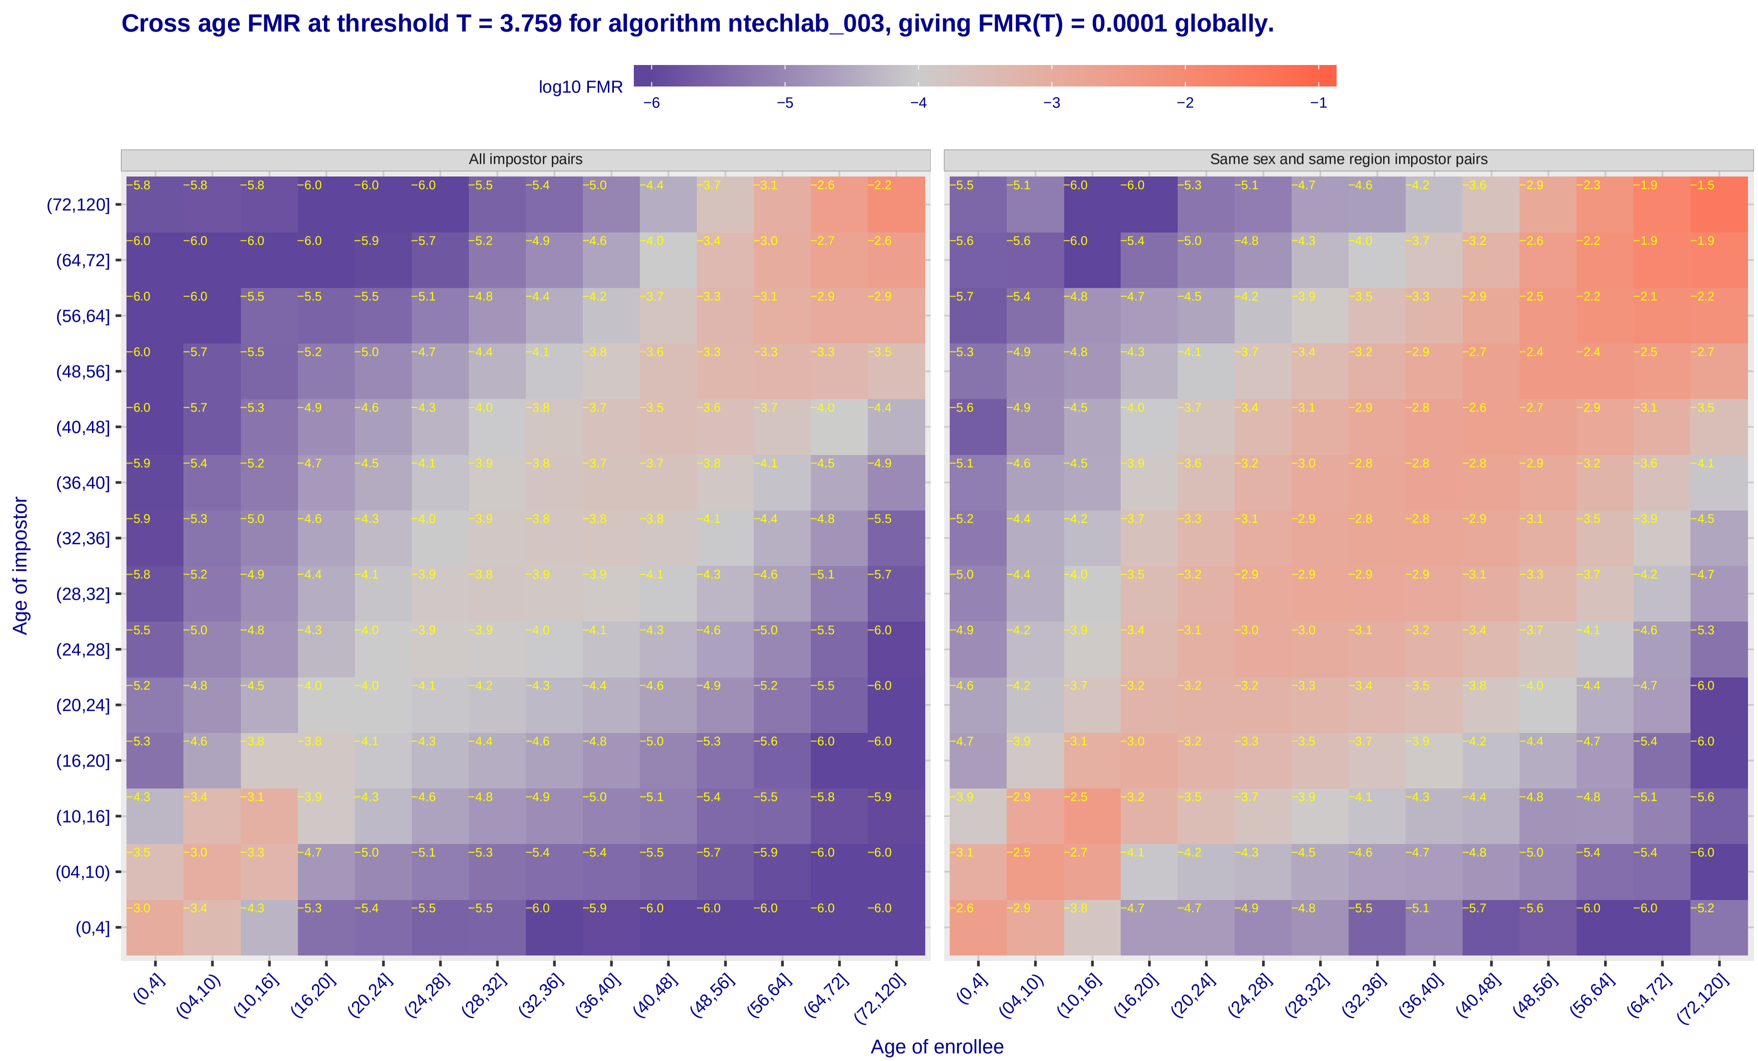

Figure 166: For algorithm ntechlab-003 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all $O(10^{10})$ impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

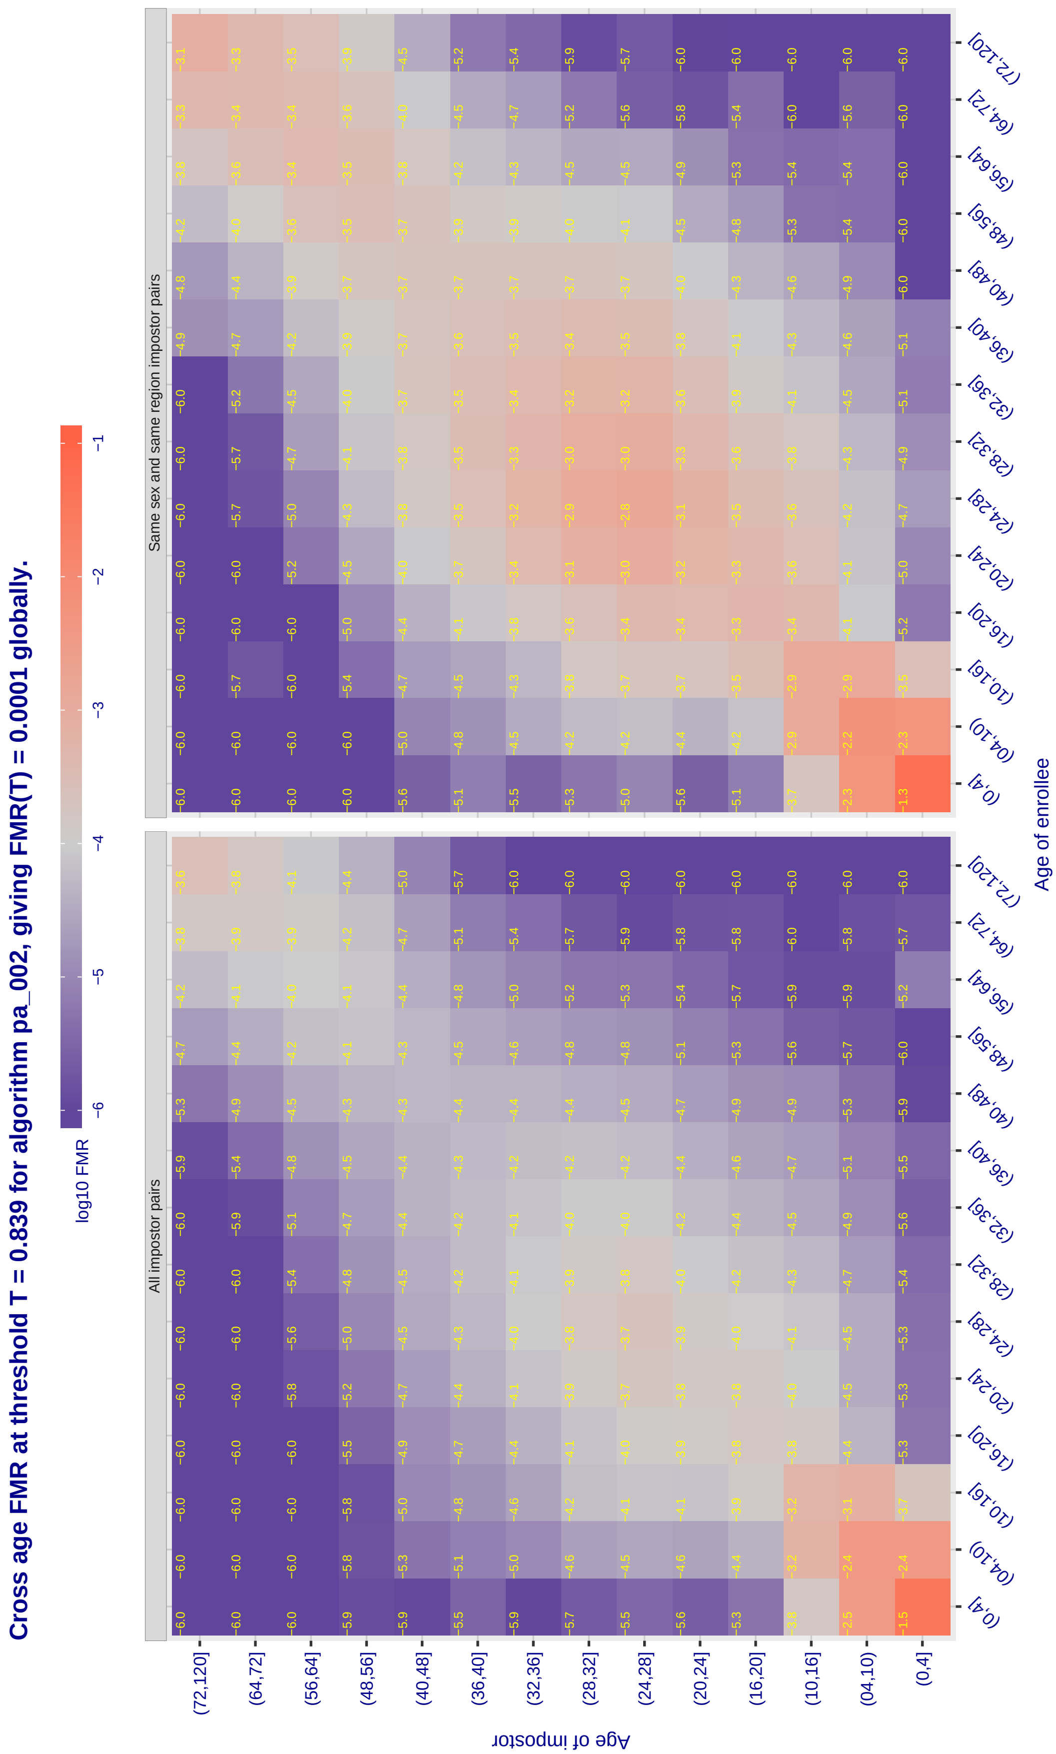

Figure 167: For algorithm pa-002 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

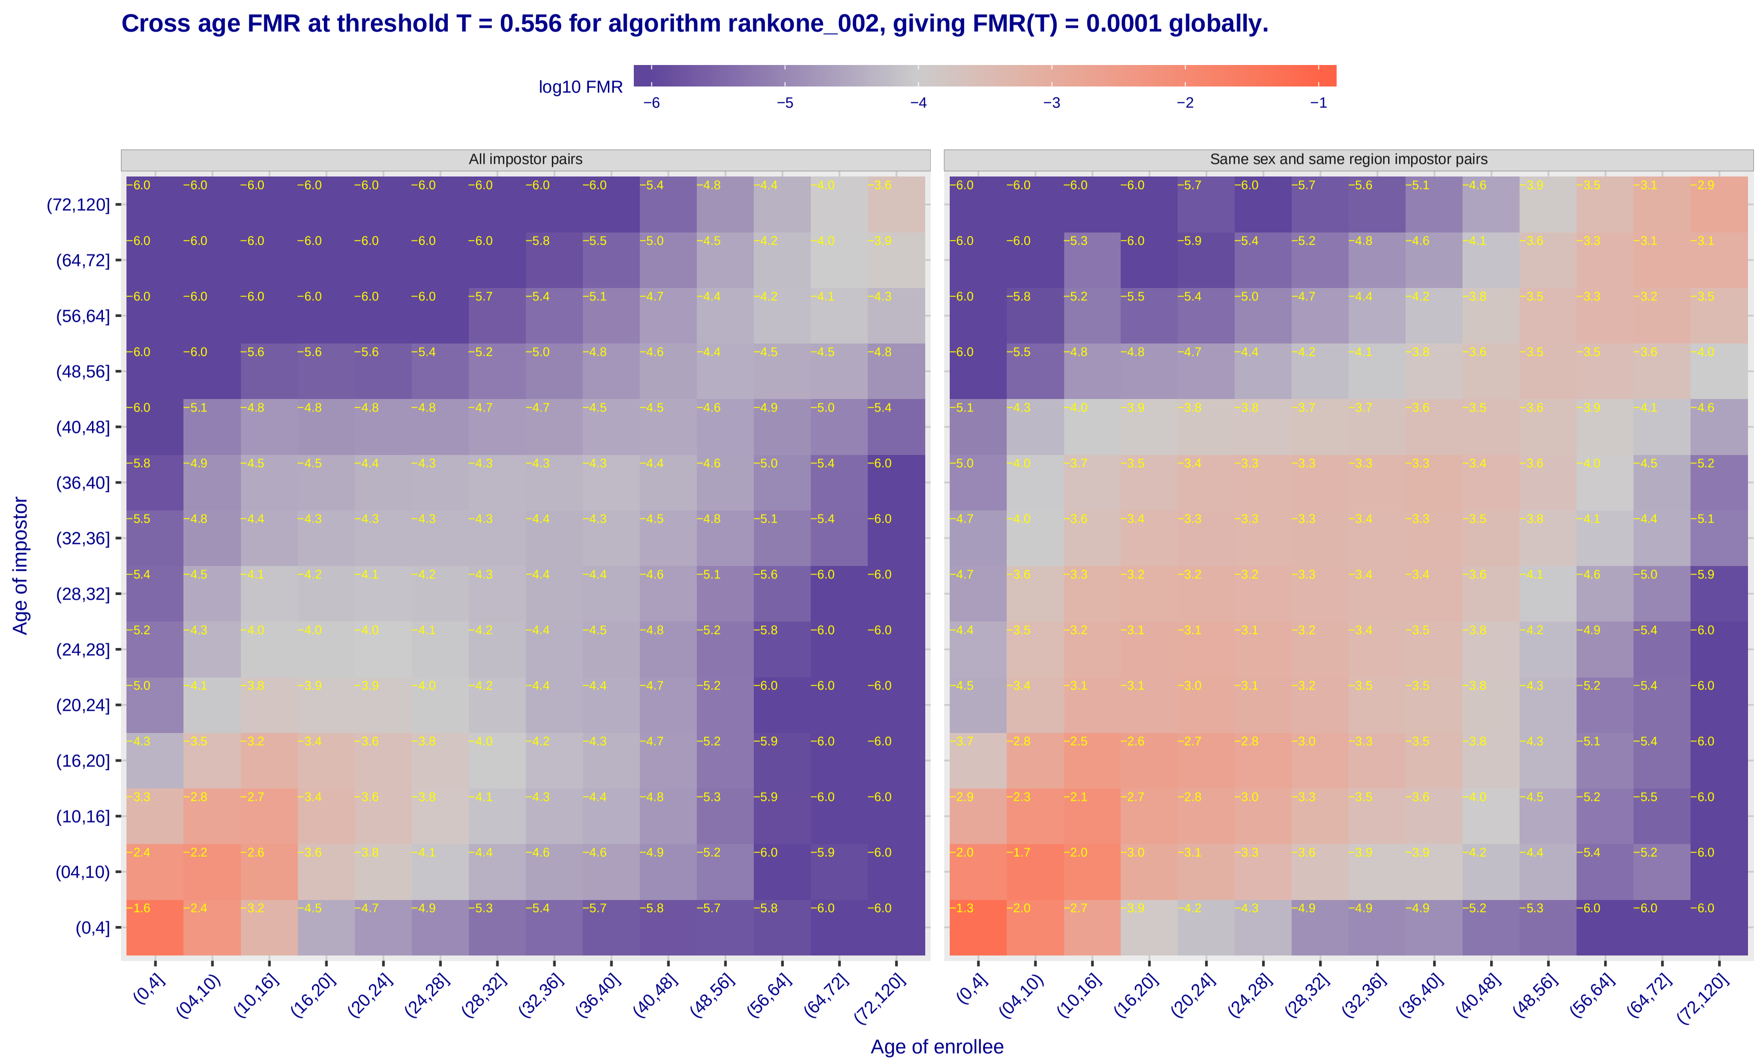

Figure 168: For algorithm rankone-002 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

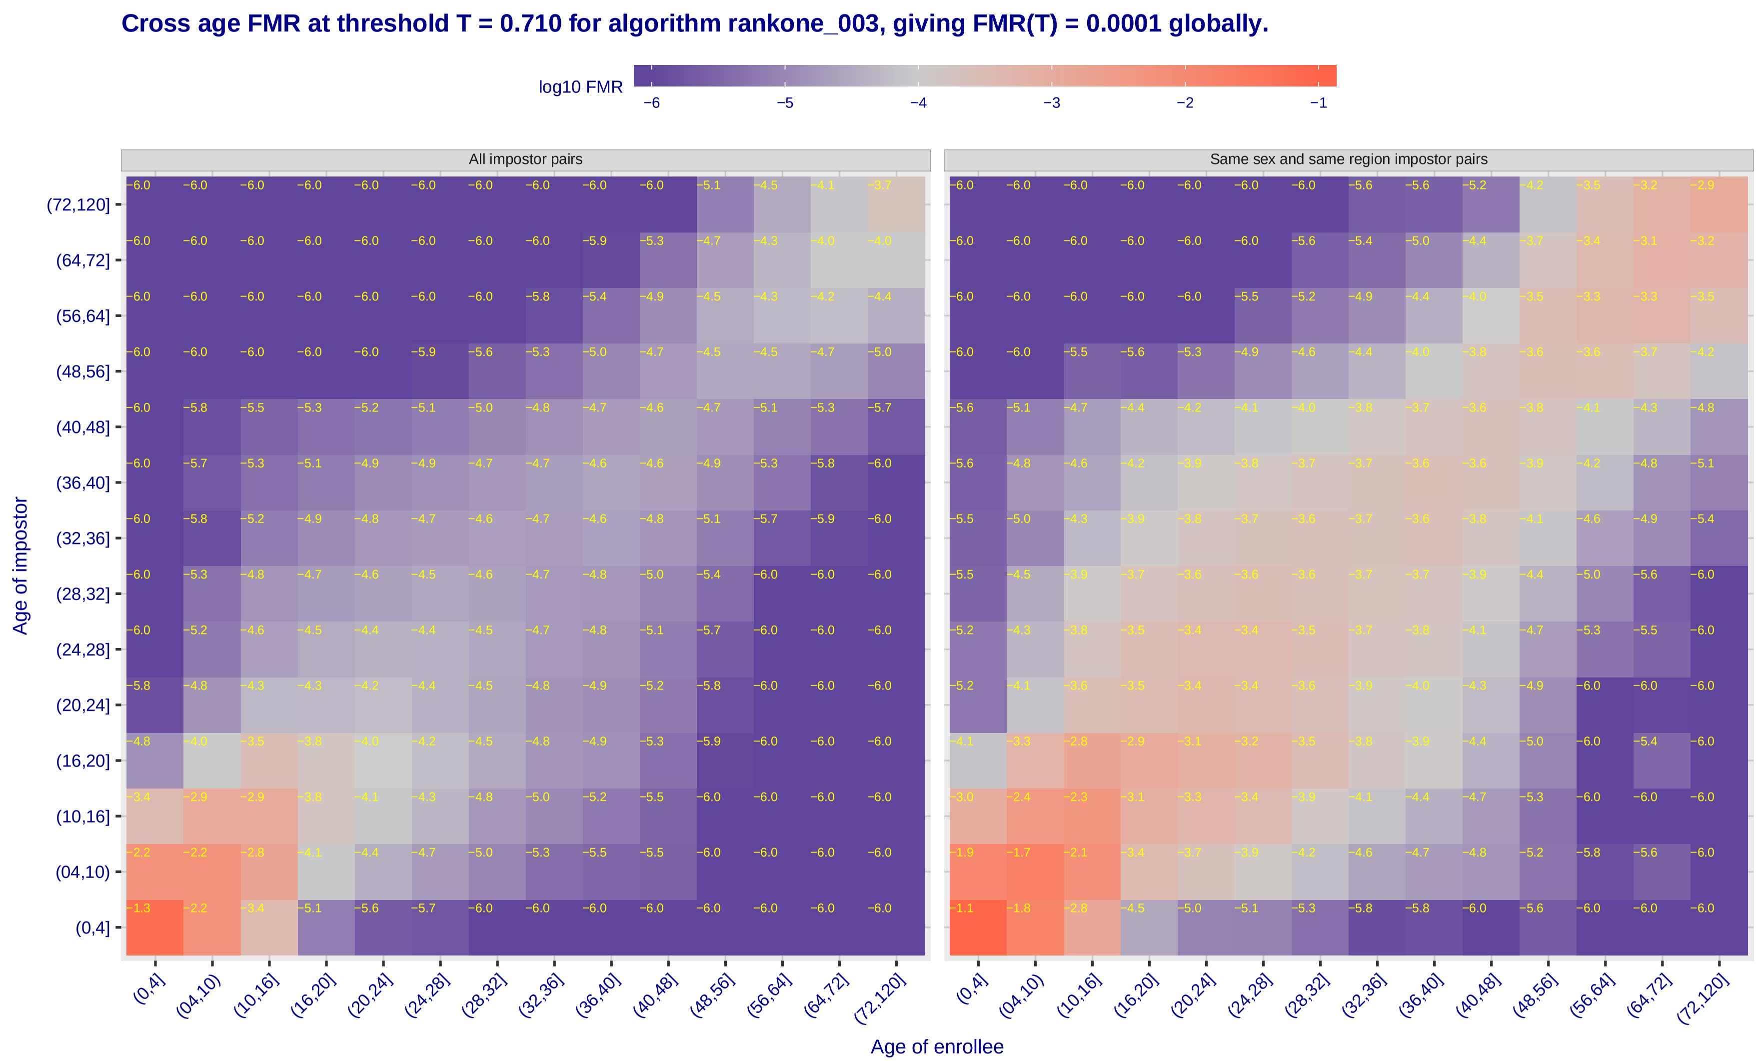

Figure 169: For algorithm rankone-003 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

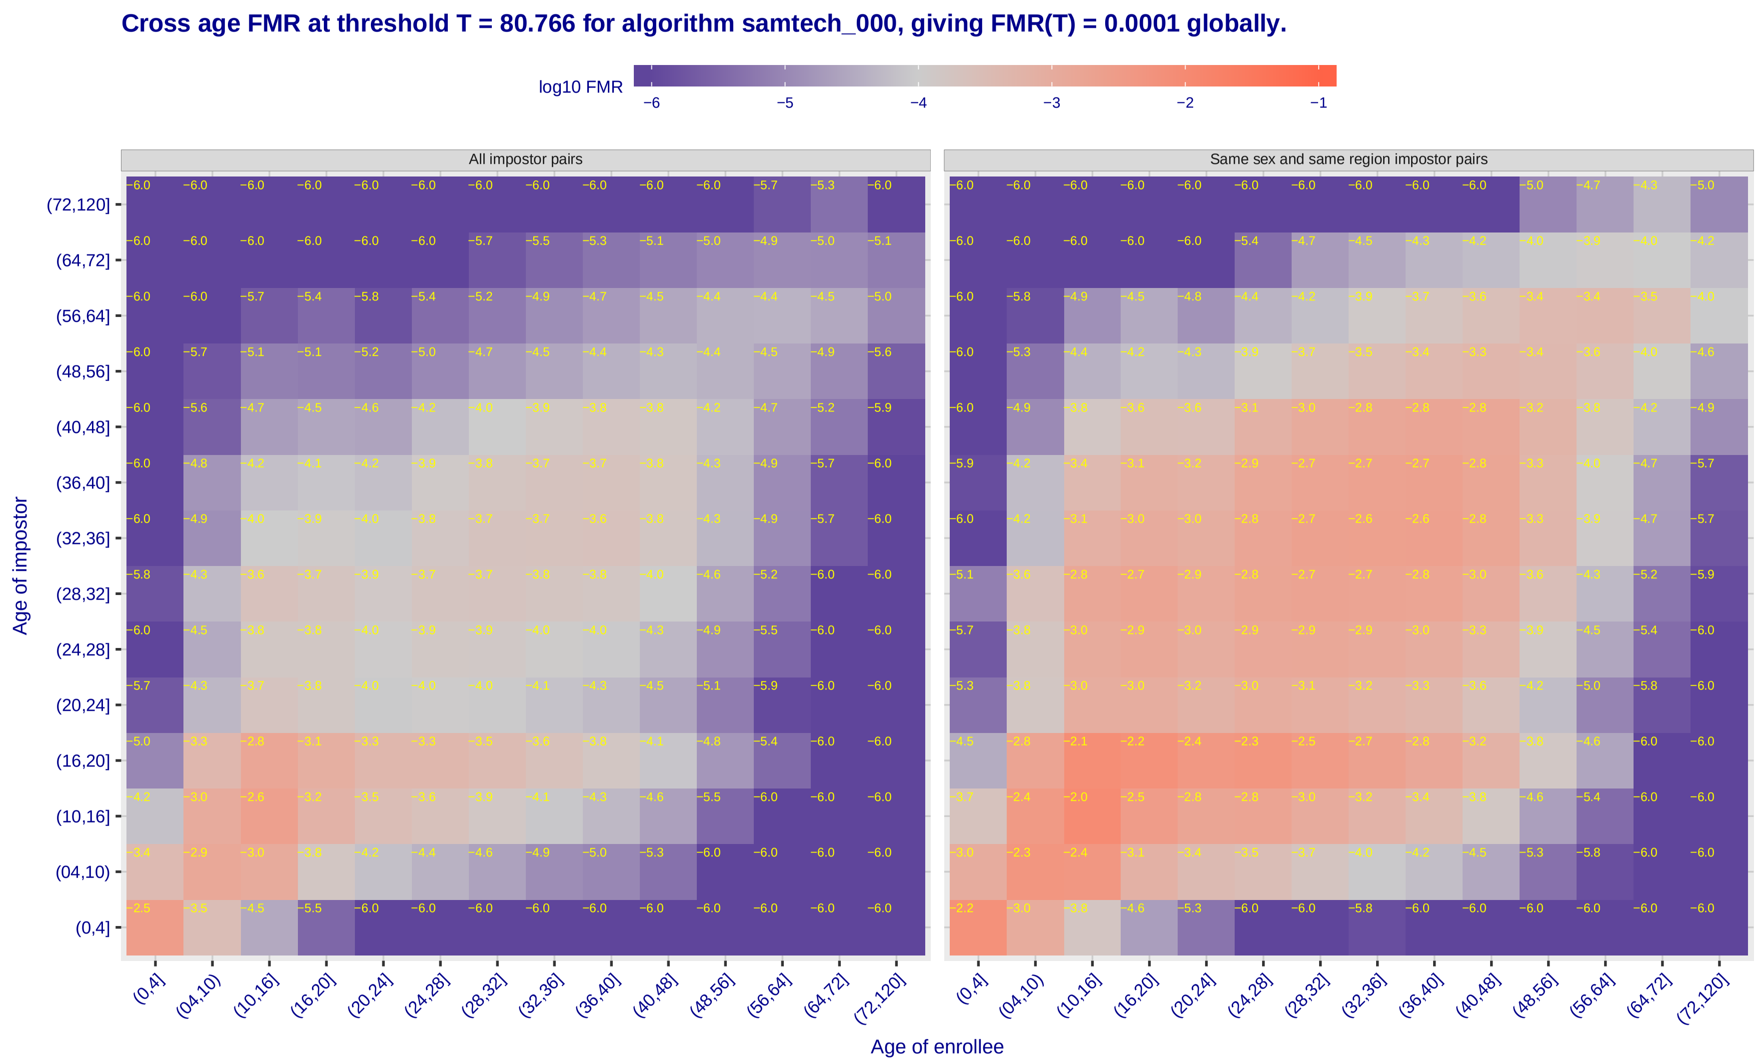

Figure 170: For algorithm samtech-000 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all $O(10^{10})$ impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

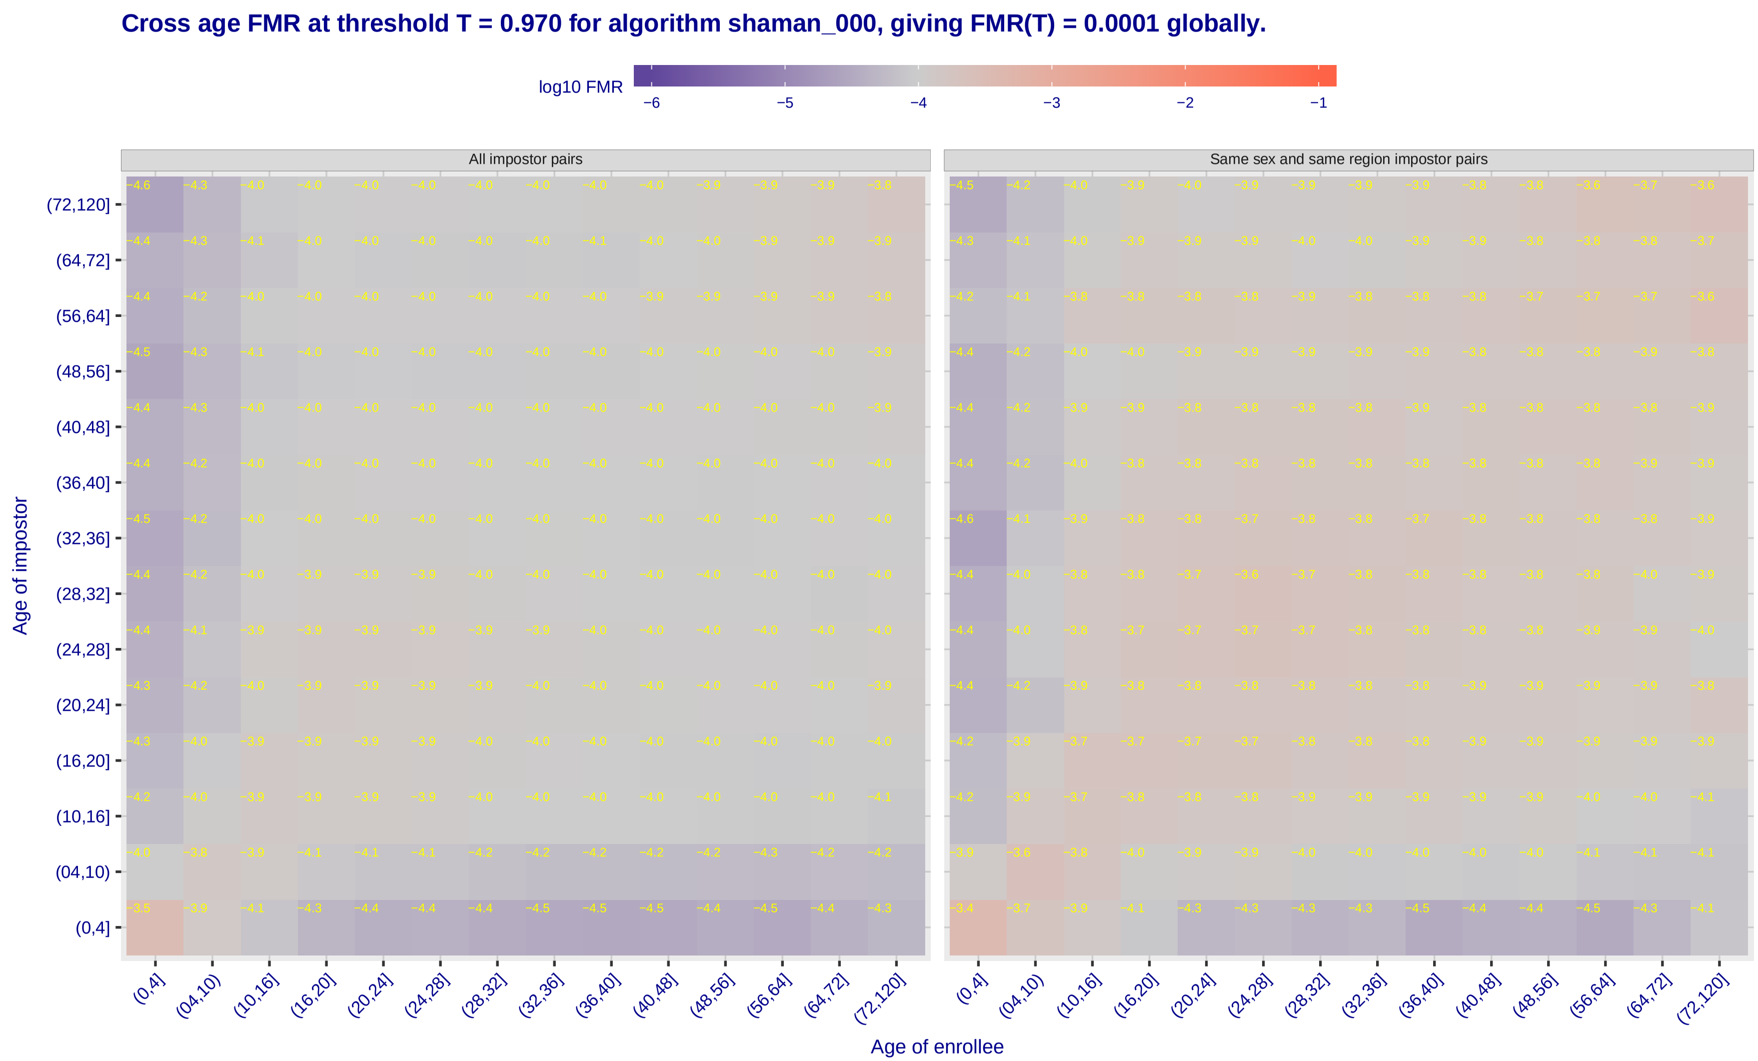

Figure 171: For algorithm shaman-000 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

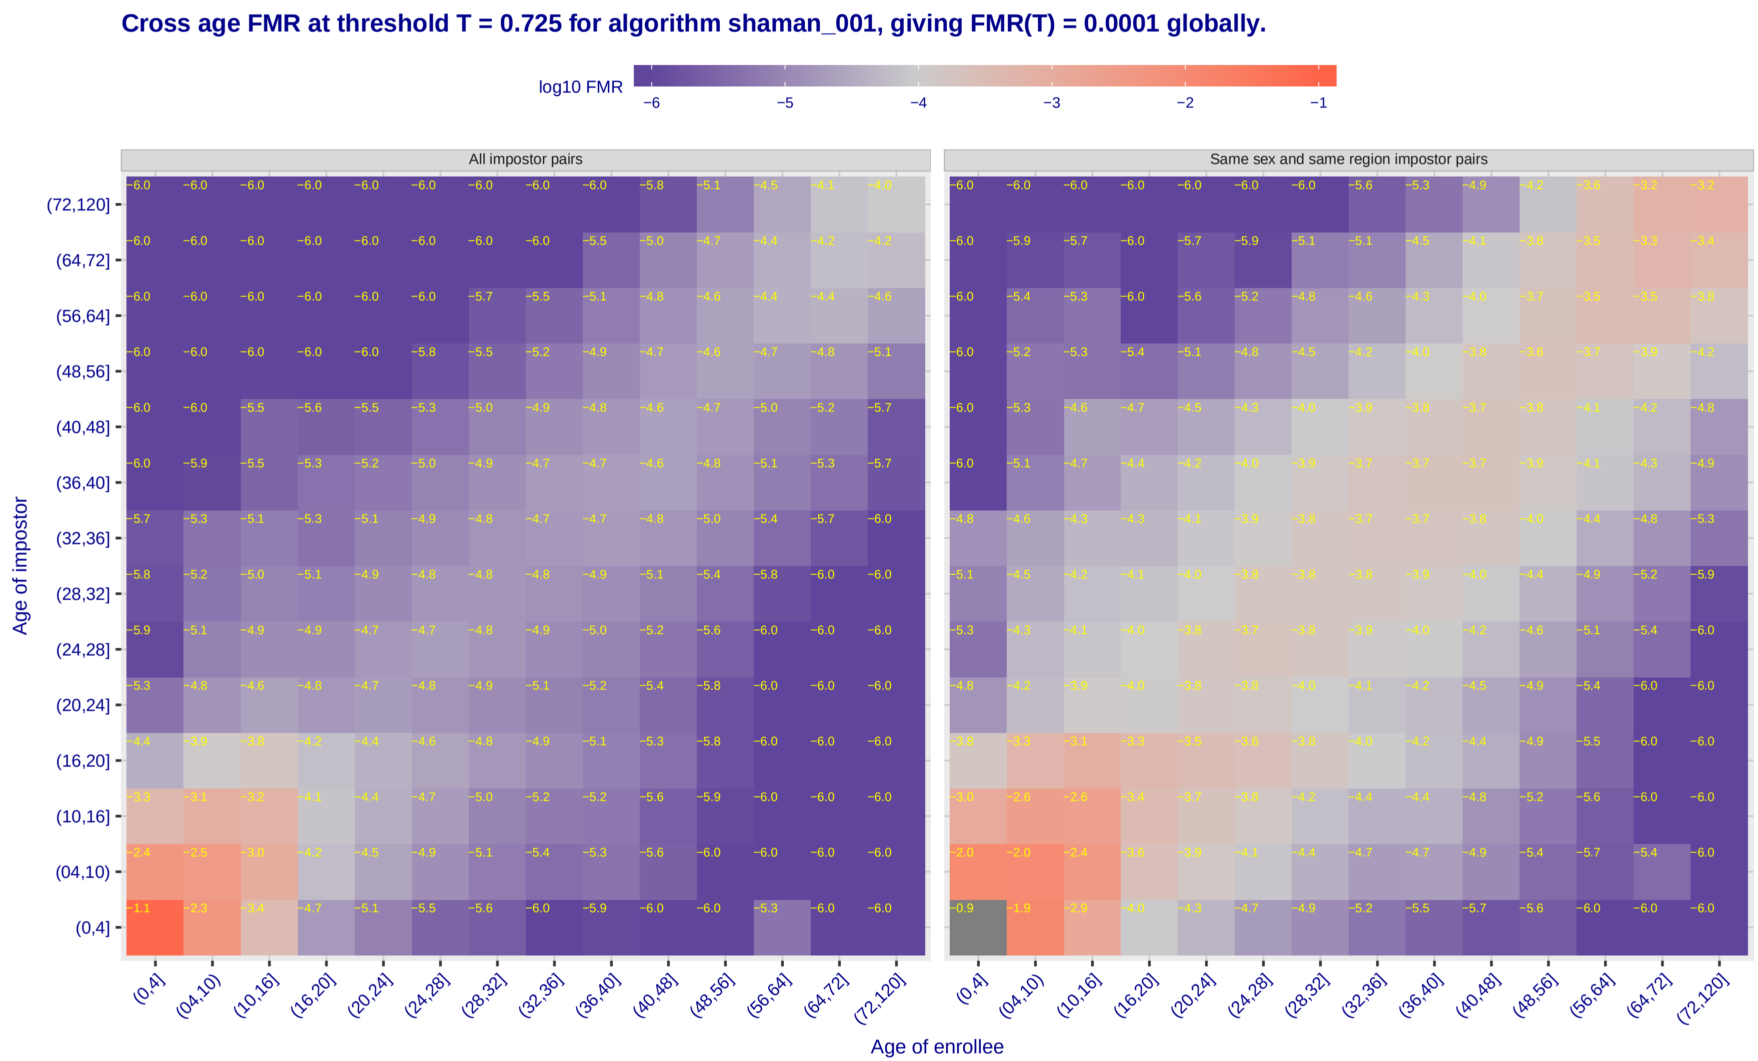

Figure 172: For algorithm shaman-001 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all $O(10^{10})$ impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

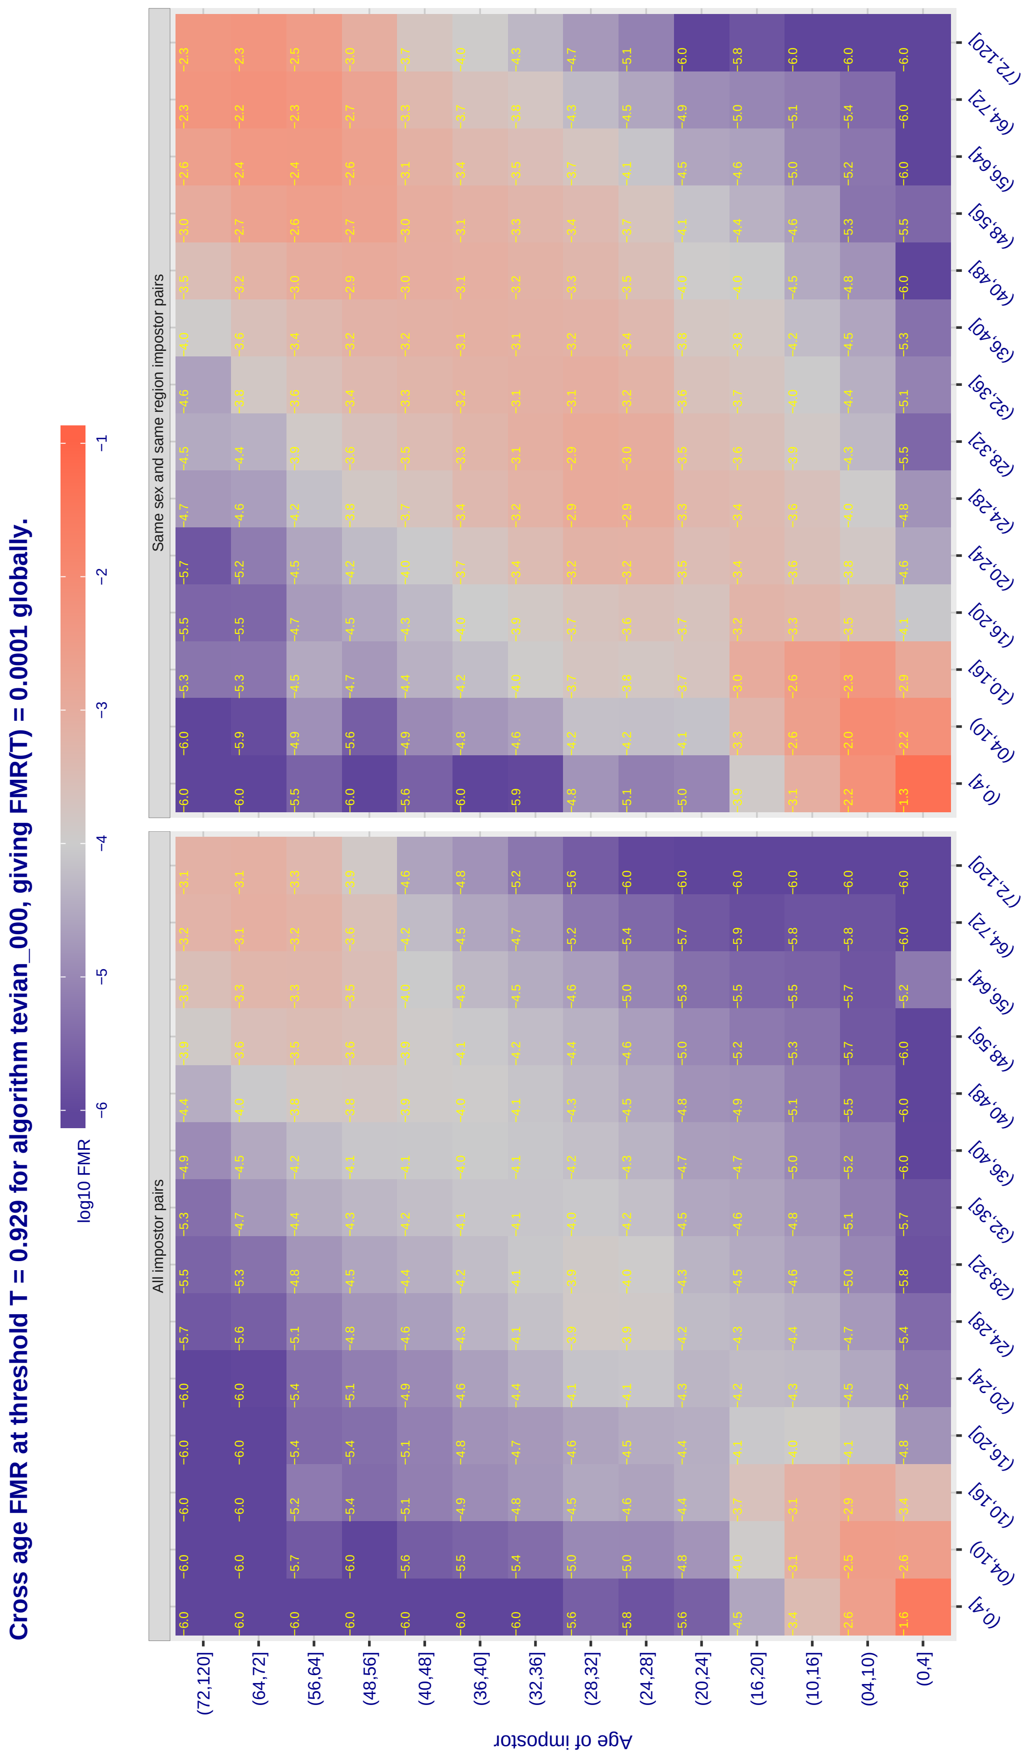

Figure 173: For algorithm tevian-000 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

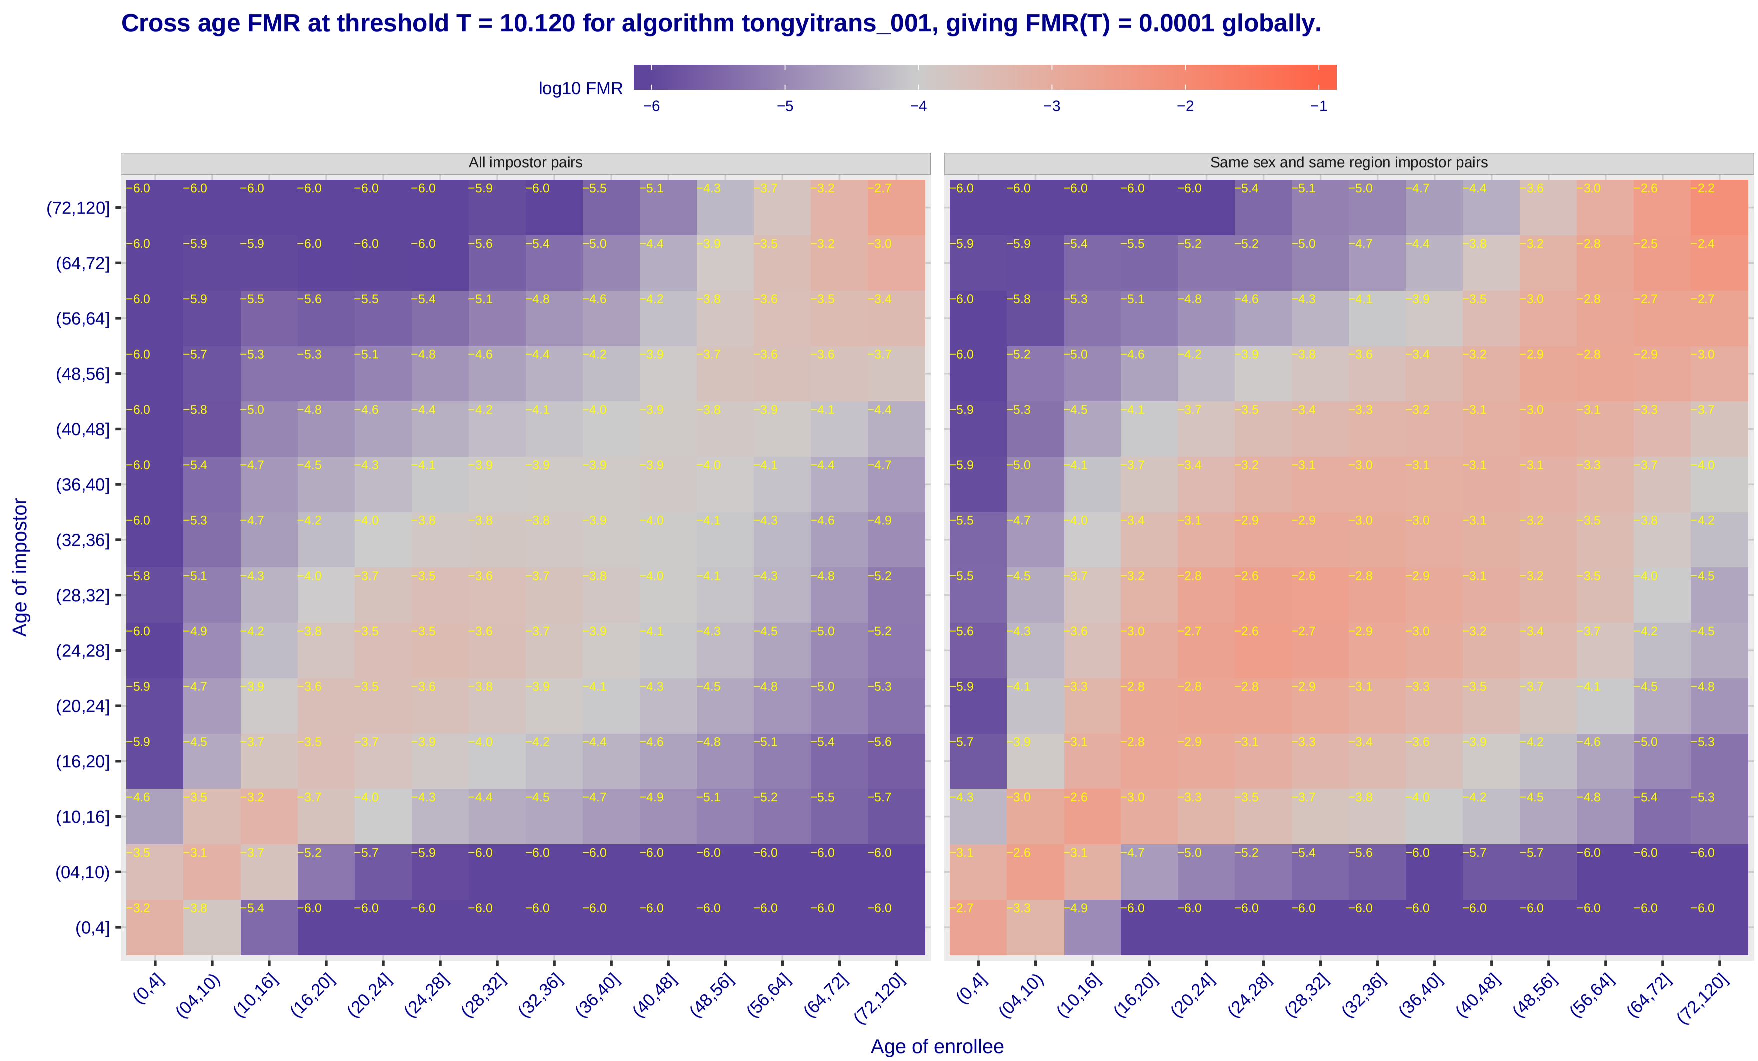

Figure 174: For algorithm tongyitrans-001 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

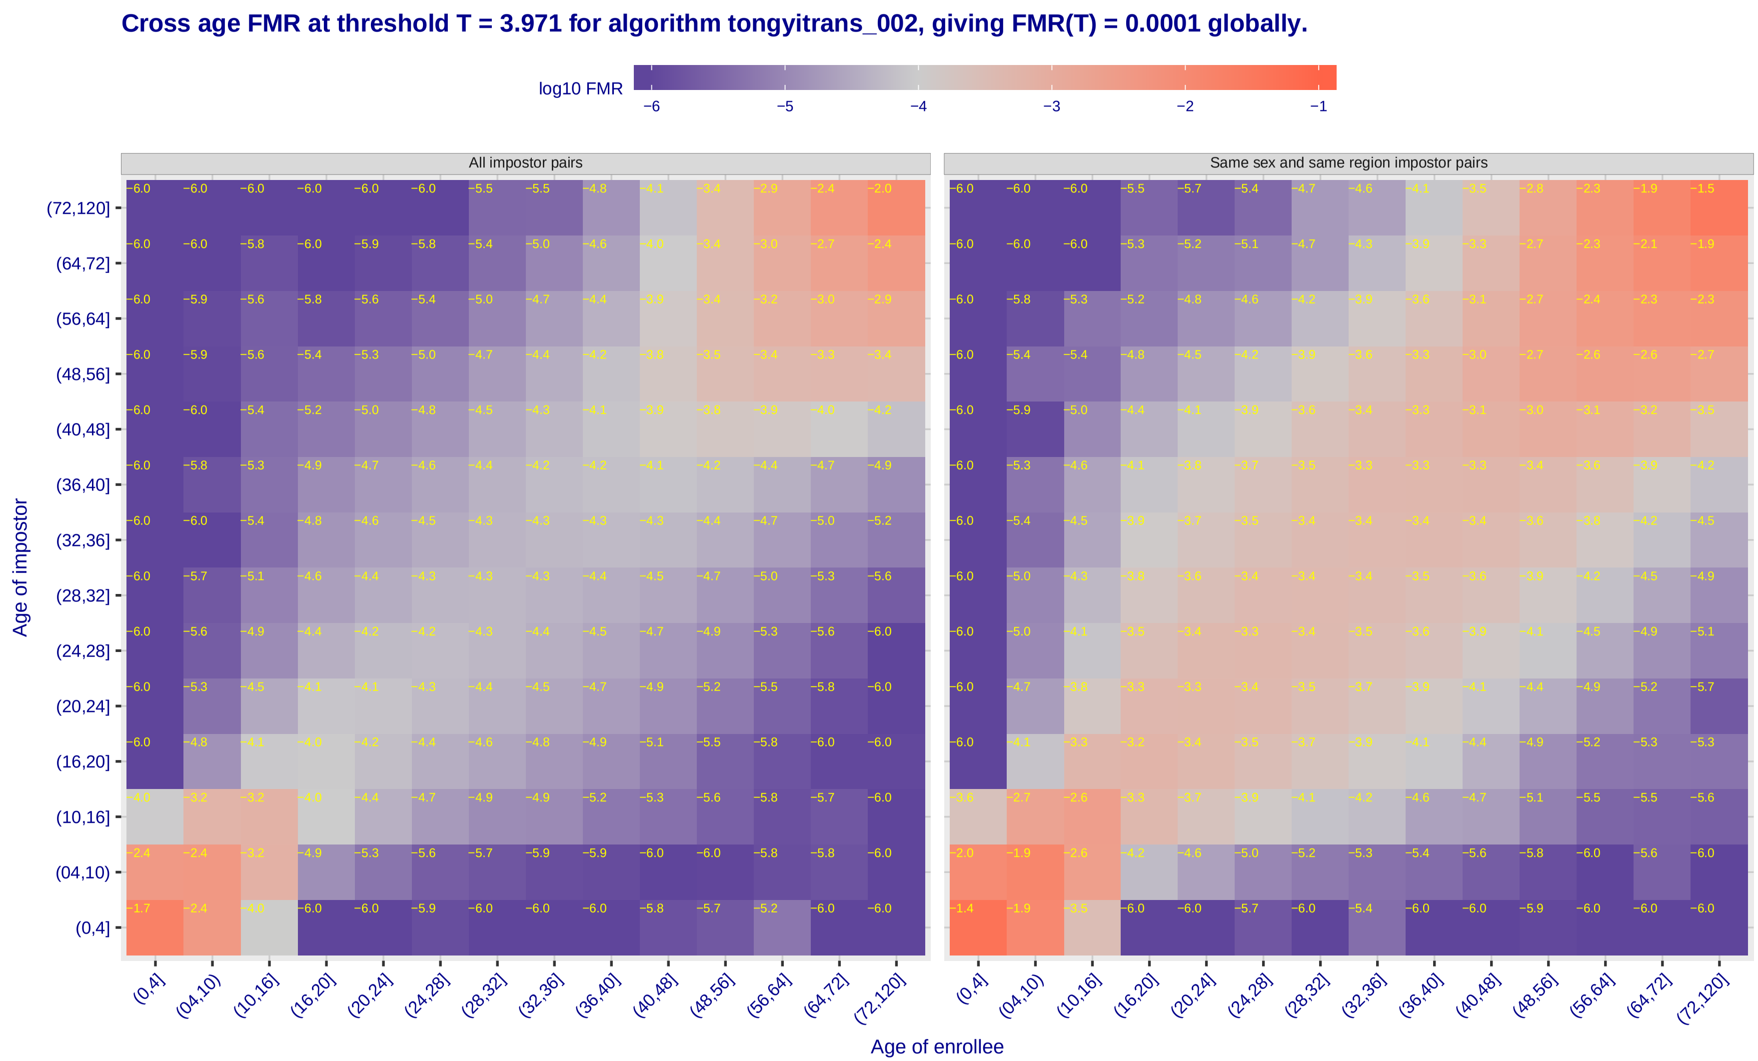

Figure 175: For algorithm tongyitrans-002 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

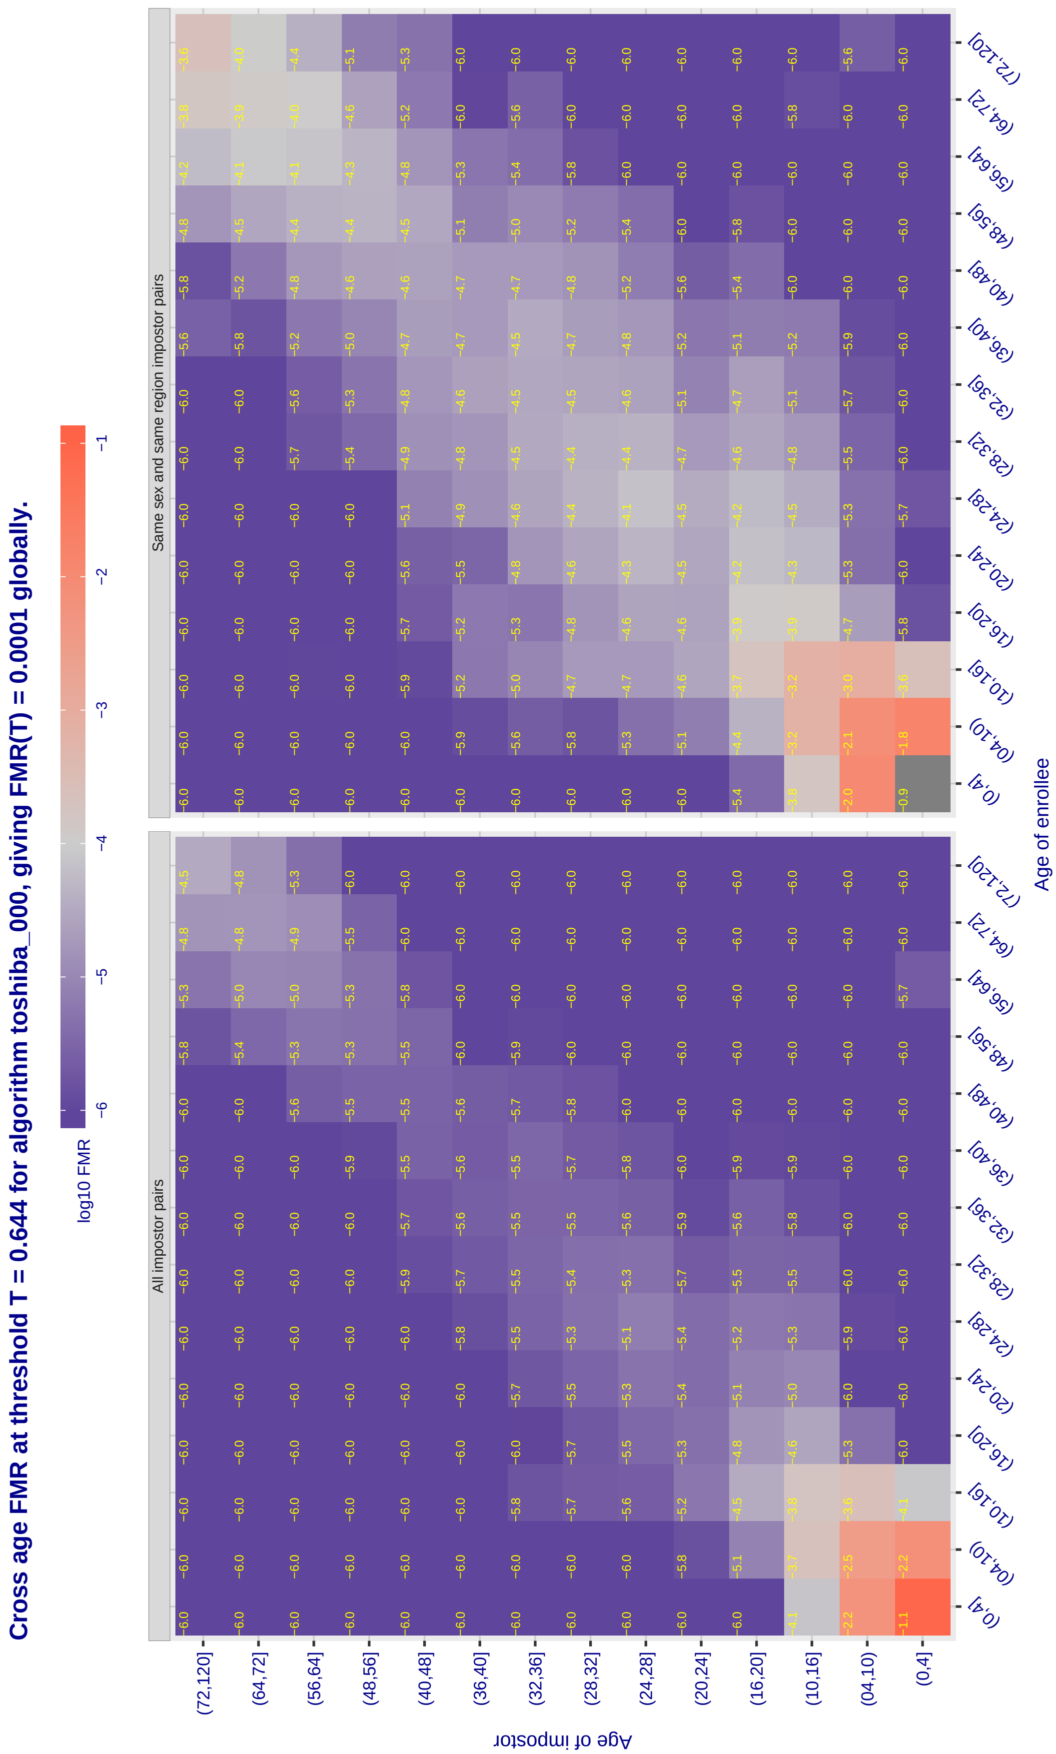

Figure 176: For algorithm toshiba-000 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

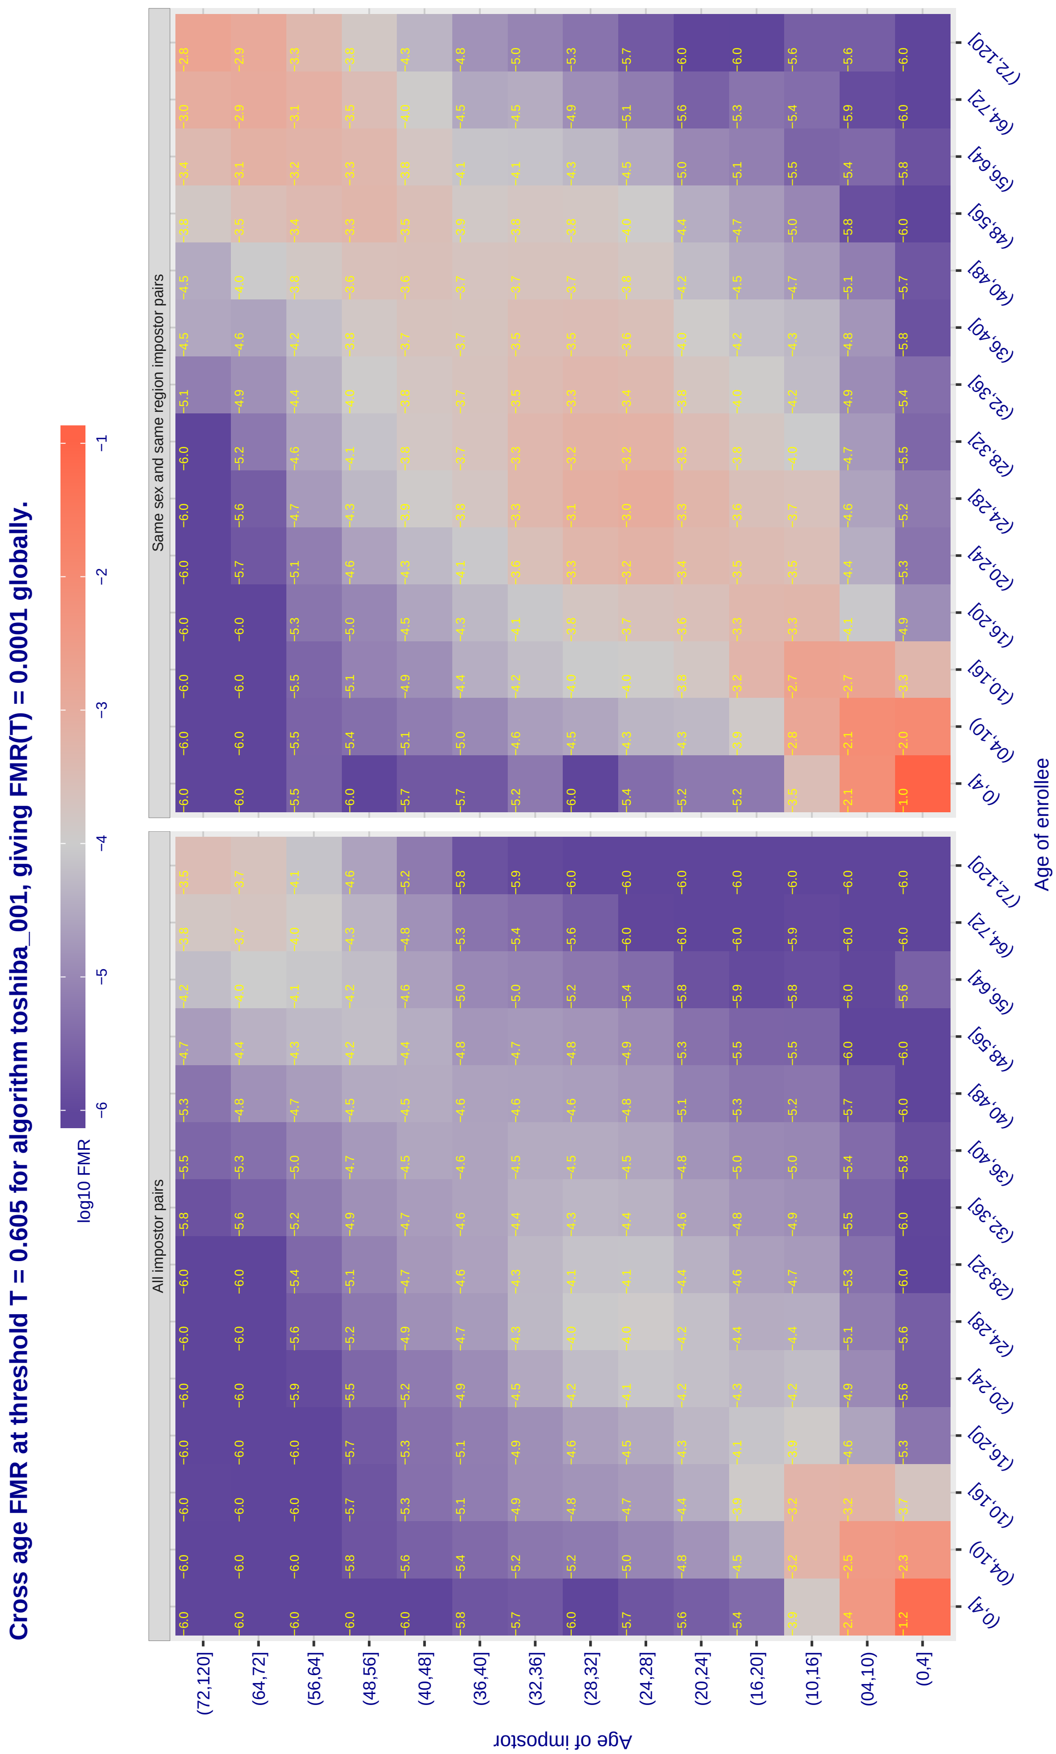

Figure 177: For algorithm toshiba-001 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

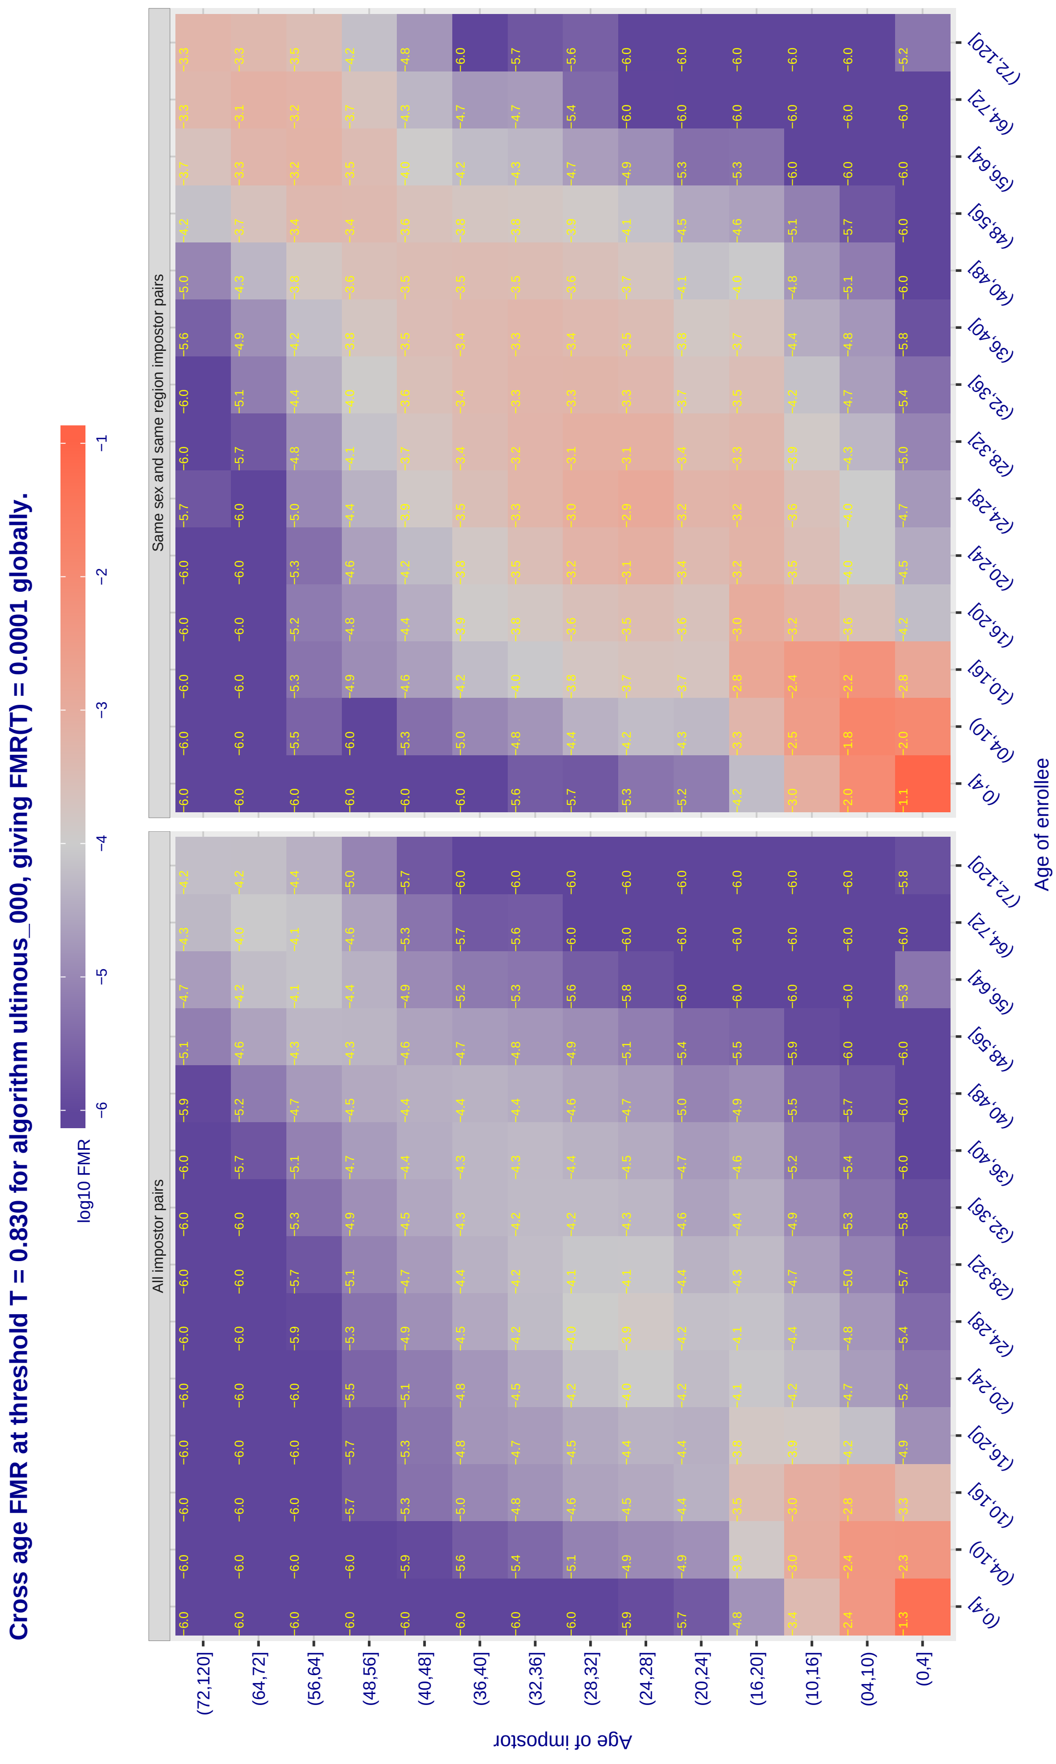

Figure 178: For algorithm ultinous-000 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

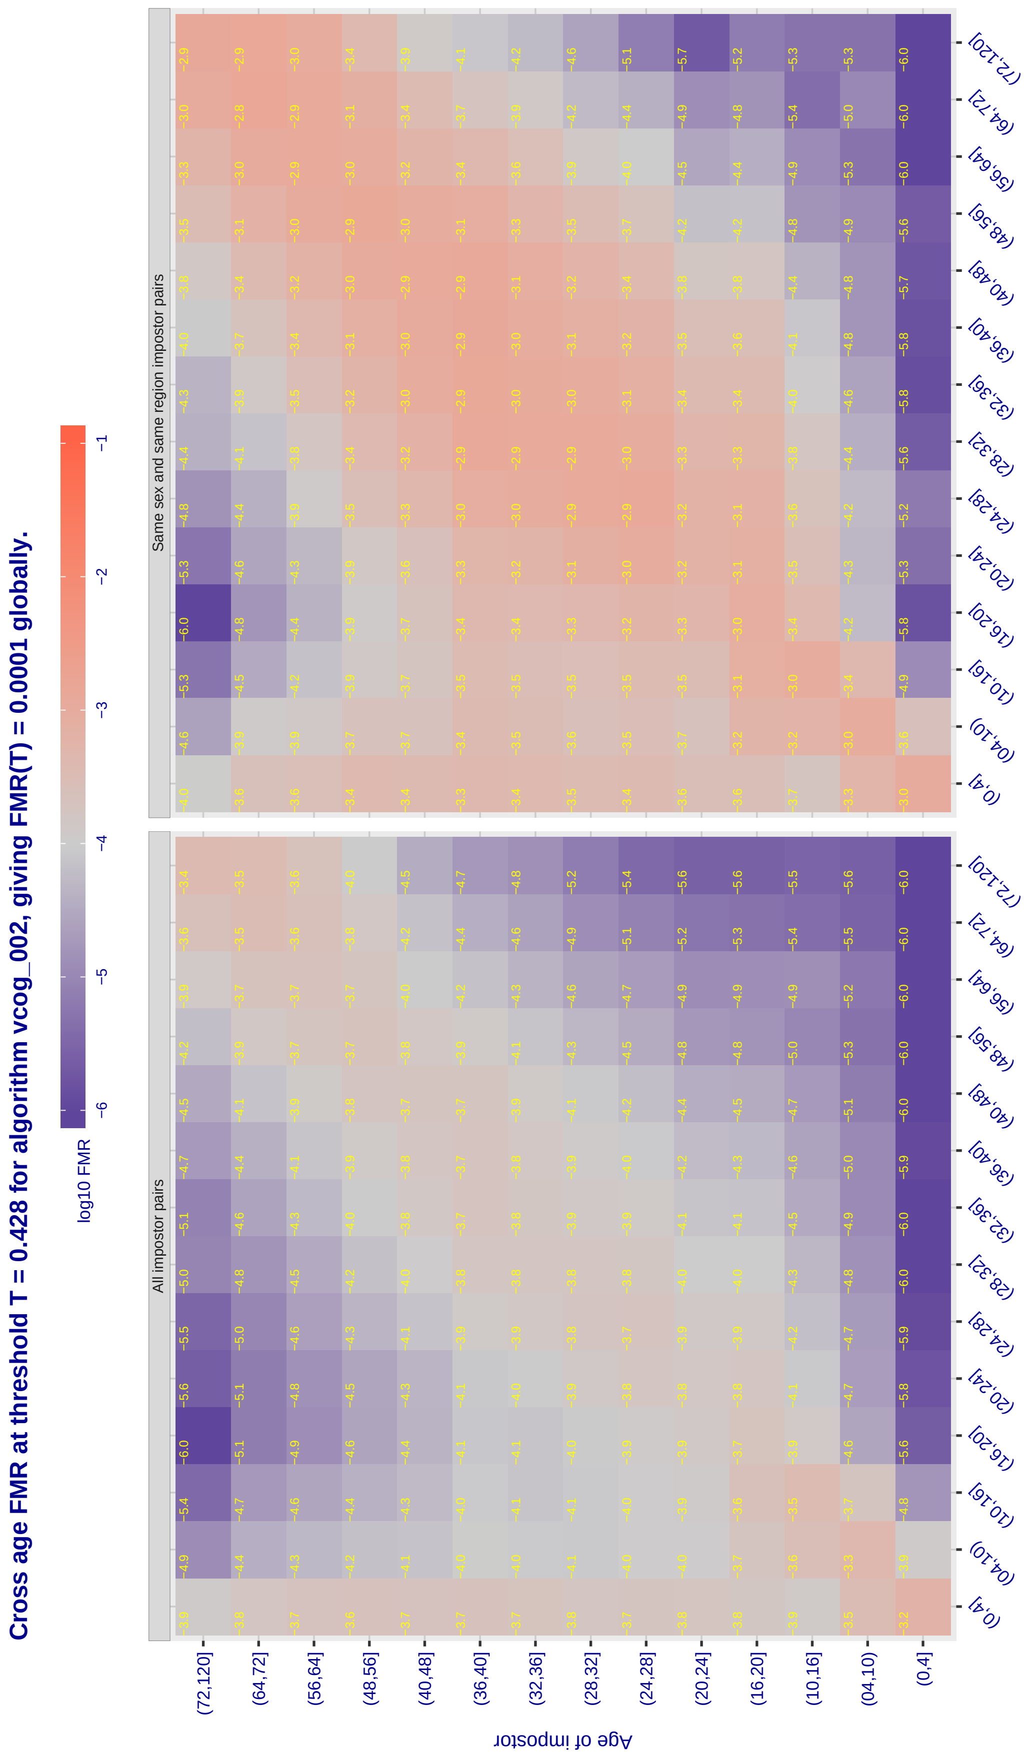

Figure 179: *For algorithm vcog-002 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.*

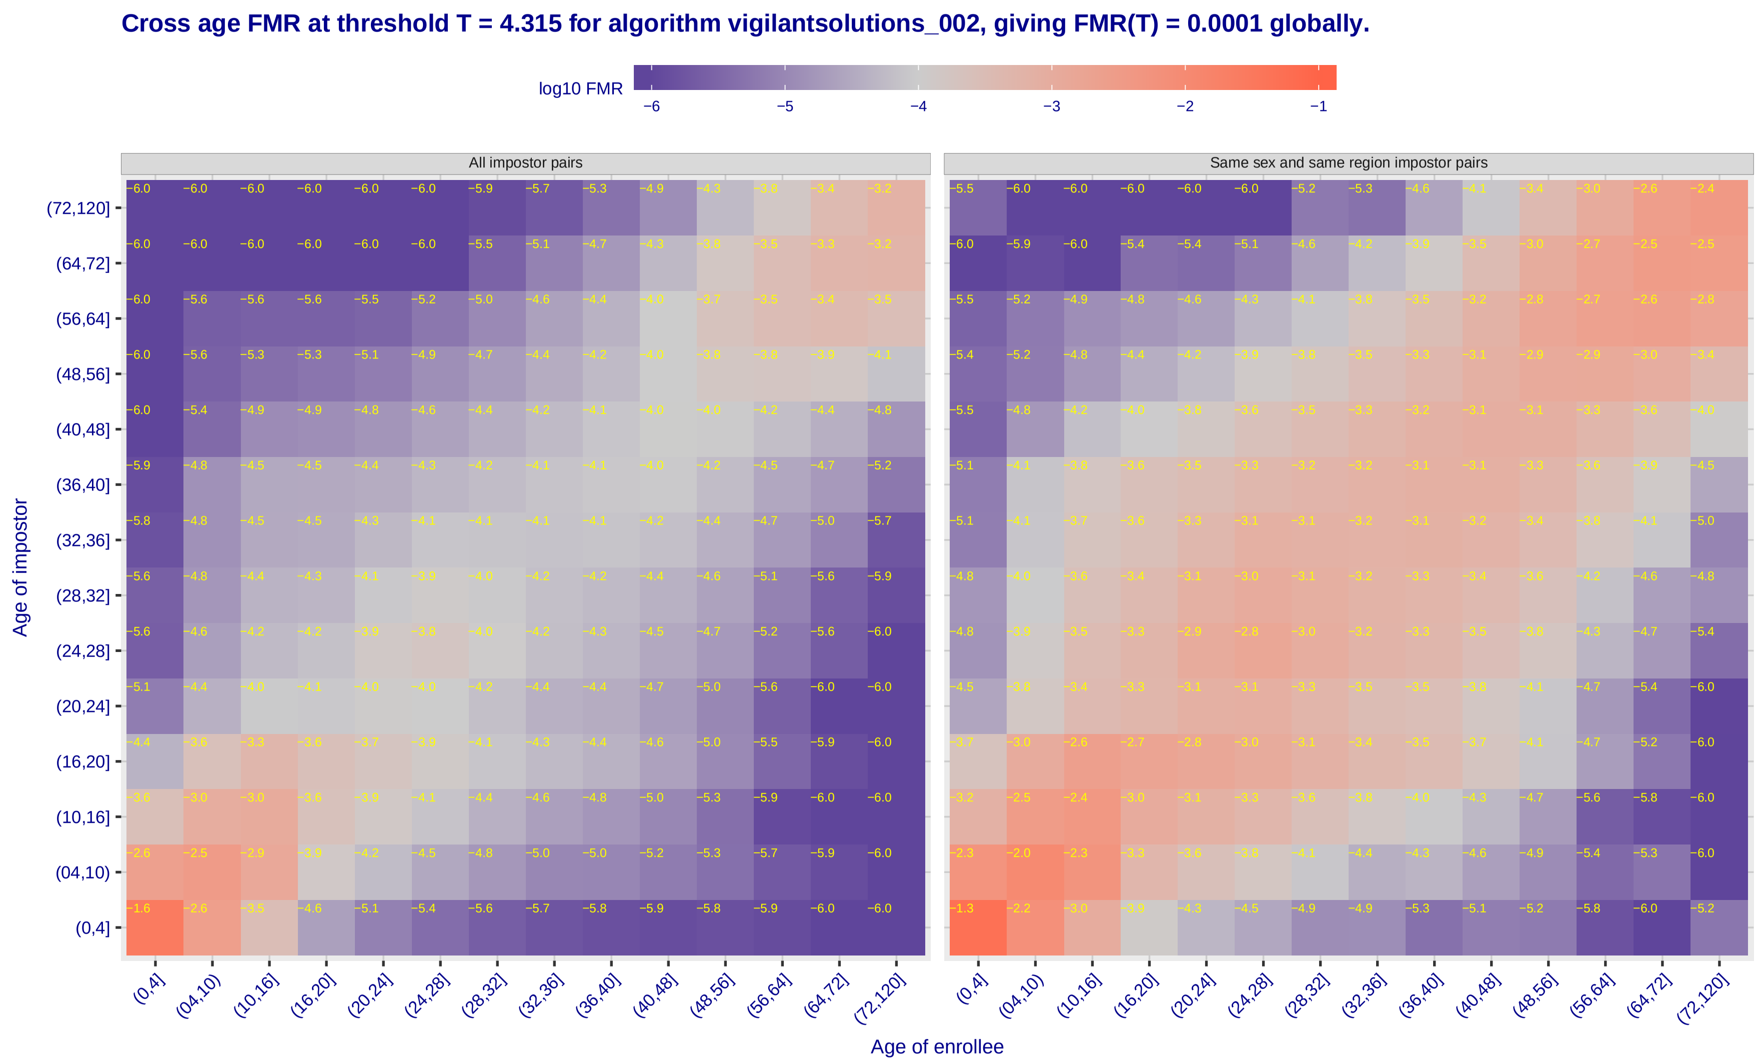

Figure 180: For algorithm vigilantsolutions-002 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

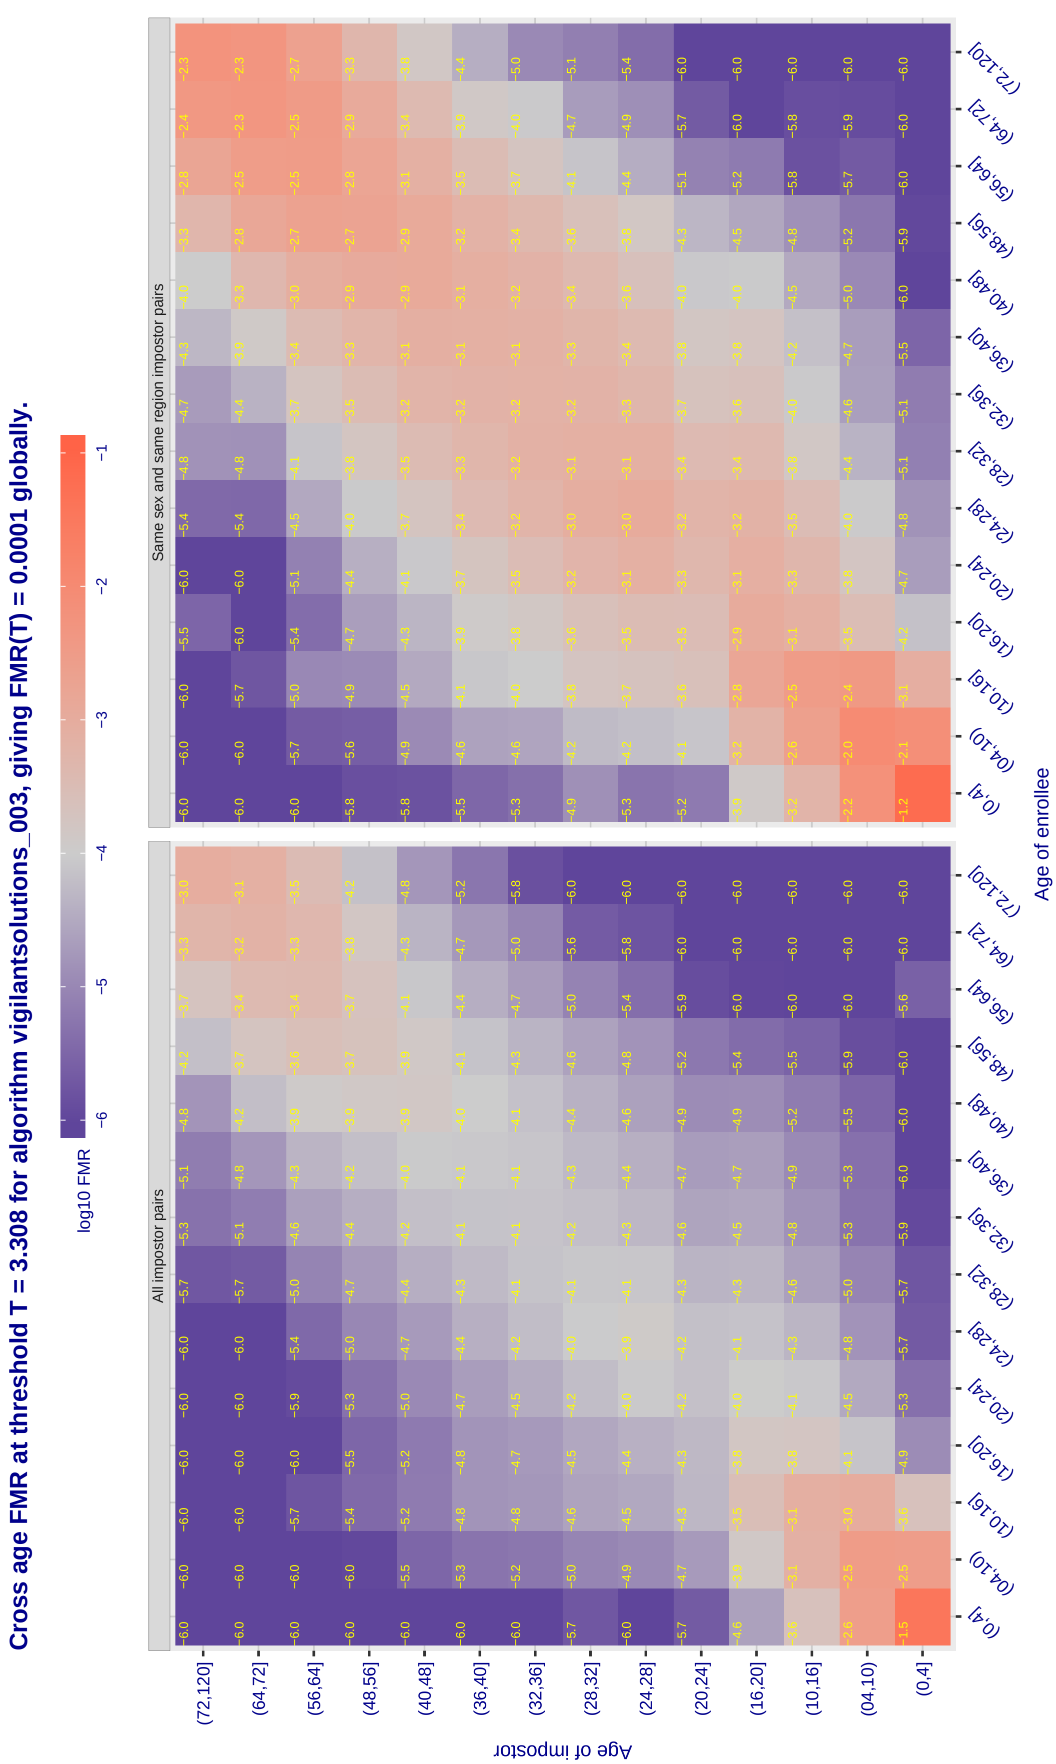

Figure 181: For algorithm vigilantsolutions-003 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all $O(10^{10})$ impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

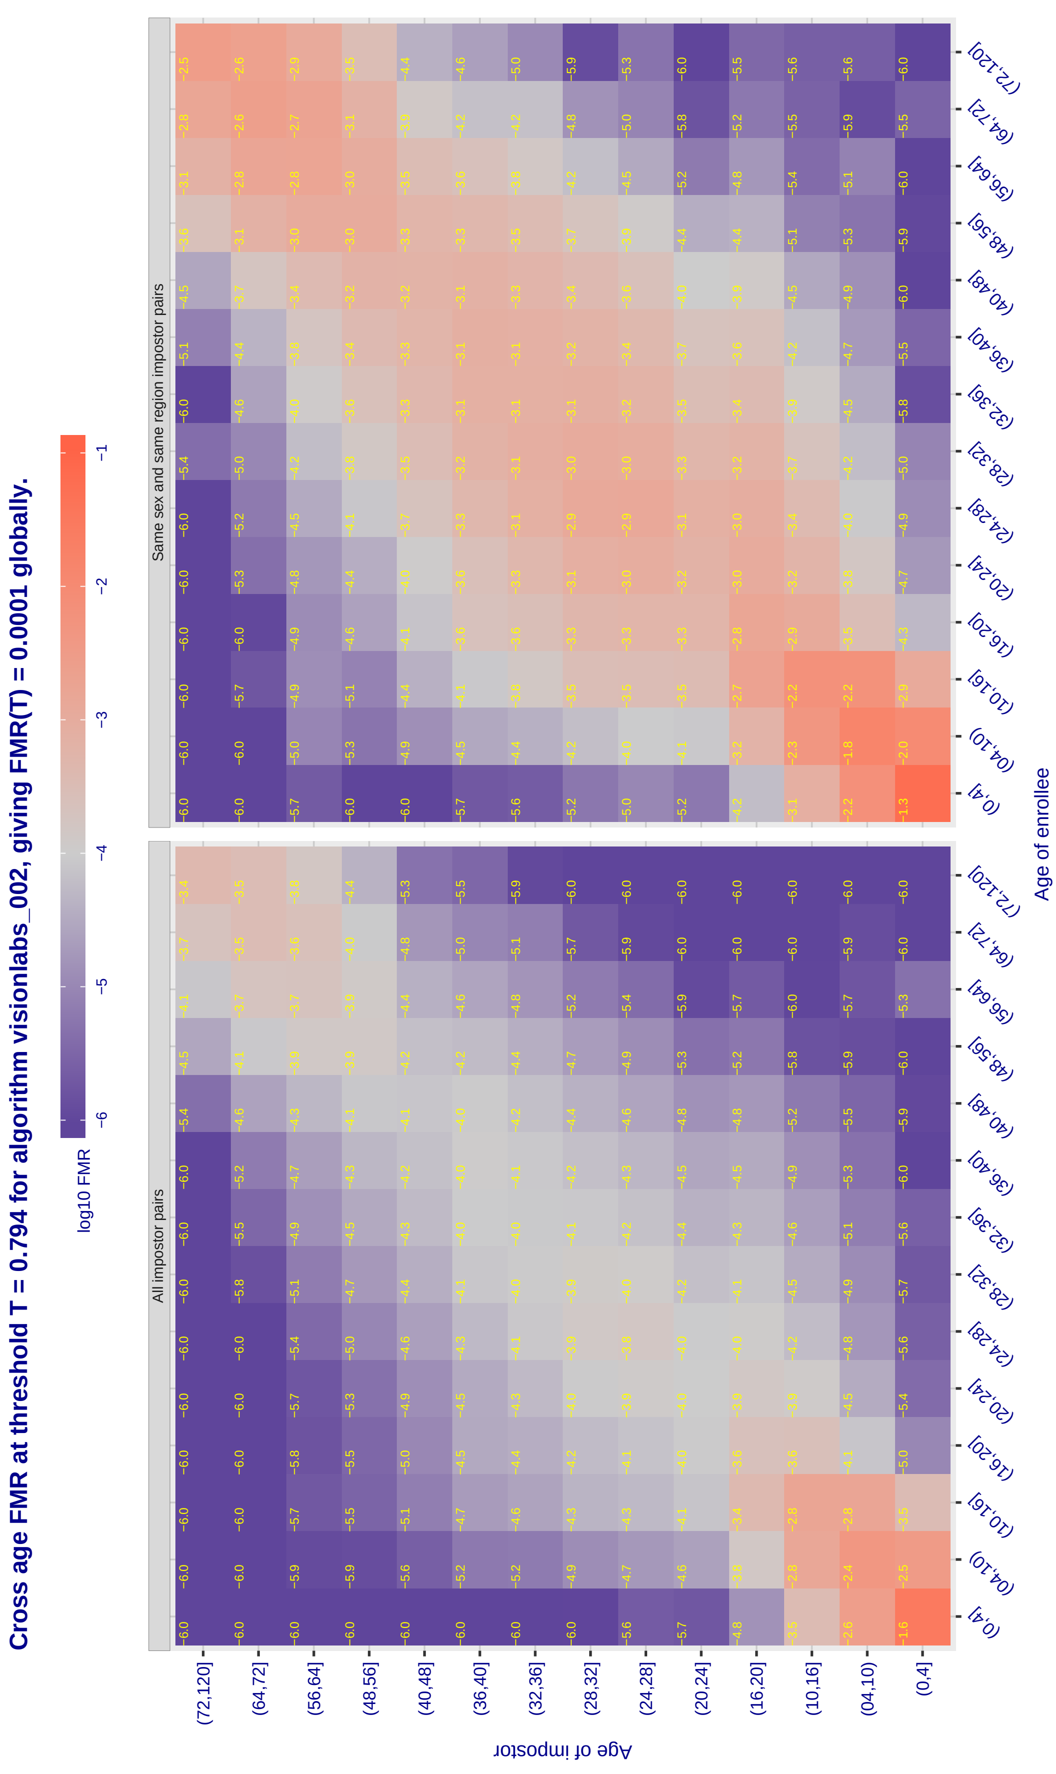

Figure 182: For algorithm visionlabs-002 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

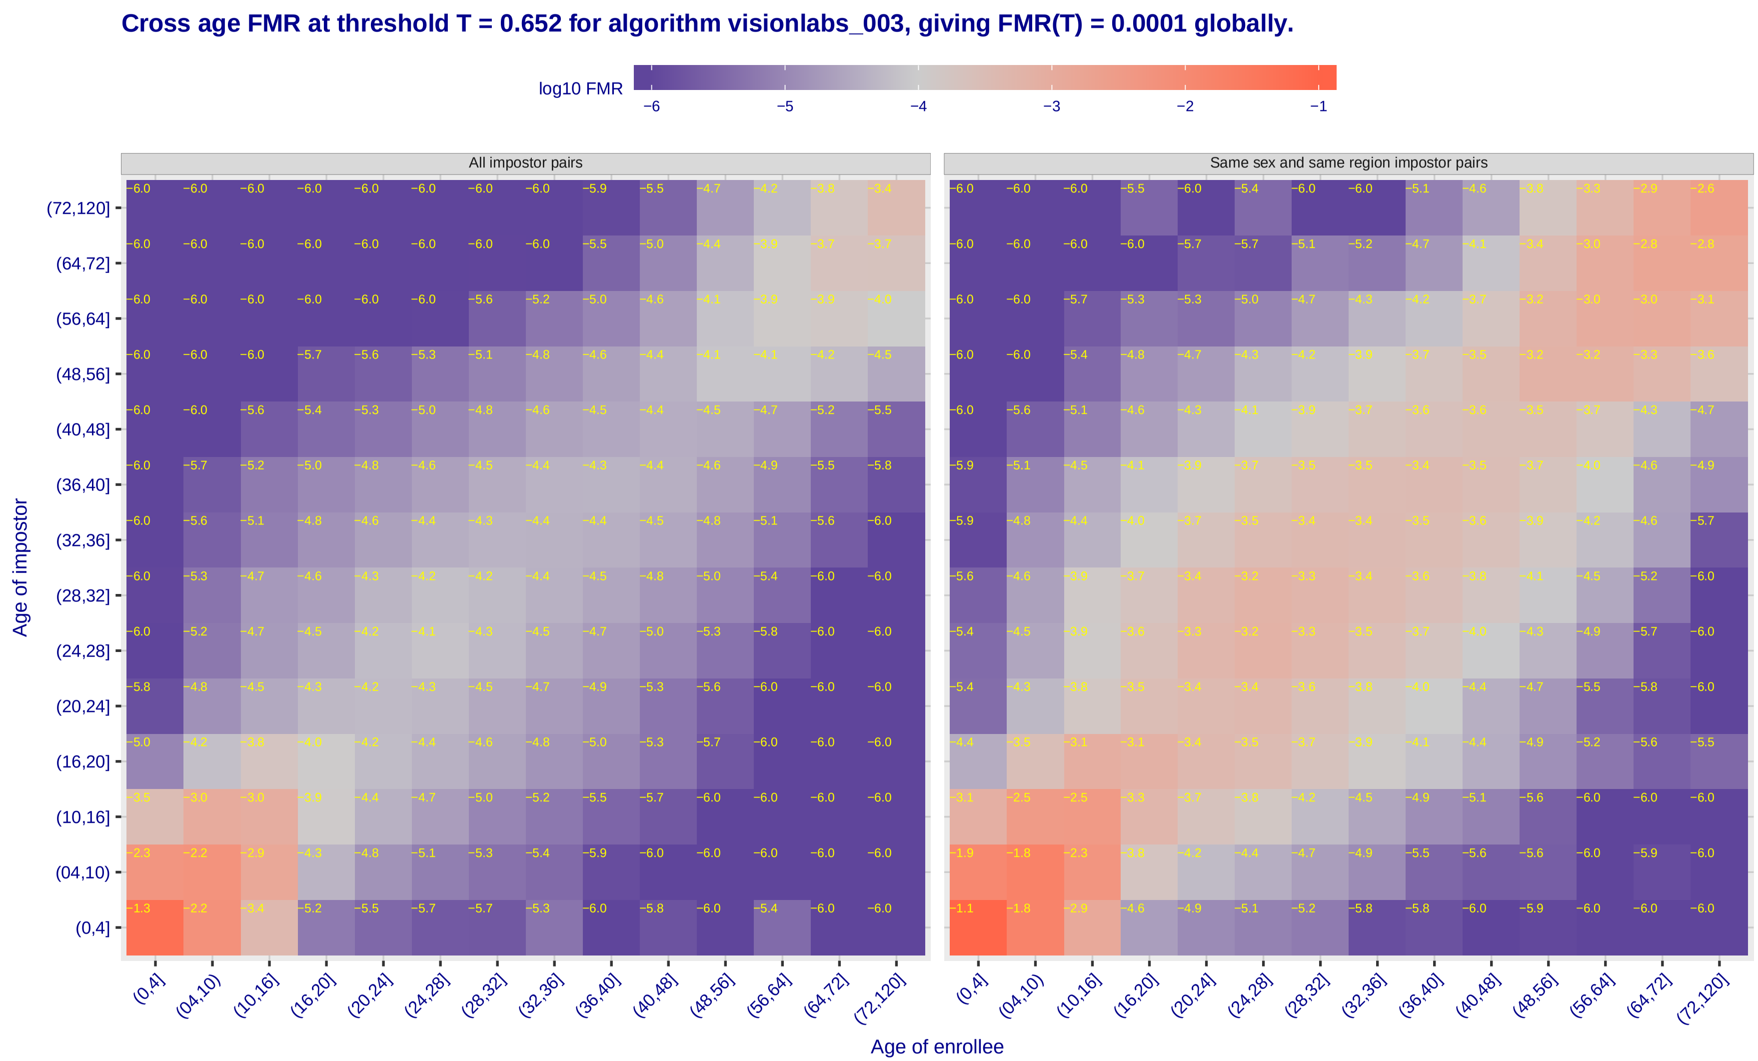

Figure 183: For algorithm visionlabs-003 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all $O(10^{10})$ impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

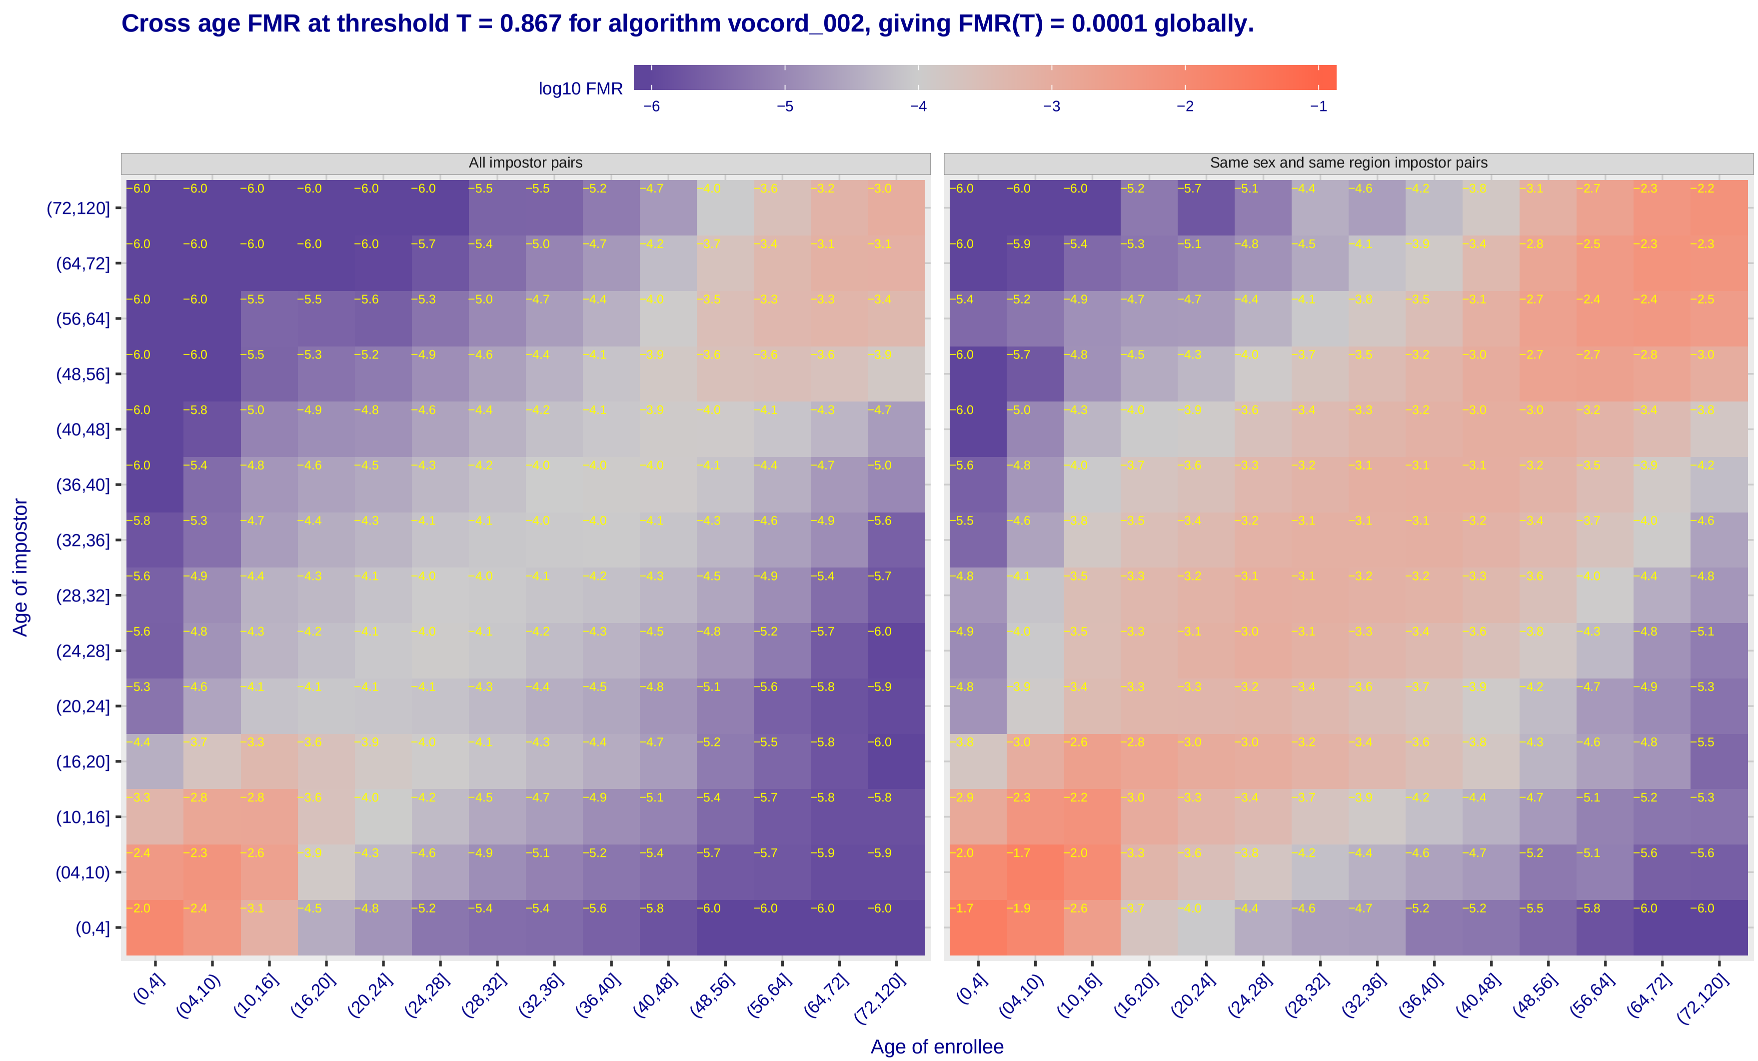

Figure 184: For algorithm vocord-002 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all $O(10^{10})$ impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

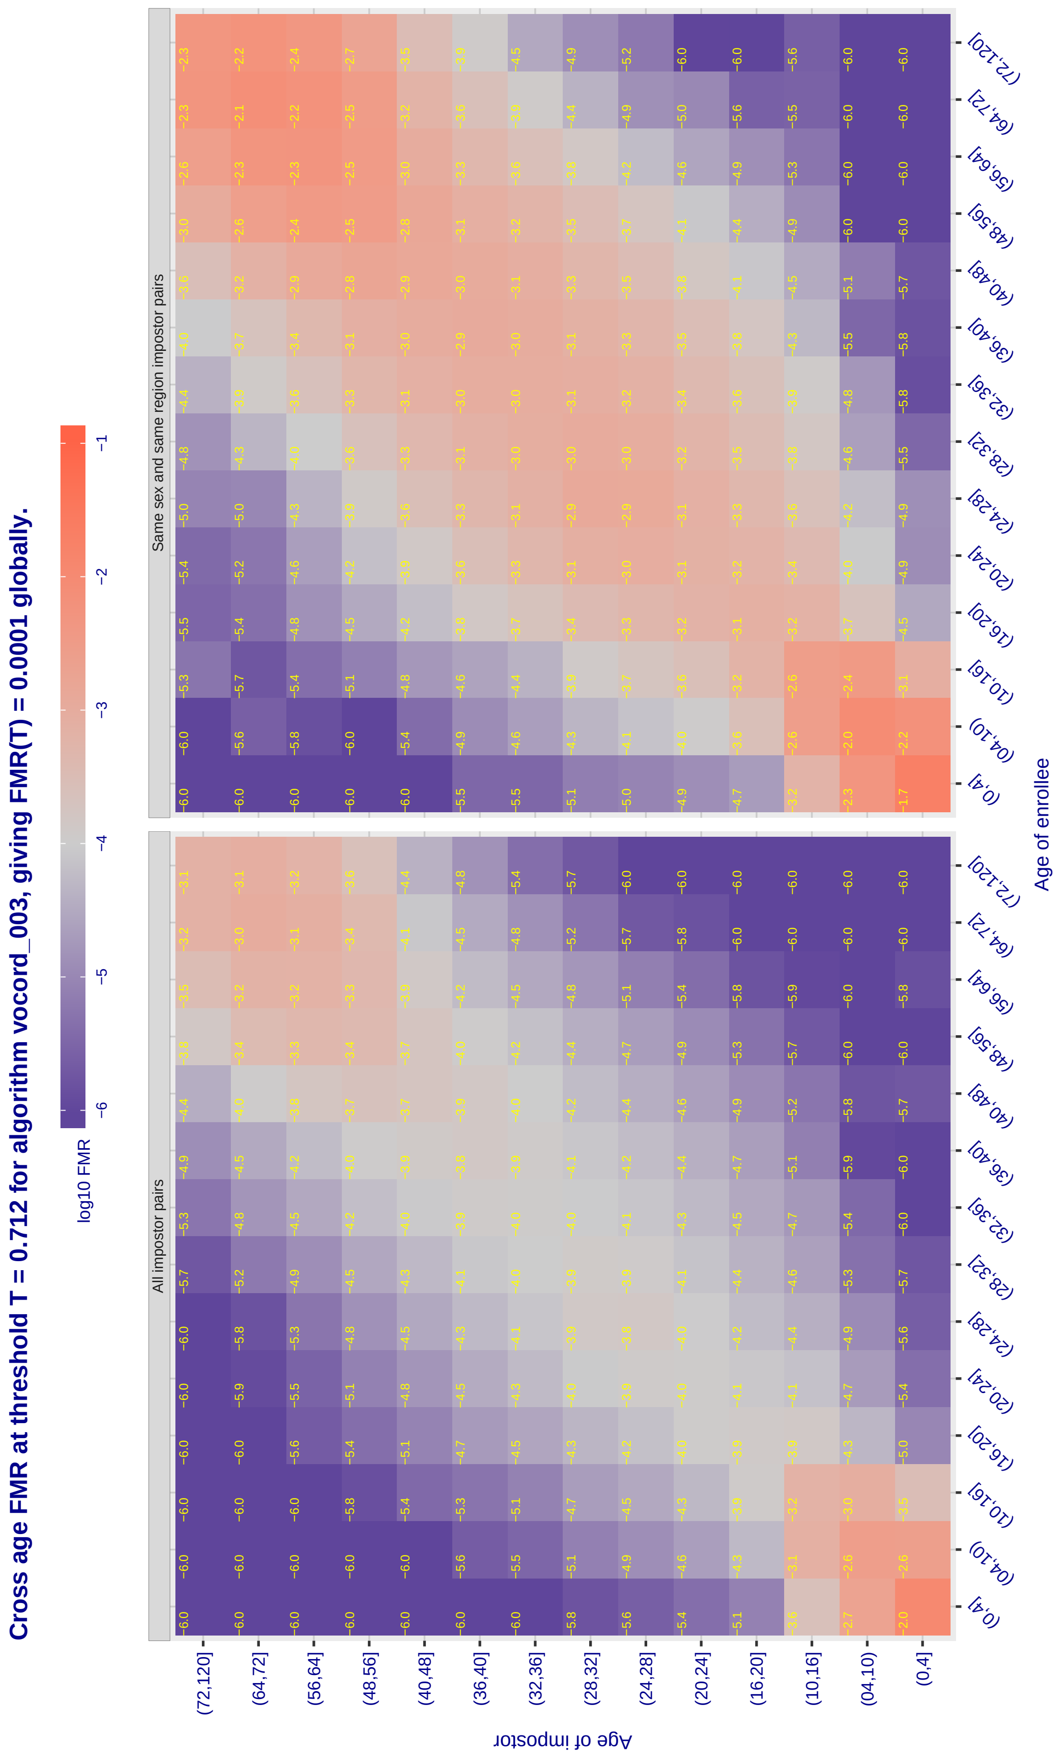

Figure 185: For algorithm vocord-003 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

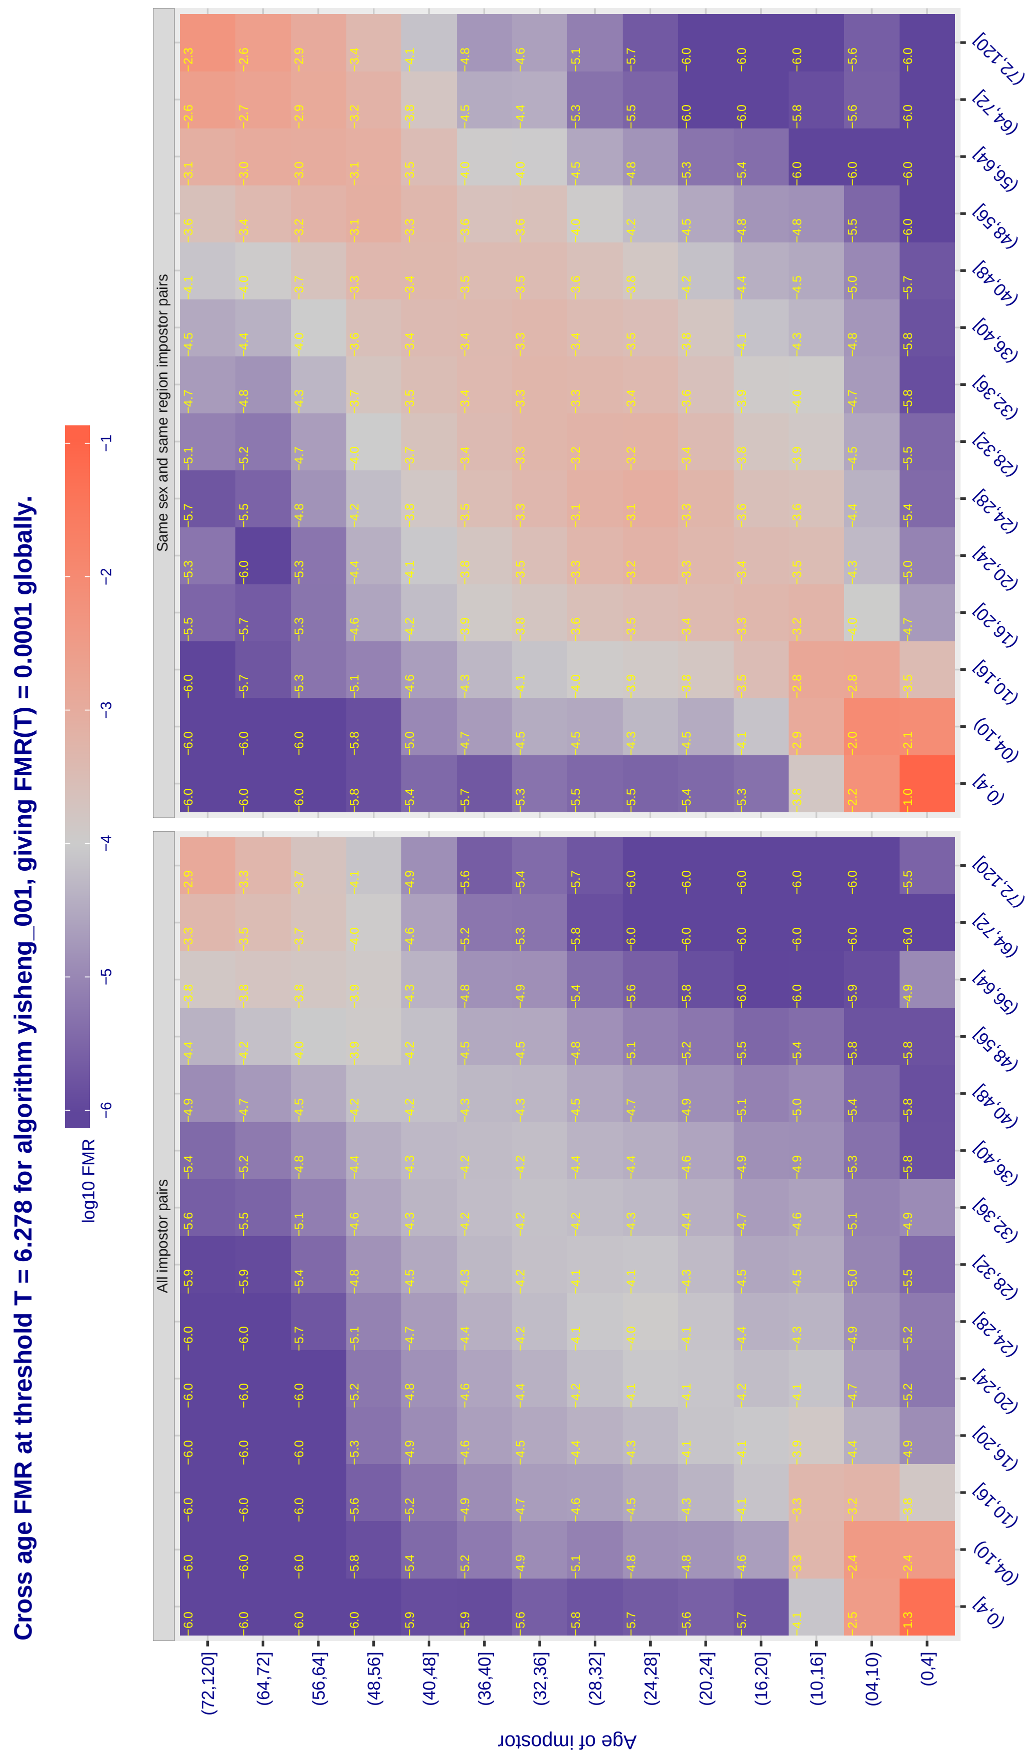

Figure 186: For algorithm yisheng-001 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

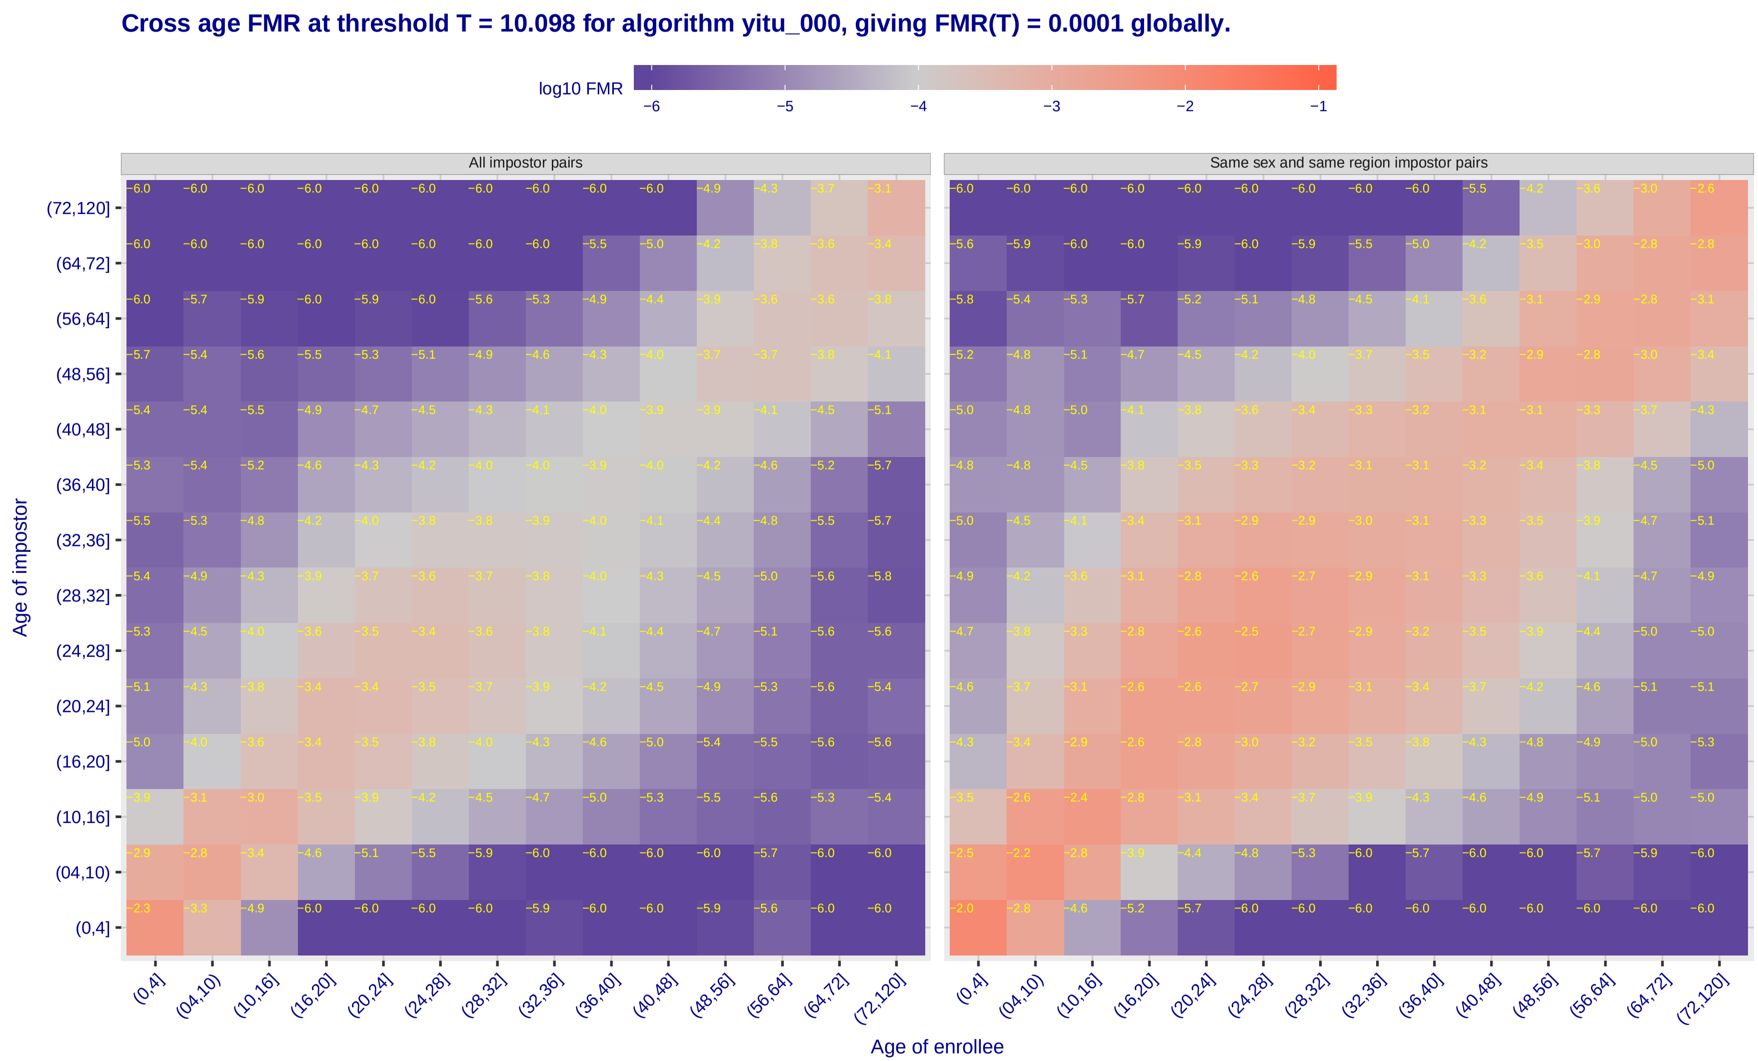

Figure 187: For algorithm yitu-000 operating on visa images, the heatmap shows false match observed over impostor comparisons of faces from different individuals who have the given age pair. False matches are counted against a recognition threshold fixed globally to give FMR = 0.001 over all O(10^{10}) impostor comparisons. The text in each box gives the same quantity as that coded by the color. Light colors present a security vulnerability to, for example, a passport gate.

Accuracy Terms + Definitions

In biometrics, Type II errors occur when two samples of one person do not match – this is called a **false negative**. Correspondingly, Type I errors occur when samples from two persons do match – this is called a **false positive**. Matches are declared by a biometric system when the native comparison score from the recognition algorithm meets some **threshold**. Comparison scores can be either **similarity scores**, in which case higher values indicate that the samples are more likely to come from the same person, or **dissimilarity scores**, in which case higher values indicate different people. Similarity scores are traditionally computed by **fingerprint** and **face** recognition algorithms, while dissimilarities are used in **iris recognition**. In some cases, the dissimilarity score is a distance; this applies only when **metric** properties are obeyed. In any case, scores can be either **mate** scores, coming from a comparison of one person's samples, or **nonmate** scores, coming from comparison of different persons' samples. The words **genuine** or **authentic** are synonyms for mate, and the word **impostor** is used a synonym for nonmate. The words mate and nonmate are traditionally used in identification applications (such as law enforcement search, or background checks) while genuine and impostor are used in verification applications (such as access control).

A **error tradeoff** characteristic represents the tradeoff between Type II and Type I classification errors. For verification this plots false non-match rate (FNMR) vs. false match rate (FMR) parametrically with T.

The error tradeoff plots are often called **detection error tradeoff (DET)** characteristics or **receiver operating characteristic** (ROC). These serve the same function but differ, for example, in plotting the complement of an error rate (e.g. TMR = 1 – FNMR) and in transforming the axes most commonly using logarithms, to show multiple decades of FMR. More rarely, the function might be the inverse Gaussian function.

More detail and generality is provided in formal biometrics testing standards, see the various parts of ISO/IEC 19795 Biometrics Testing and Reporting. More terms, including and beyond those to do with accuracy, see ISO/IEC 2382-37 Information technology -- Vocabulary -- Part 37: Harmonized biometric vocabulary

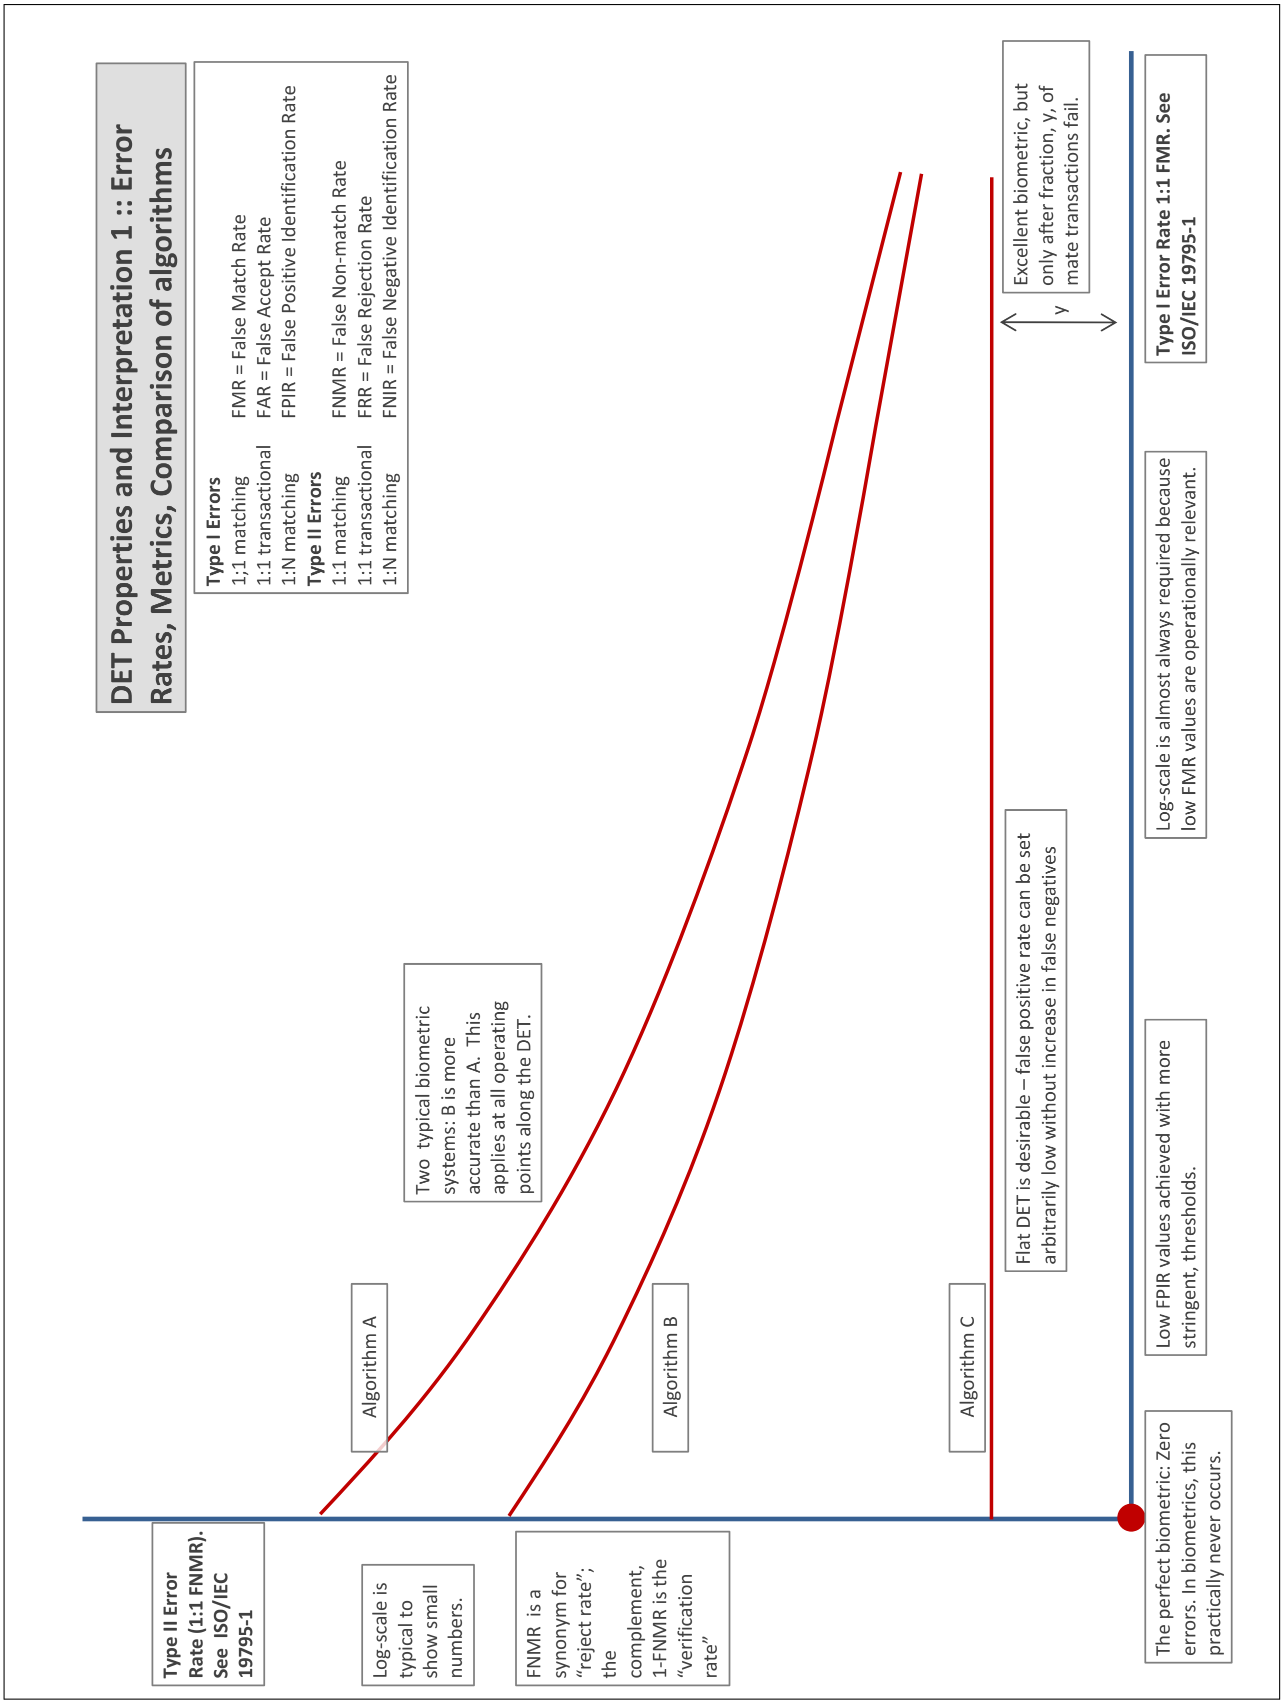
DET Properties and Interpretation 1 :: Error Rates, Metrics, Comparison of algorithms
Type I Errors
1;1 matching FMR = False Match Rate
1:1 transactional FAR = False Accept Rate
1:N matching FPIR = False Positive Identification Rate
Type II Errors
1:1 matching FNMR = False Non-match Rate
1:1 transactional FRR = False Rejection Rate
1:N matching FNIR = False Negative Identification Rate
Two typical biometric systems: B is more accurate than A. This applies at all operating points along the DET.
Algorithm A
Algorithm B
Algorithm C
Flat DET is desirable – false positive rate can be set arbitrarily low without increase in false negatives
Excellent biometric, but only after fraction, y, of mate transactions fail.
y
Type II Error Rate (1:1 FNMR). See ISO/IEC 19795-1
Log-scale is typical to show small numbers.
FNMR is a synonym for "reject rate"; the complement, 1-FNMR is the "verification rate"
The perfect biometric: Zero errors. In biometrics, this practically never occurs.
Low FPIR values achieved with more stringent, thresholds.
Log-scale is almost always required because low FMR values are operationally relevant.
Type I Error Rate 1:1 FMR. See ISO/IEC 19795-1

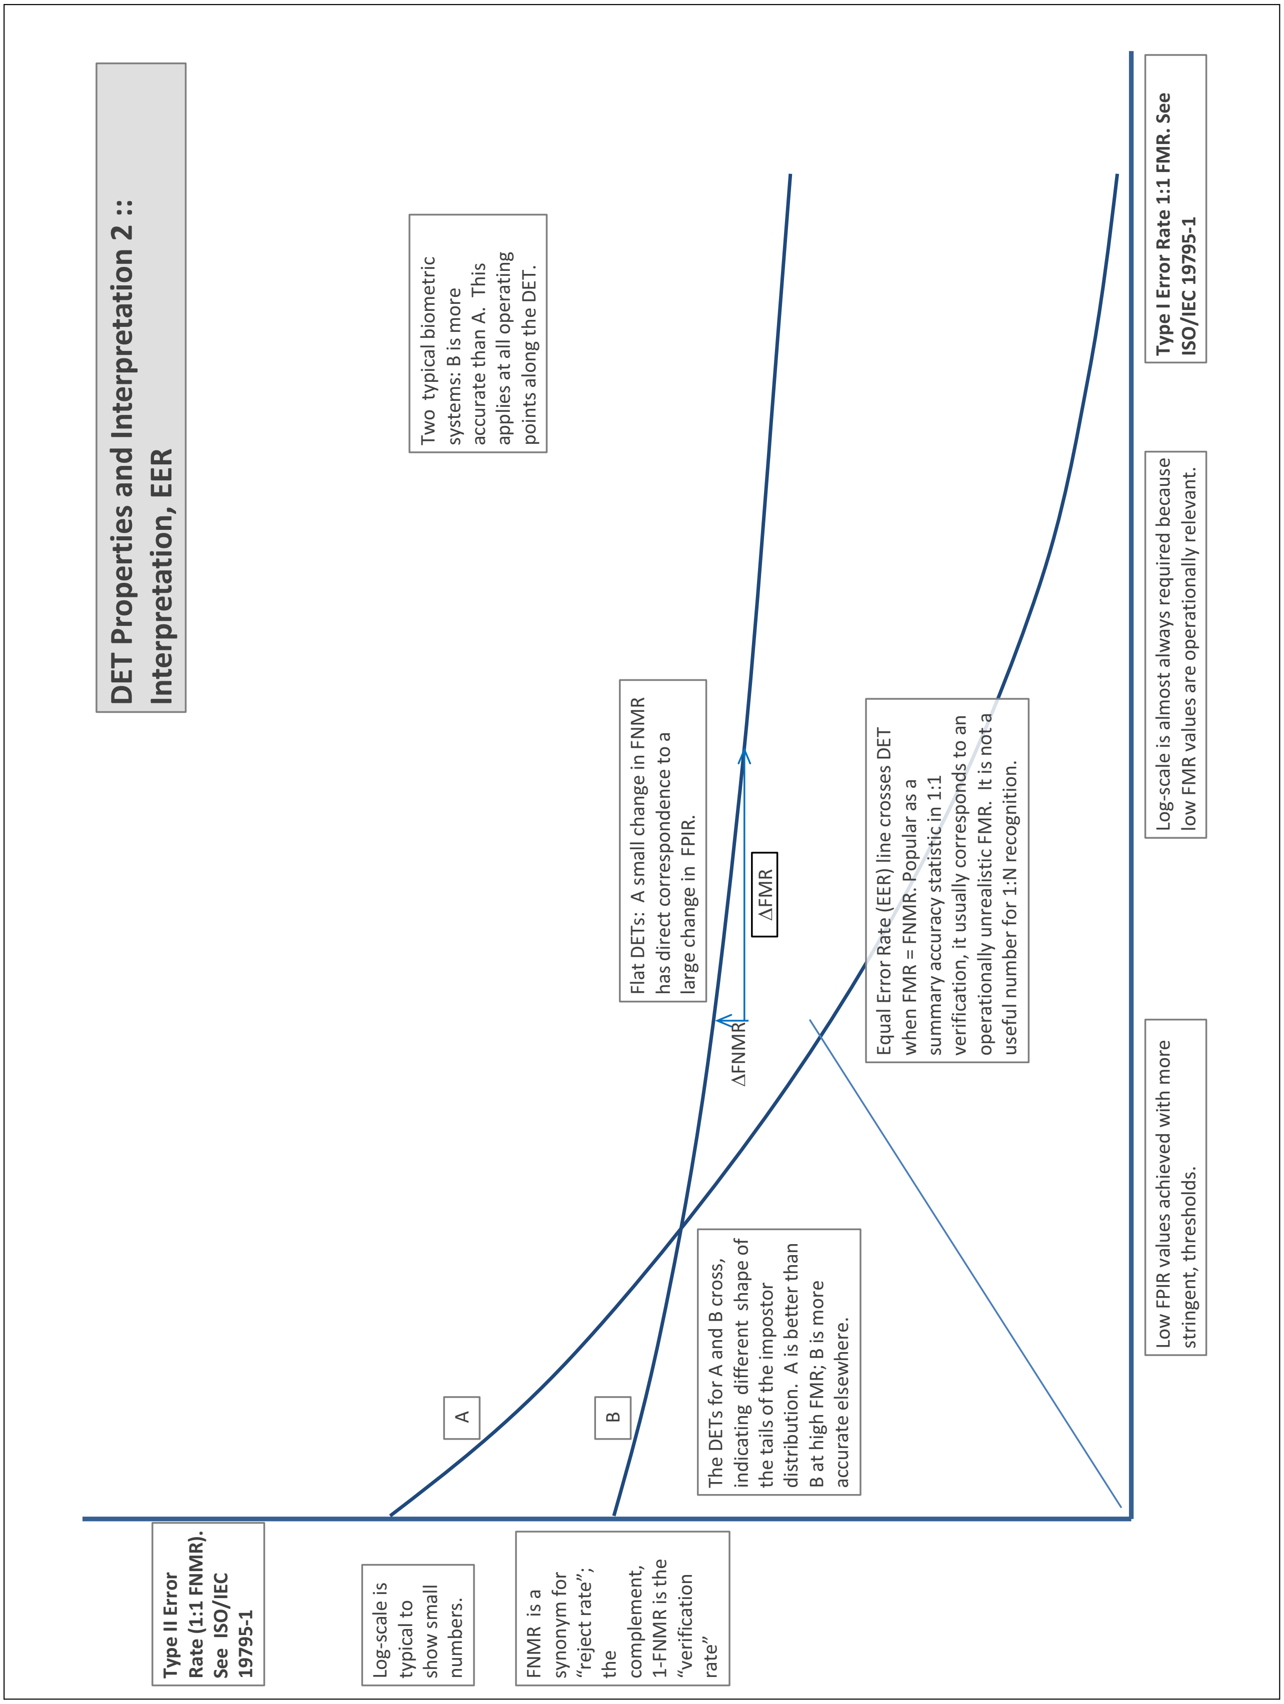
DET Properties and Interpretation 2 :: Interpretation, EER
Two typical biometric systems: B is more accurate than A. This applies at all operating points along the DET.
Type I Error Rate 1:1 FMR. See ISO/IEC 19795-1
Flat DETs: A small change in FNMR has direct correspondence to a large change in FPIR.
ΔFMR
ΔFNMR
Equal Error Rate (EER) line crosses DET when FMR = FNMR. Popular as a summary accuracy statistic in 1:1 verification, it usually corresponds to an operationally unrealistic FMR. It is not a useful number for 1:N recognition.
Log-scale is almost always required because low FMR values are operationally relevant.
Low FPIR values achieved with more stringent, thresholds.
A
B
The DETs for A and B cross, indicating different shape of the tails of the impostor distribution. A is better than B at high FMR; B is more accurate elsewhere.
Type II Error Rate (1:1 FNMR). See ISO/IEC 19795-1
Log-scale is typical to show small numbers.
FNMR is a synonym for "reject rate"; the complement, 1-FNMR is the "verification rate"

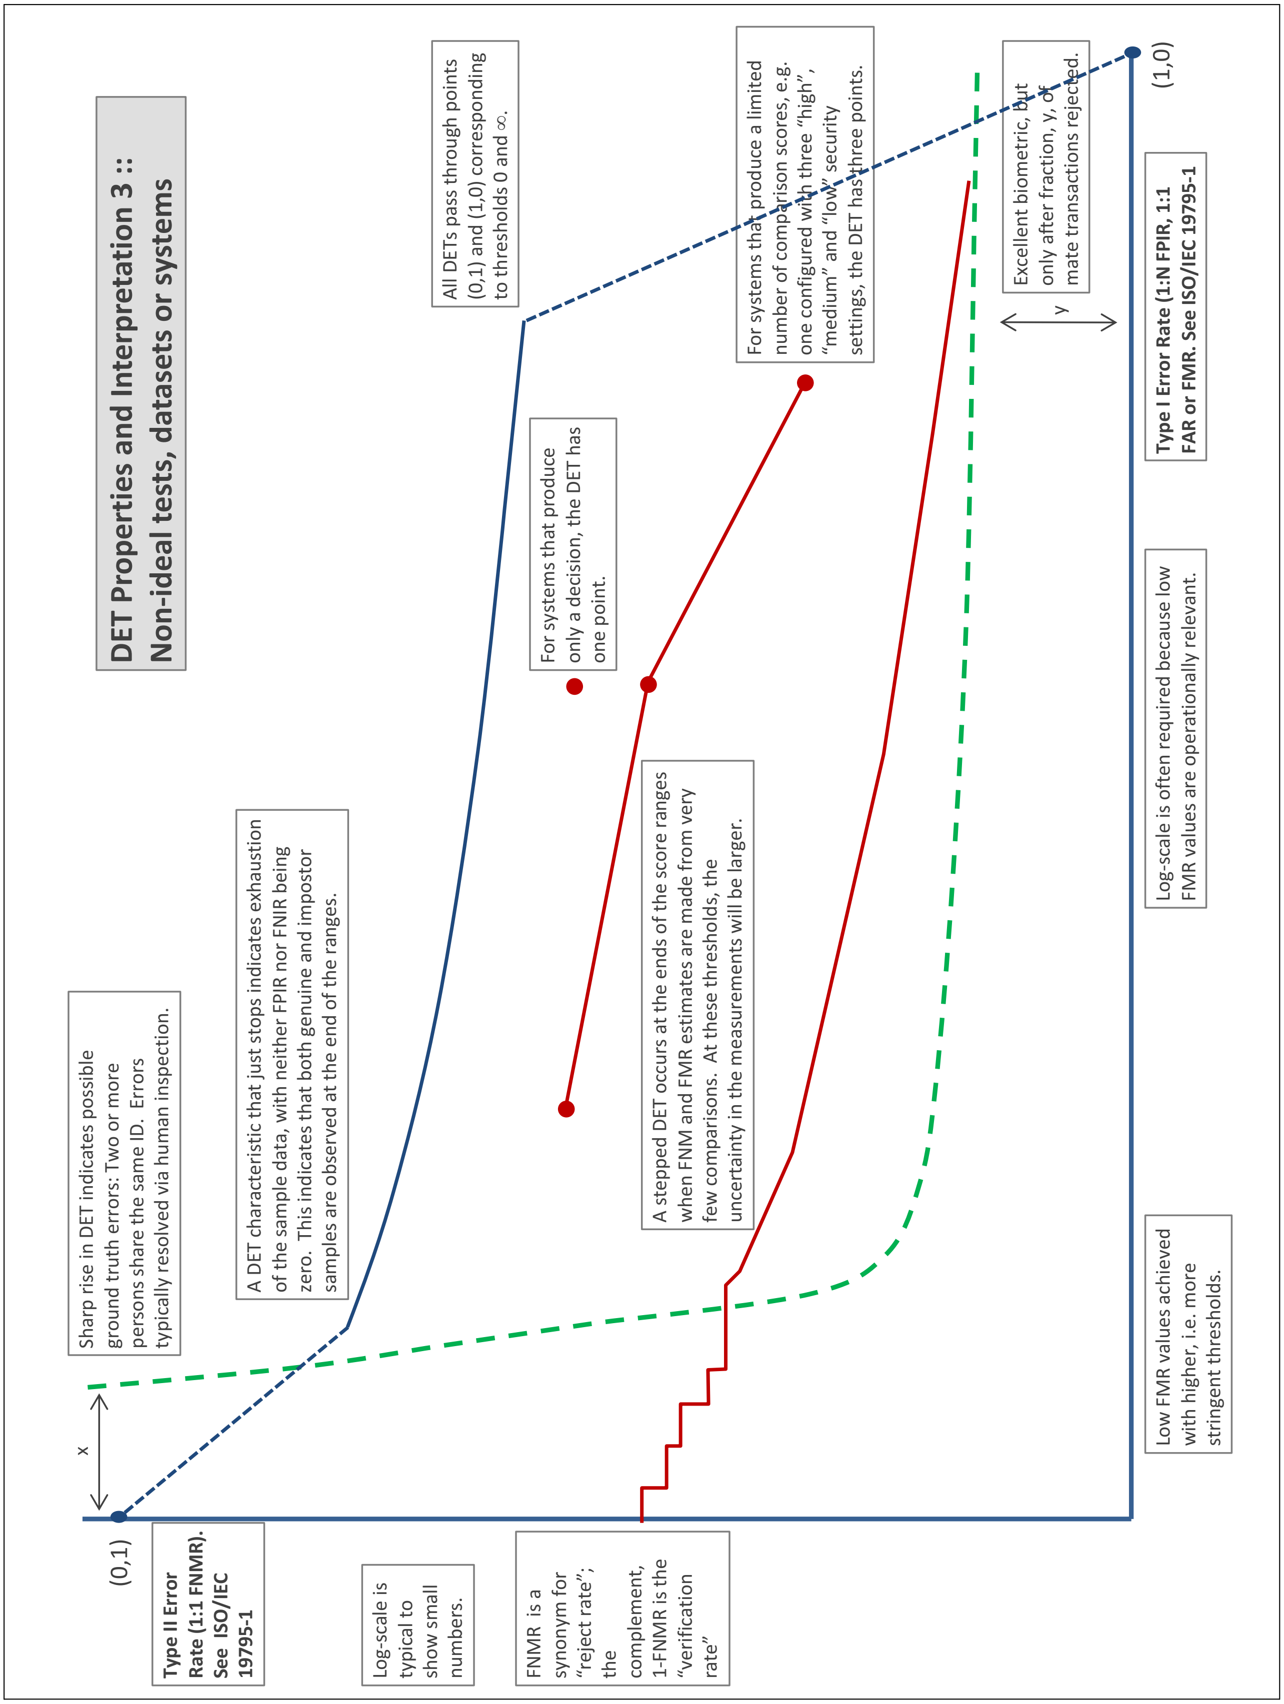

FNMR(T) "False non-match rate"
FMR(T) "False match rate"

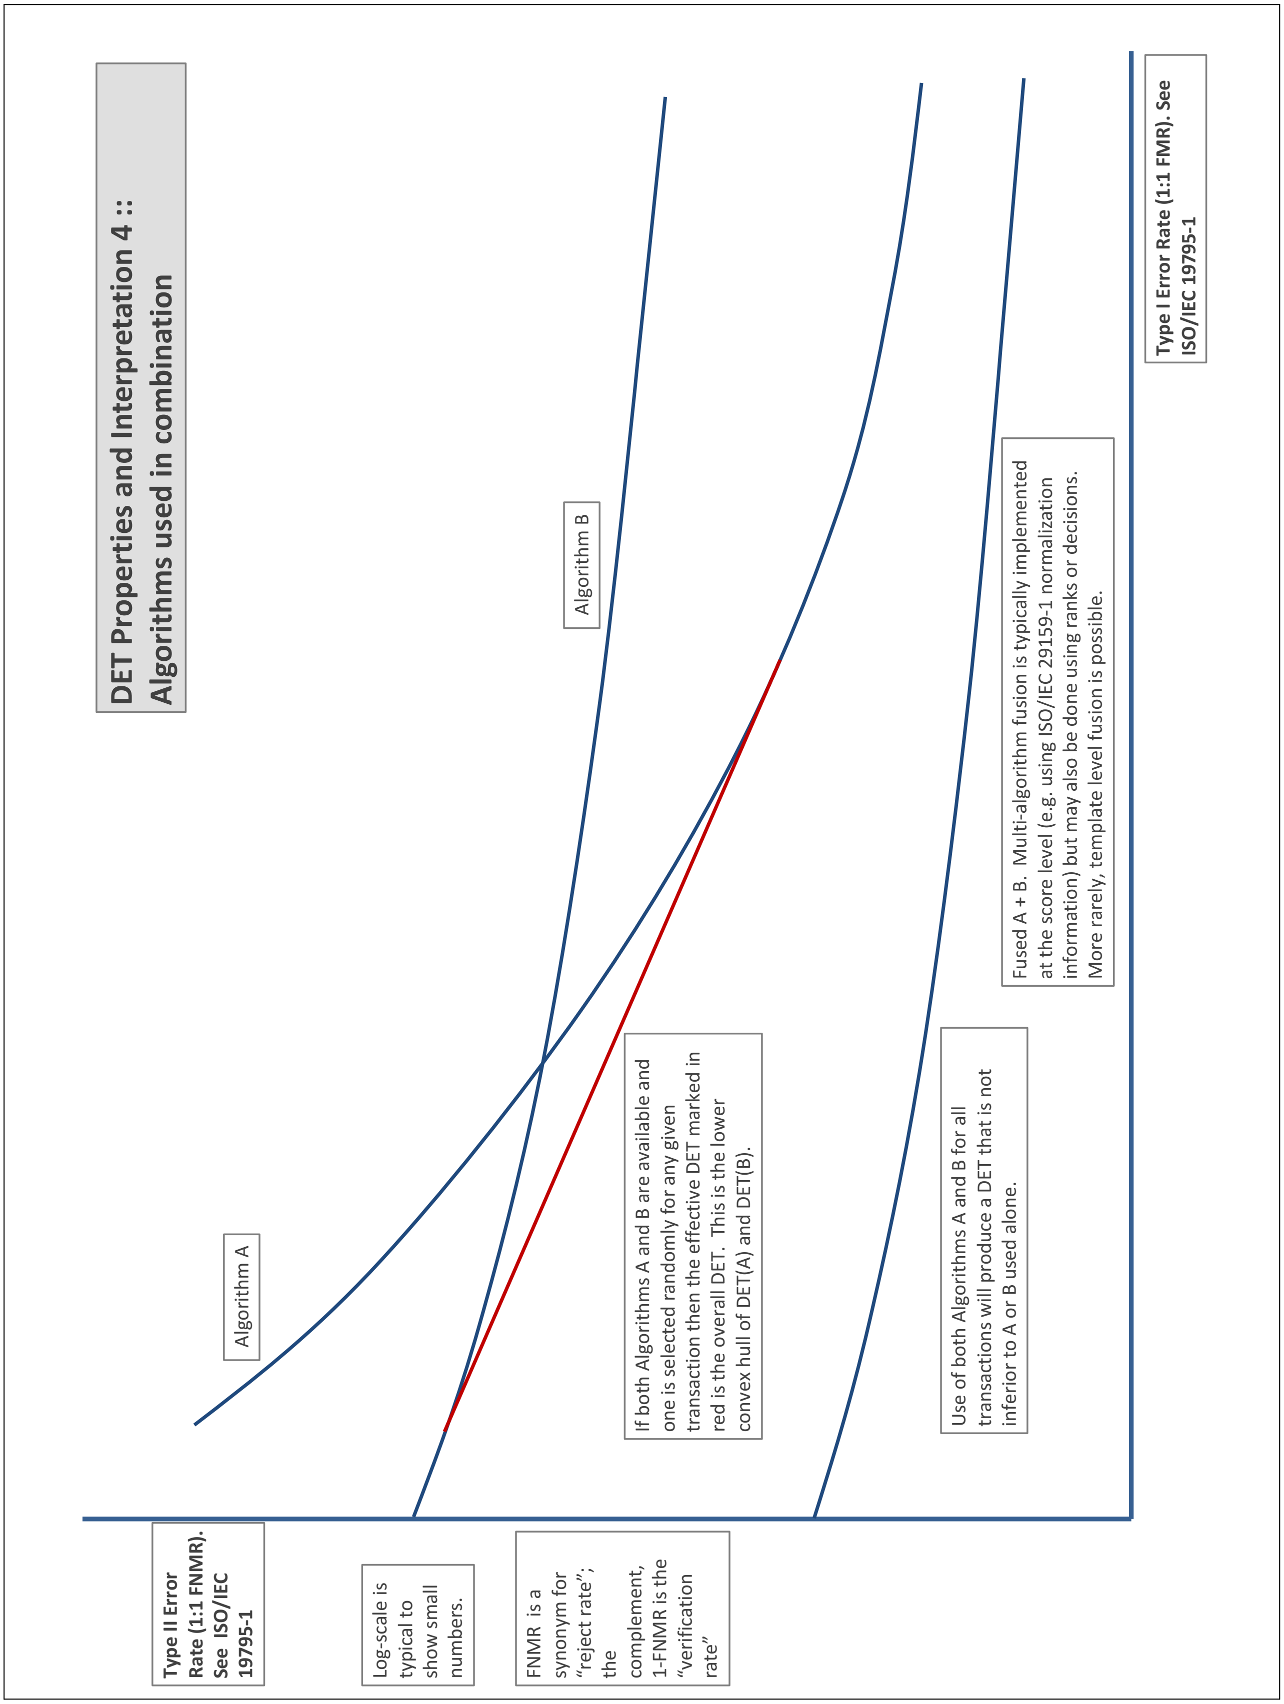

FNMR(T) “False non-match rate”
FMR(T) “False match rate”

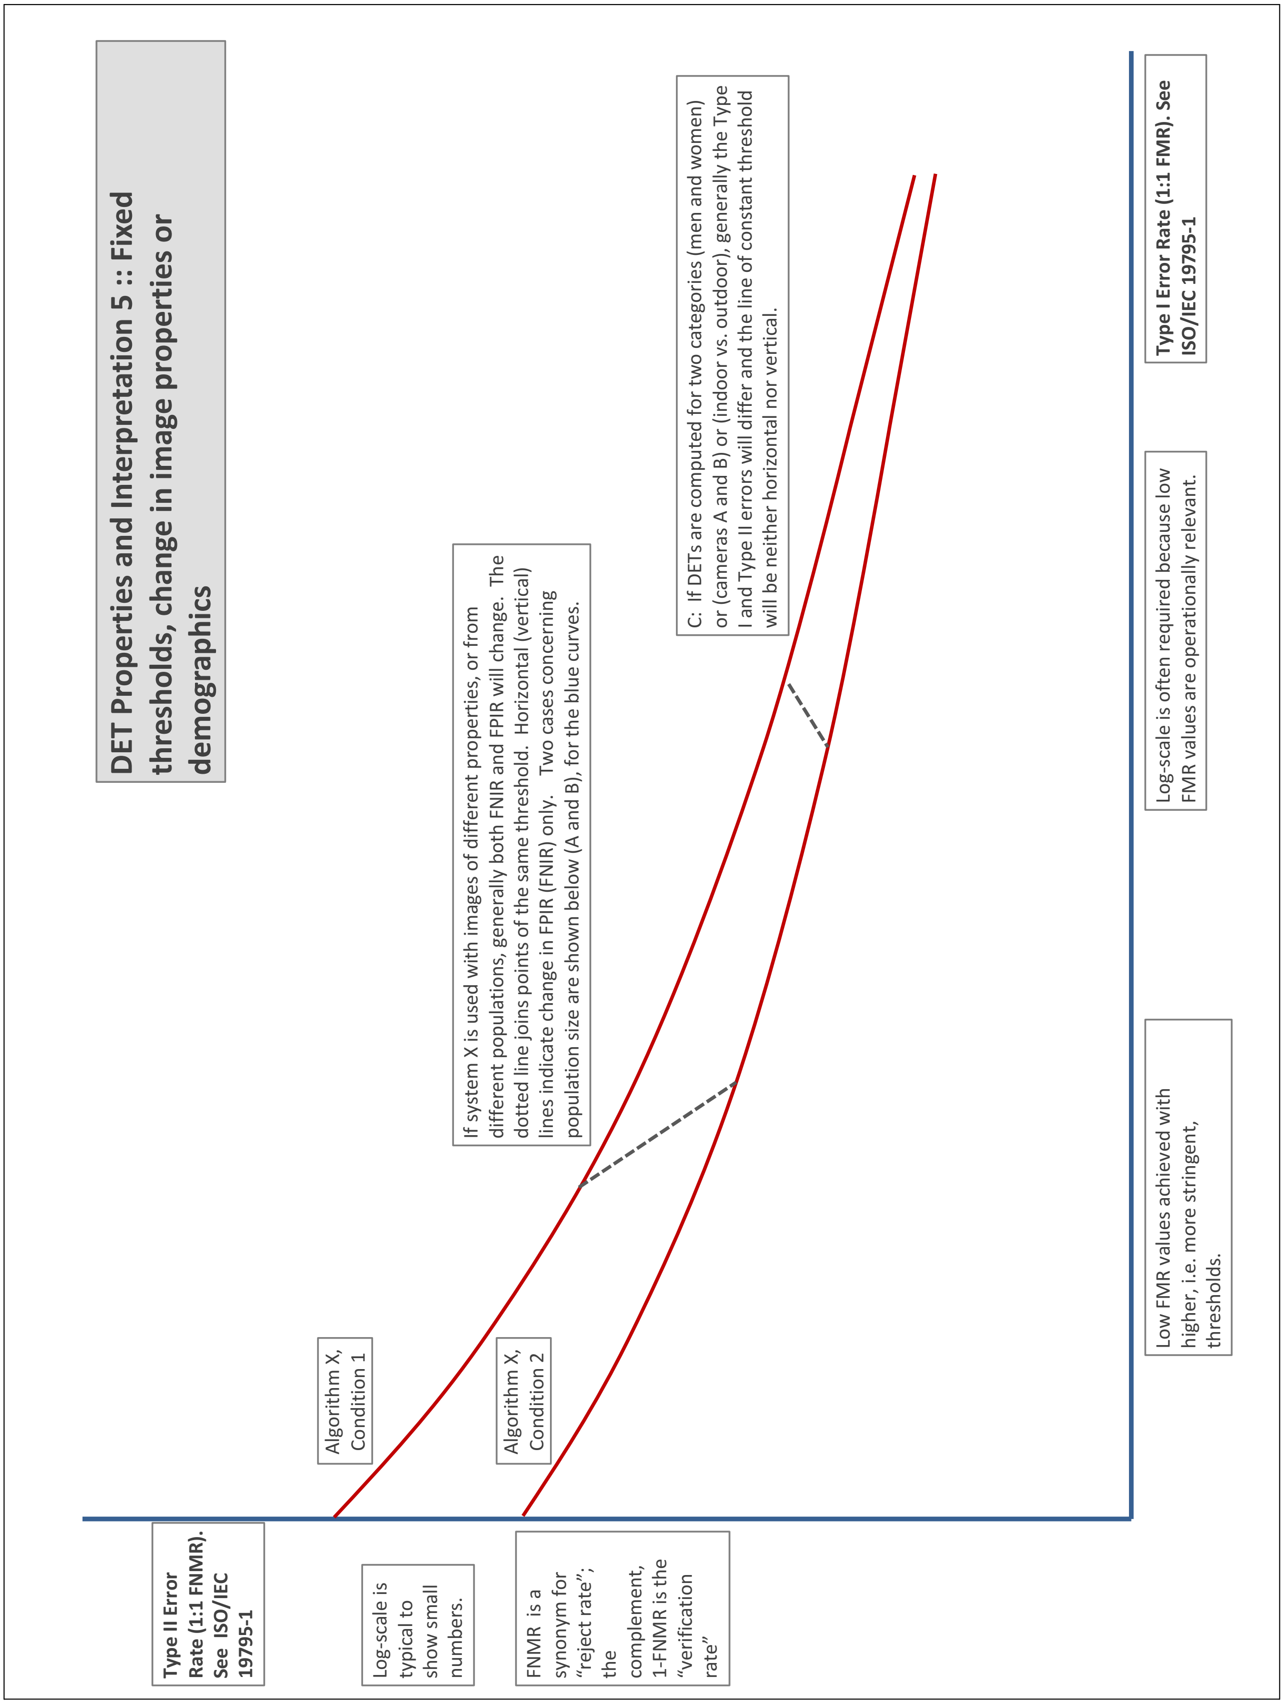
DET Properties and Interpretation 5 :: Fixed thresholds, change in image properties or demographics
If system X is used with images of different properties, or from different populations, generally both FNIR and FPIR will change. The dotted line joins points of the same threshold. Horizontal (vertical) lines indicate change in FPIR (FNIR) only. Two cases concerning population size are shown below (A and B), for the blue curves.
C: If DETs are computed for two categories (men and women) or (cameras A and B) or (indoor vs. outdoor), generally the Type I and Type II errors will differ and the line of constant threshold will be neither horizontal nor vertical.
Algorithm X, Condition 1
Algorithm X, Condition 2
Type II Error Rate (1:1 FNMR). See ISO/IEC 19795-1
Log-scale is typical to show small numbers.
FNMR is a synonym for "reject rate", the complement, 1-FNMR is the "verification rate"
Low FMR values achieved with higher, i.e. more stringent, thresholds.
Log-scale is often required because low FMR values are operationally relevant.
Type I Error Rate (1:1 FMR). See ISO/IEC 19795-1