

Figure A.1: All generators used in the MMCGAN architecture. While the image heads and most of the shared generator are fully convolutional, the tabular generator uses residual blocks consisting of fully connected layers.

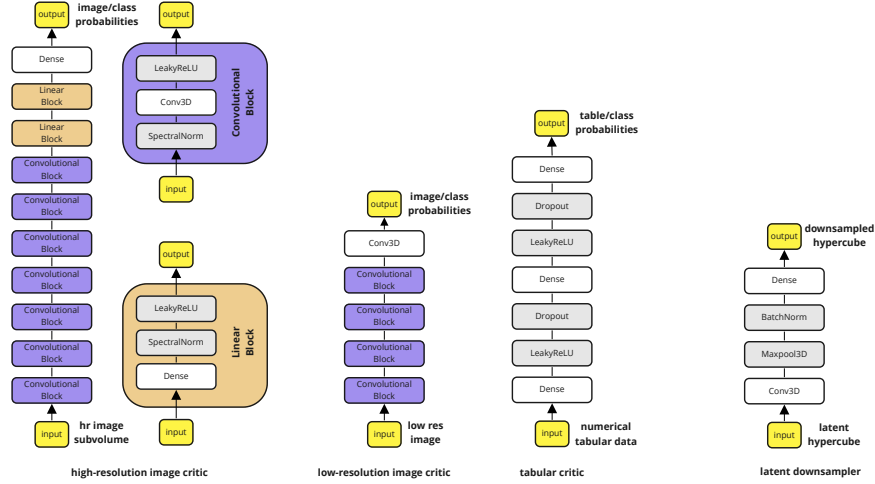


Figure A.2: All critics and the downsampling module used in the MMCGAN architecture.

Table 1: Optimization parameters for individual model components in MMCGAN architecture.

| model component | learning rate | β_1 | β_2 |
|-----------------------|--------------------|-----------|-----------|
| shared generator | 1×10^{-4} | 0 | 0.999 |
| image generators | 1×10^{-4} | 0 | 0.999 |
| tabular generator | 2×10^{-4} | 0.5 | 0.9 |
| image discriminators | 4×10^{-4} | 0 | 0.999 |
| tabular discriminator | 2×10^{-4} | 0.5 | 0.9 |
| tabular bottleneck | 2×10^{-4} | 0.5 | 0.9 |




Figure A.3: Random image samples generated using the MMCGAN model. Horizontal axis shows 2×12 different virtual patients, vertical axis shows slices through main axes.




Figure A.4: Real and synthetic data are compared with respect to percentage of scan occupied by fat, tissue, and water, compared to Body Mass Index of subjects.

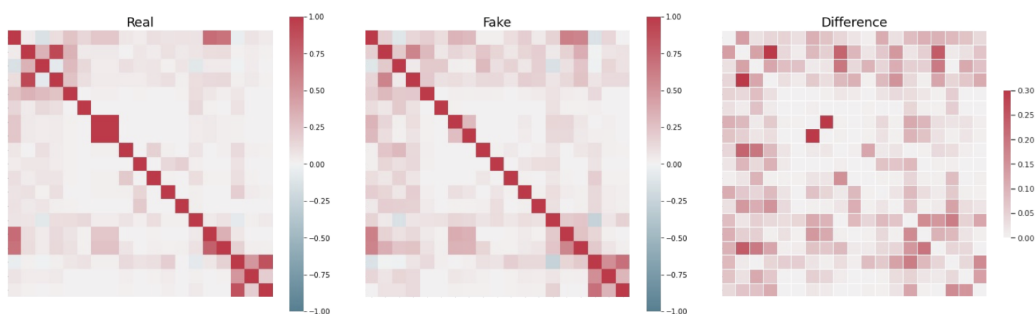


Figure A.5: Comparison of bivariate cross-correlation between real (left) and synthetic (middle) data. The difference signal is shown in the rightmost image. Much of the data set's correlation is retained in the synthetic surrogate.