1. presents the complete workflow, from paired-volume preprocessing and baseline ultrasound-to-MRI translation to radiomics-guided FiLM conditioning and multi-level evaluation.

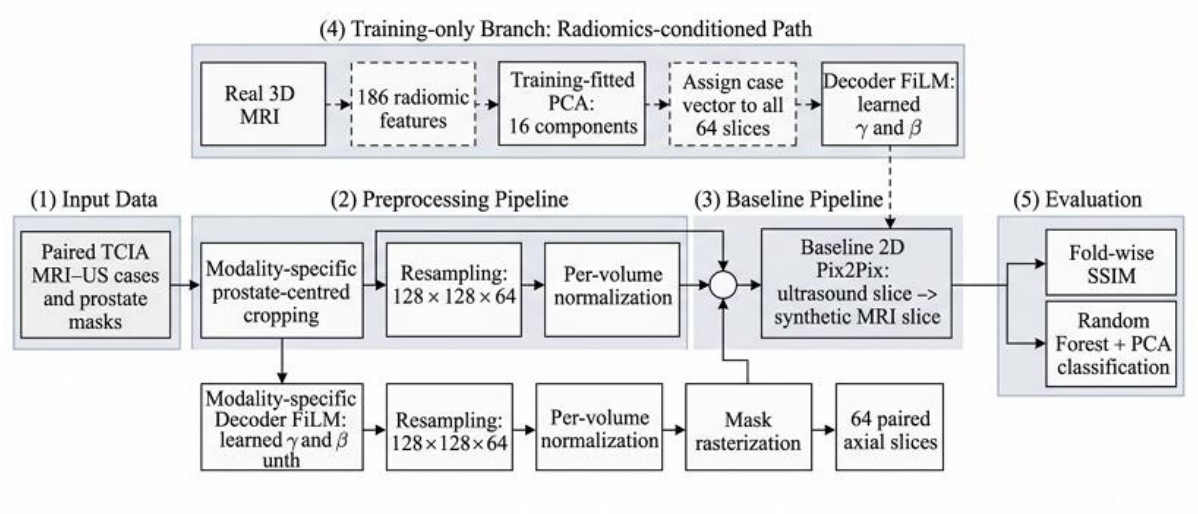

Overall architecture of the proposed radiomics-aware ultrasound-to-MRI translation framework. Case-level three-dimensional radiomic descriptors are reduced using training-fitted principal component analysis (PCA) and used to condition the decoder of a two-dimensional Pix2Pix generator through feature-wise linear modulation (FiLM).

2. shows a representative central pair produced by this pipeline.

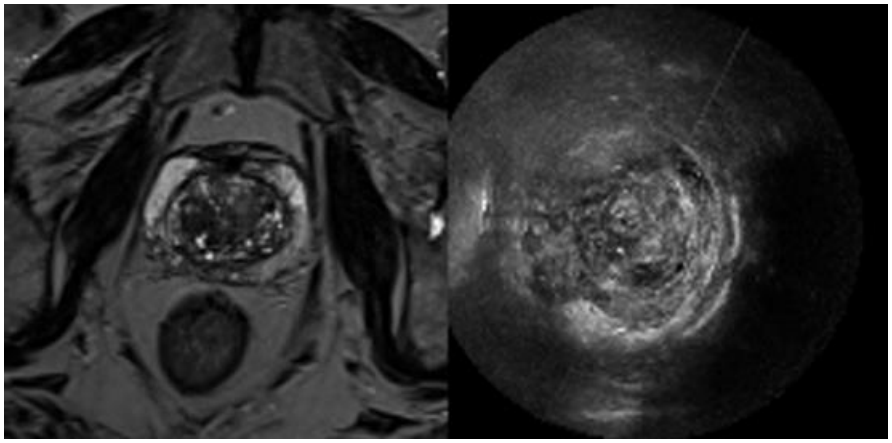

Representative central axial slice pair after preprocessing. The MRI slice is shown on the left, and the corresponding ultrasound slice is shown on the right. Both images originate from volumes cropped around their modality-specific prostate centres and resampled to the common 128 × 128 × 64 grid.

3. summarizes the verified U-Net generator, conditional PatchGAN discriminator, and adversarial and reconstruction objectives used in the baseline model.

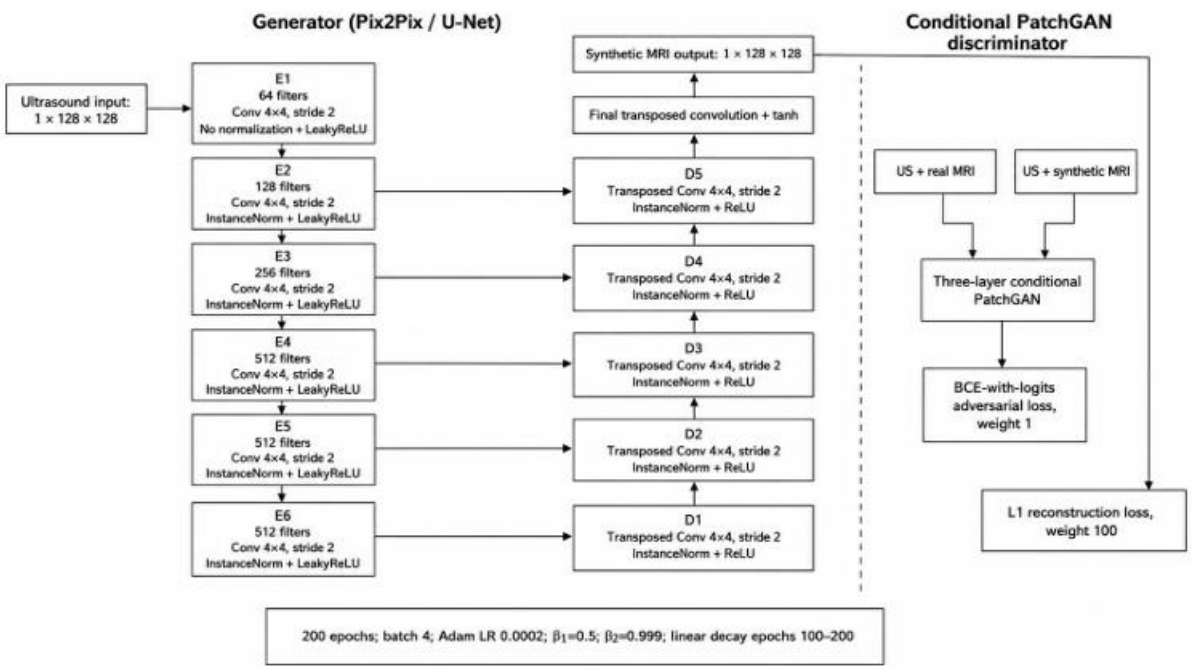

Verified baseline Pix2Pix architecture for paired ultrasound-to-MRI translation. The U-Net generator is optimized using adversarial and pixel-wise reconstruction objectives, while the conditional PatchGAN discriminator evaluates paired ultrasound–MRI inputs.

4. summarizes the case-level extraction and dimensionality-reduction pathways used to obtain radiomic representations for classification and FiLM conditioning.

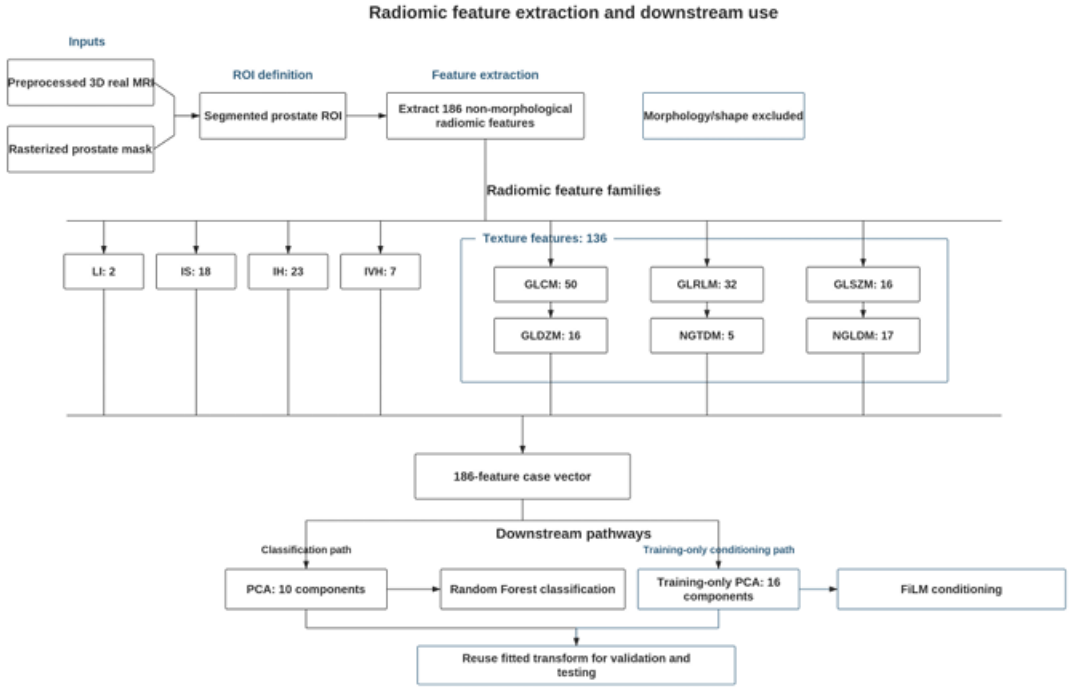

Three-dimensional radiomic feature extraction and PCA pathways. The segmented MRI prostate volume produces 186 non-morphological radiomic features, with separate PCA configurations supporting downstream classification and generator conditioning.

5. the case-level radiomic vector is transformed into learned scaling and shifting parameters that modulate decoder feature maps during MRI synthesis.

Radiomics-guided FiLM conditioning within the Pix2Pix generator. A 16-component case-level PCA representation is mapped to learned channel-wise scale and shift parameters, which are subsequently applied to decoder feature maps.

Table 3.1. Verified baseline Pix2Pix training configuration.

Parameter	**Value**
Input and target size	128×128, grayscale
Epochs	200
Batch size	4
Initial learning rate	2×10^{-4}
Adam coefficients	$\beta_1 = 0.5, \beta_2 = 0.999$
Adversarial loss weight	$\lambda_{\text{GAN}} = 1$
Reconstruction loss weight	$\lambda_{L1} = 100$
Generator base filters	64
Generator downsampling blocks	6
Discriminator convolutional layers	3
Random seed	42