Diffusion posterior sampling: methodology and applications to ECG reconstruction

Lisa Bedin

Ecole polytechnique

Joint work with:

Yazid Janati, Gabriel Cardoso, Badr Moufad, Alain Durmus, Randal Douc, Jimmy Olsson, Eric Moulines

Context

ECGs

Electrodes Placement

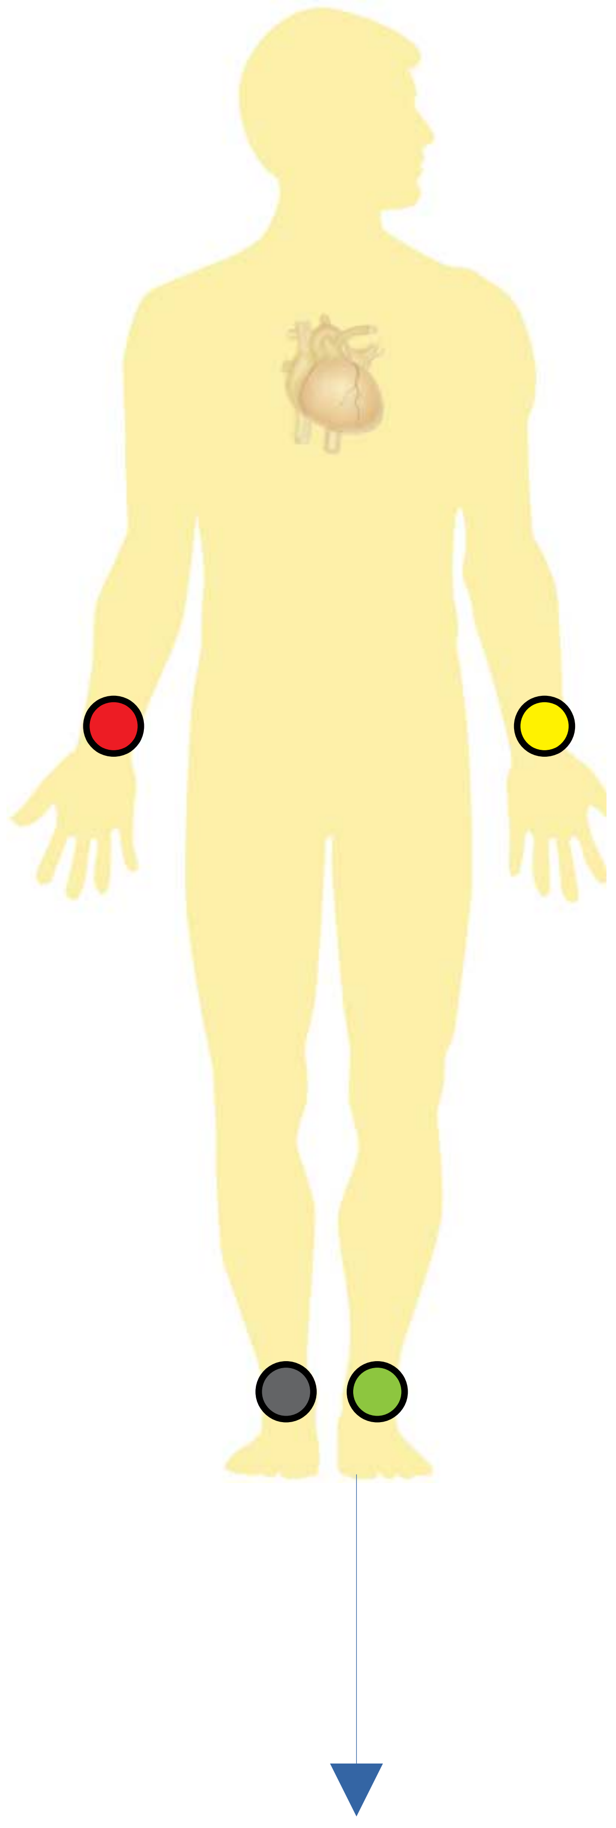

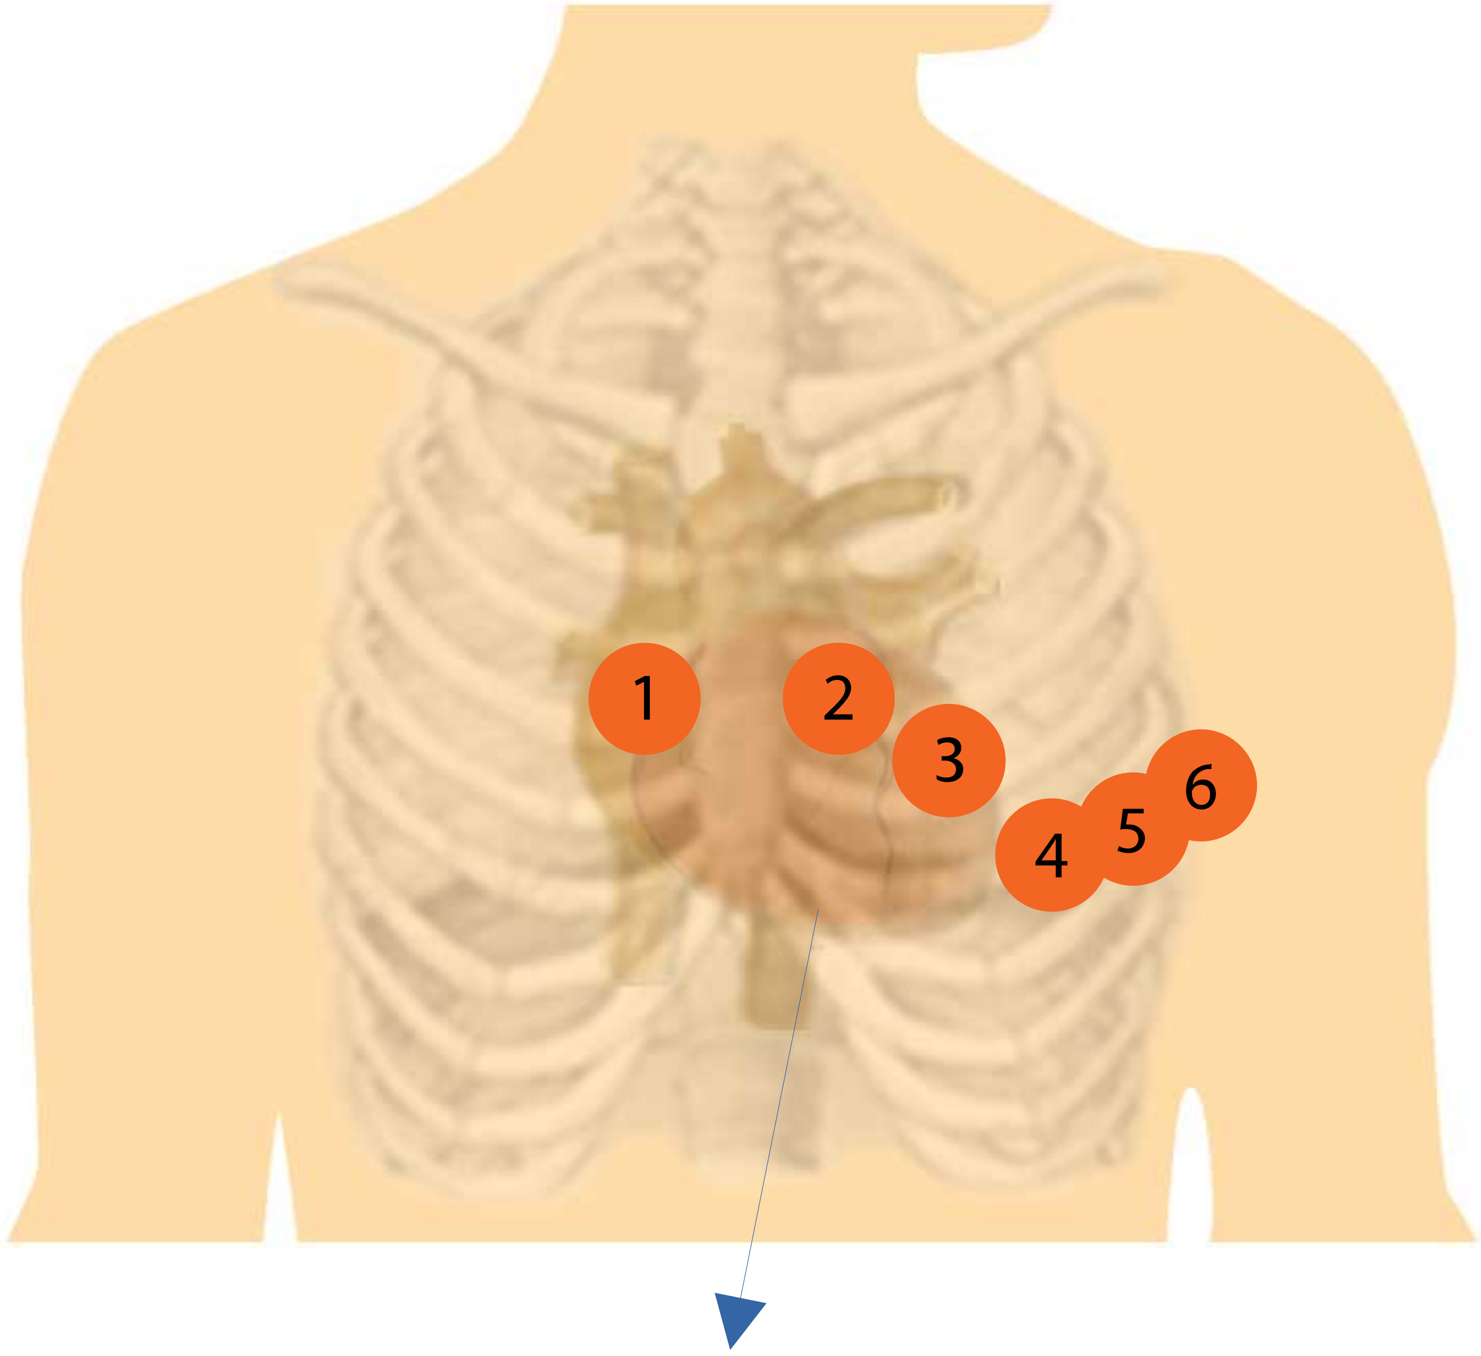

Precordial Leads : V1 – V6
(closer to the heart)

Limbs Leads :
aVL, aVR, aVF

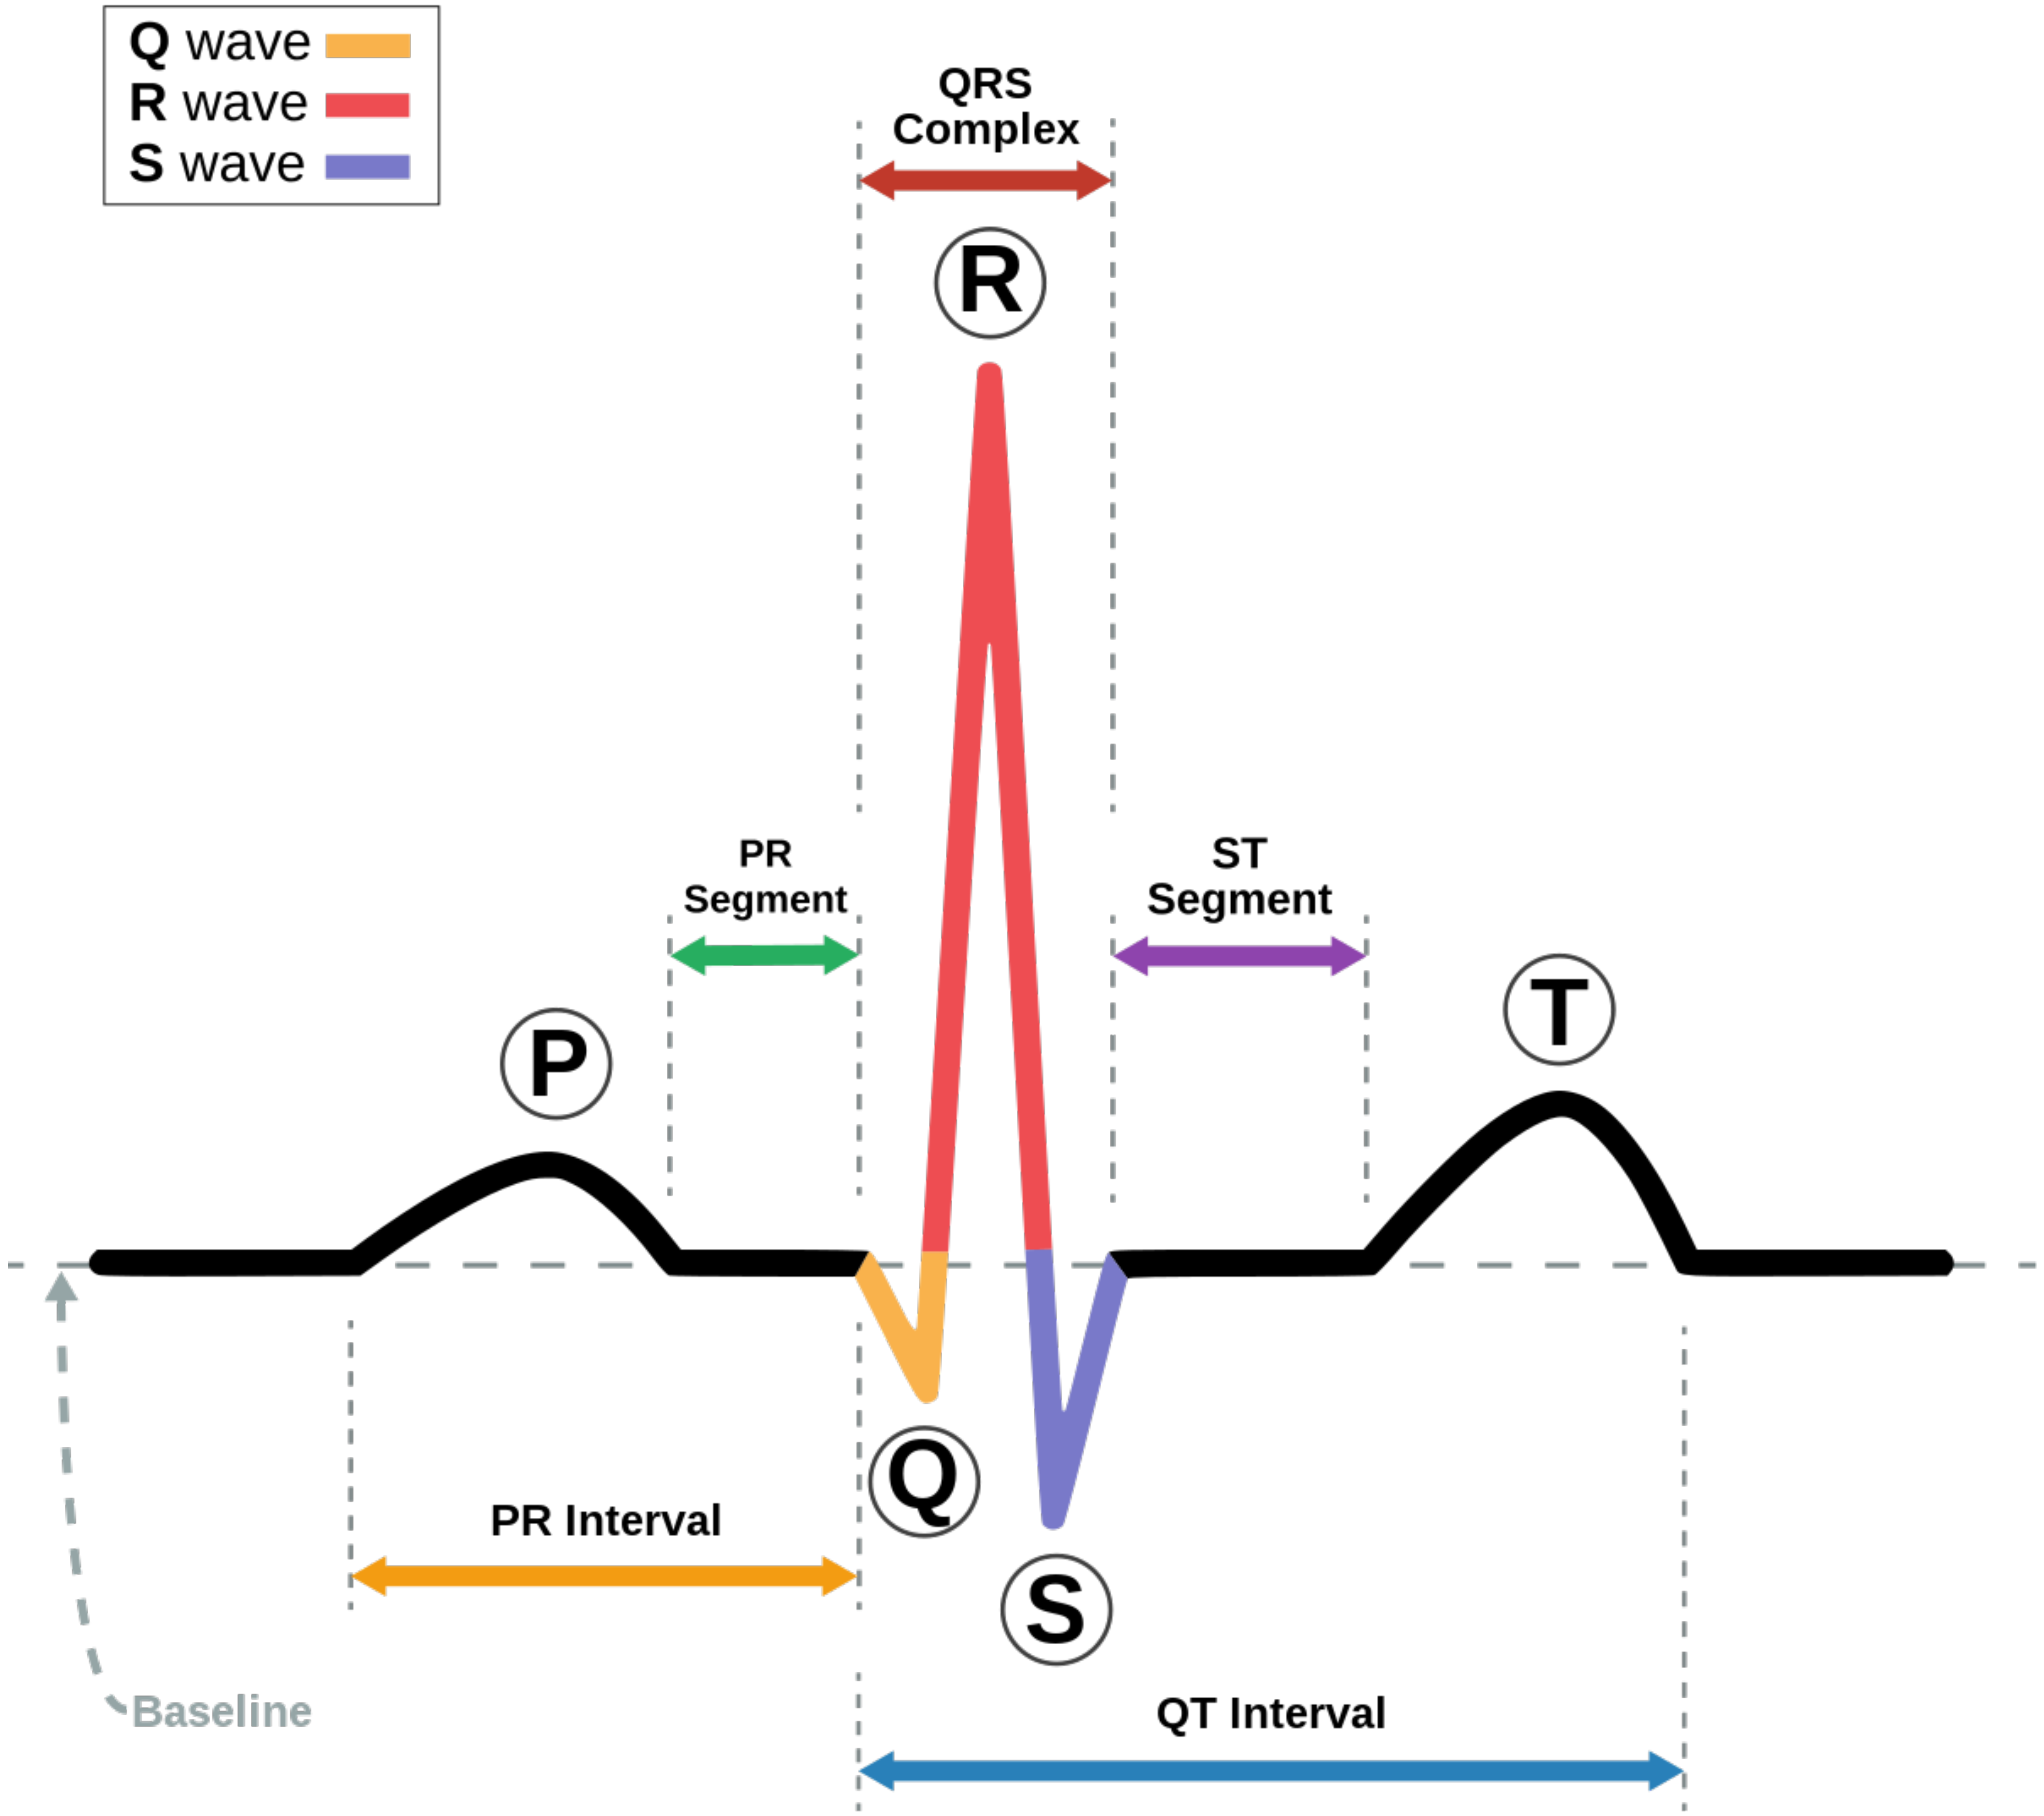

ECGs

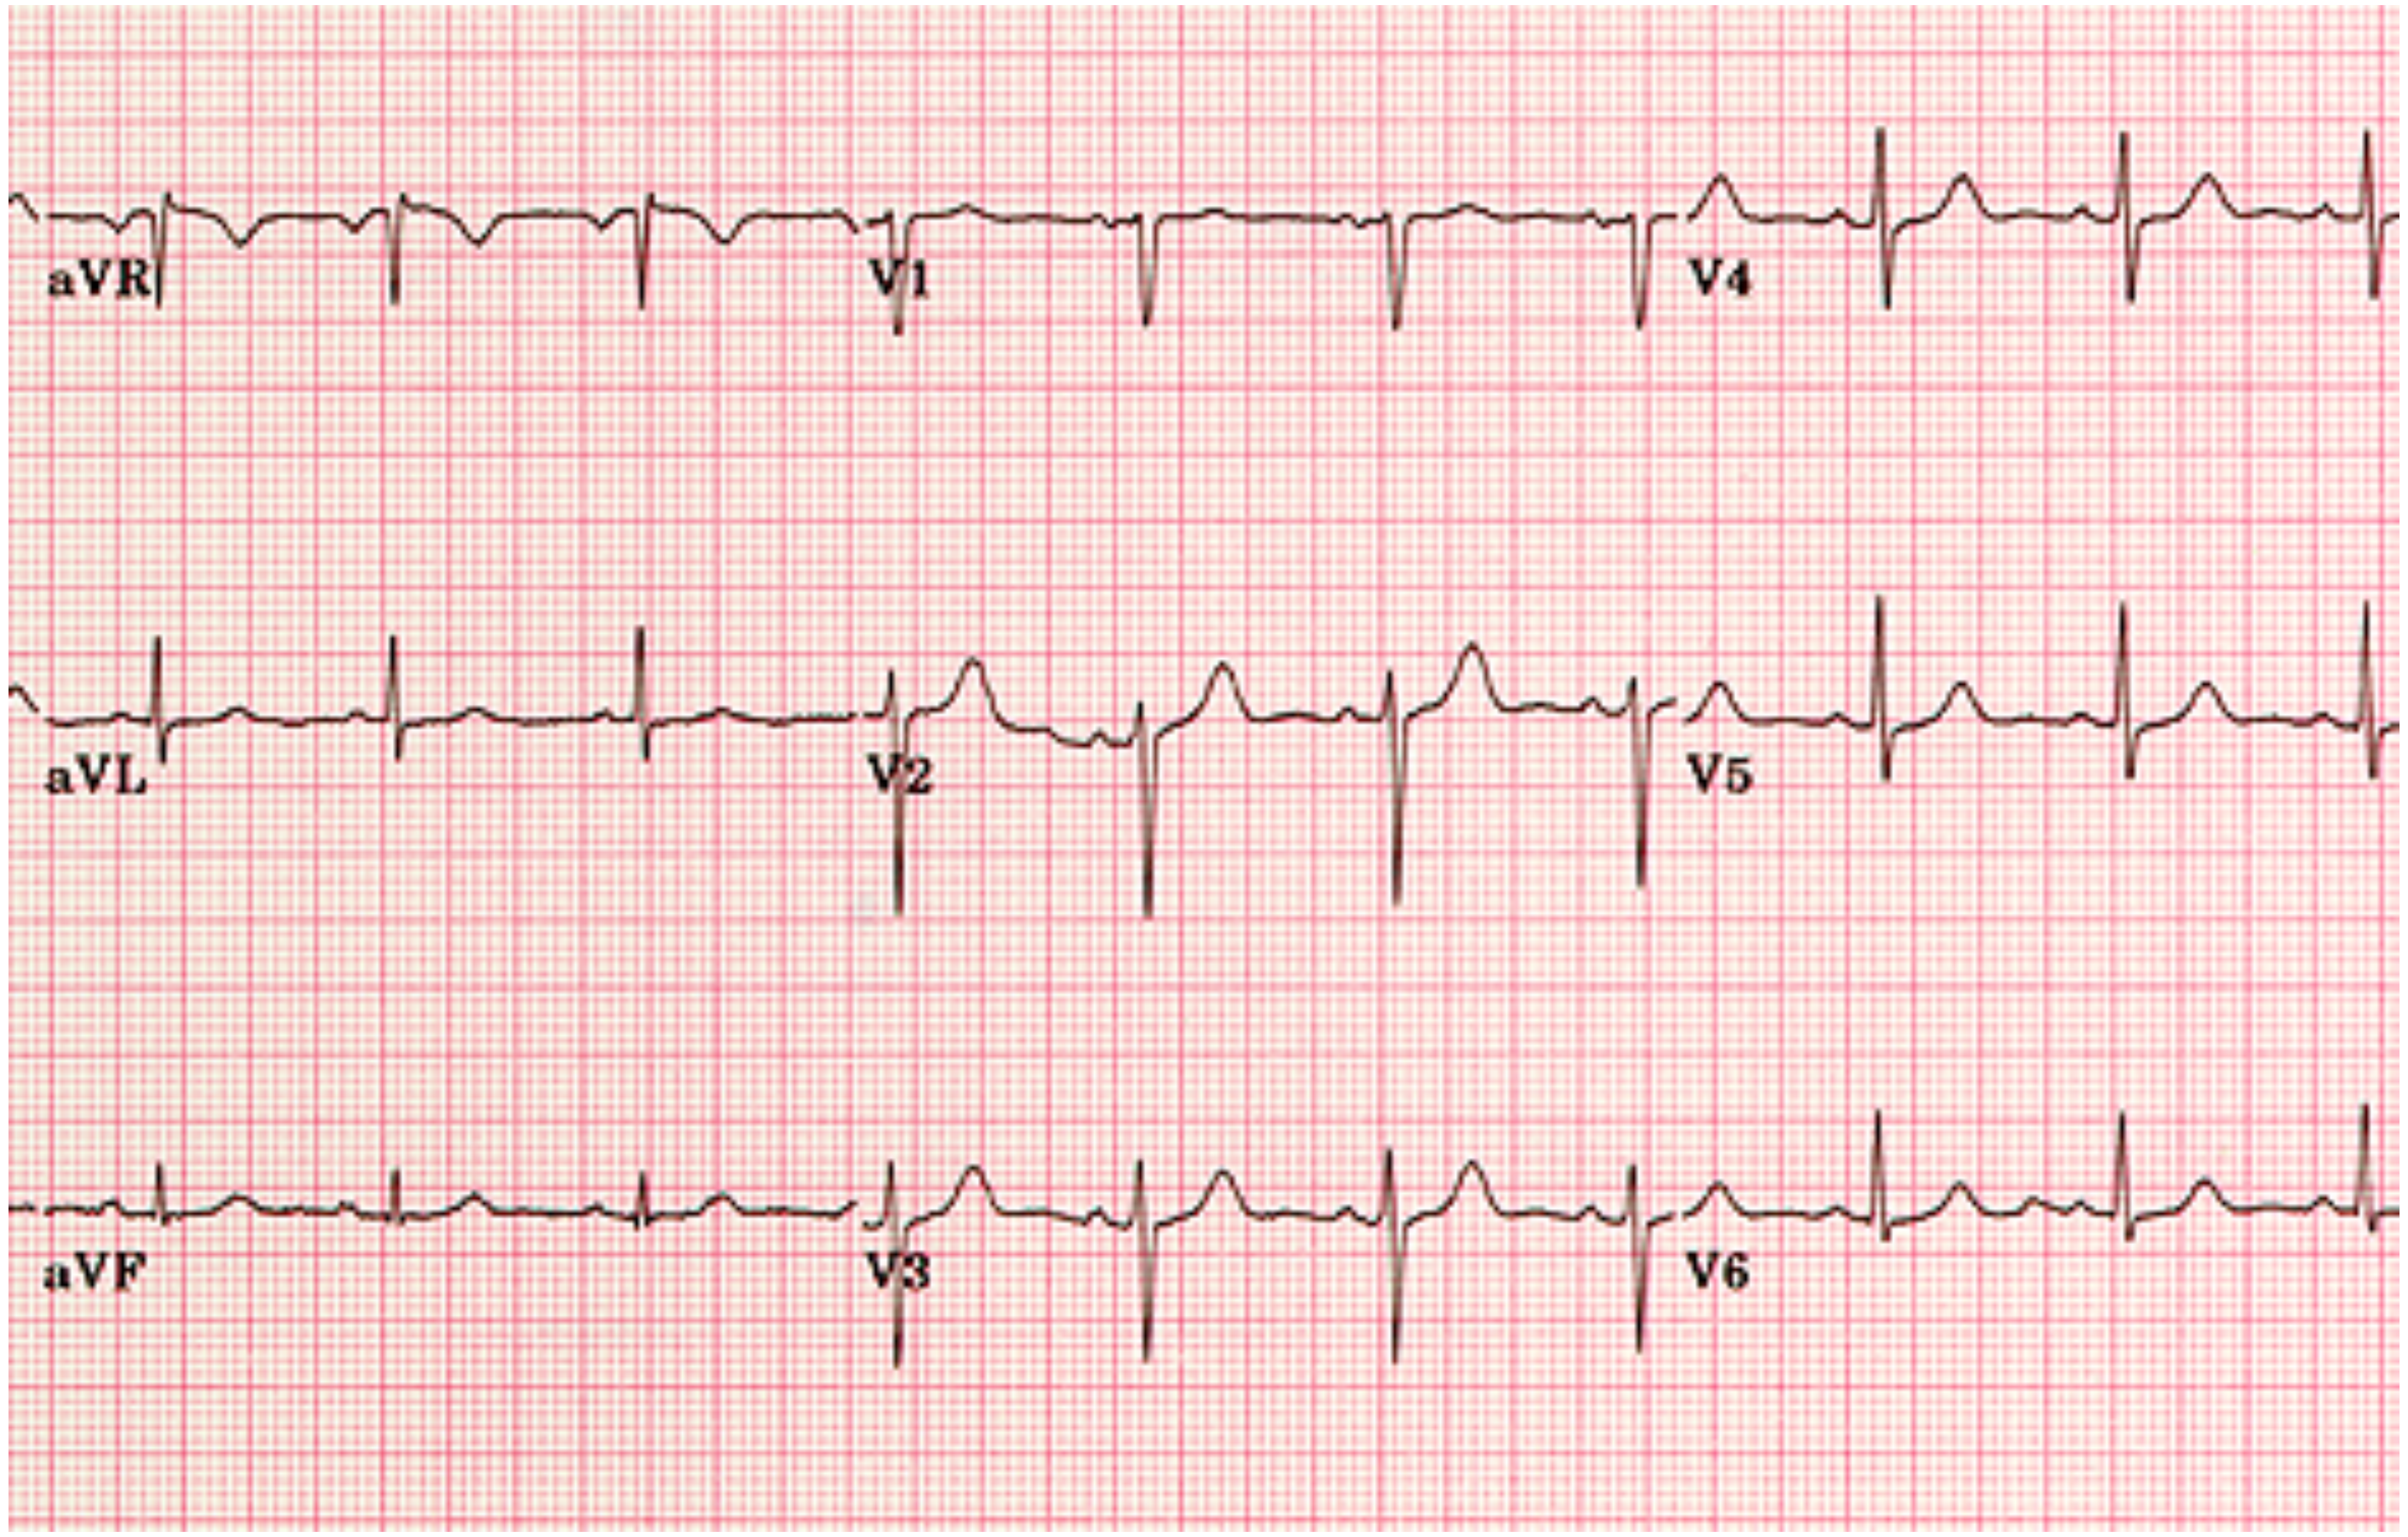

- Acquisition noise : baseline wander, electrode motion
- Incomplete ECG e.g., missing leads in portative devices Kardia, Smart Watch ...

Question : how to recover an ECG from noisy data / from incomplete data ?

Denoising Diffusion models

Diffusion models

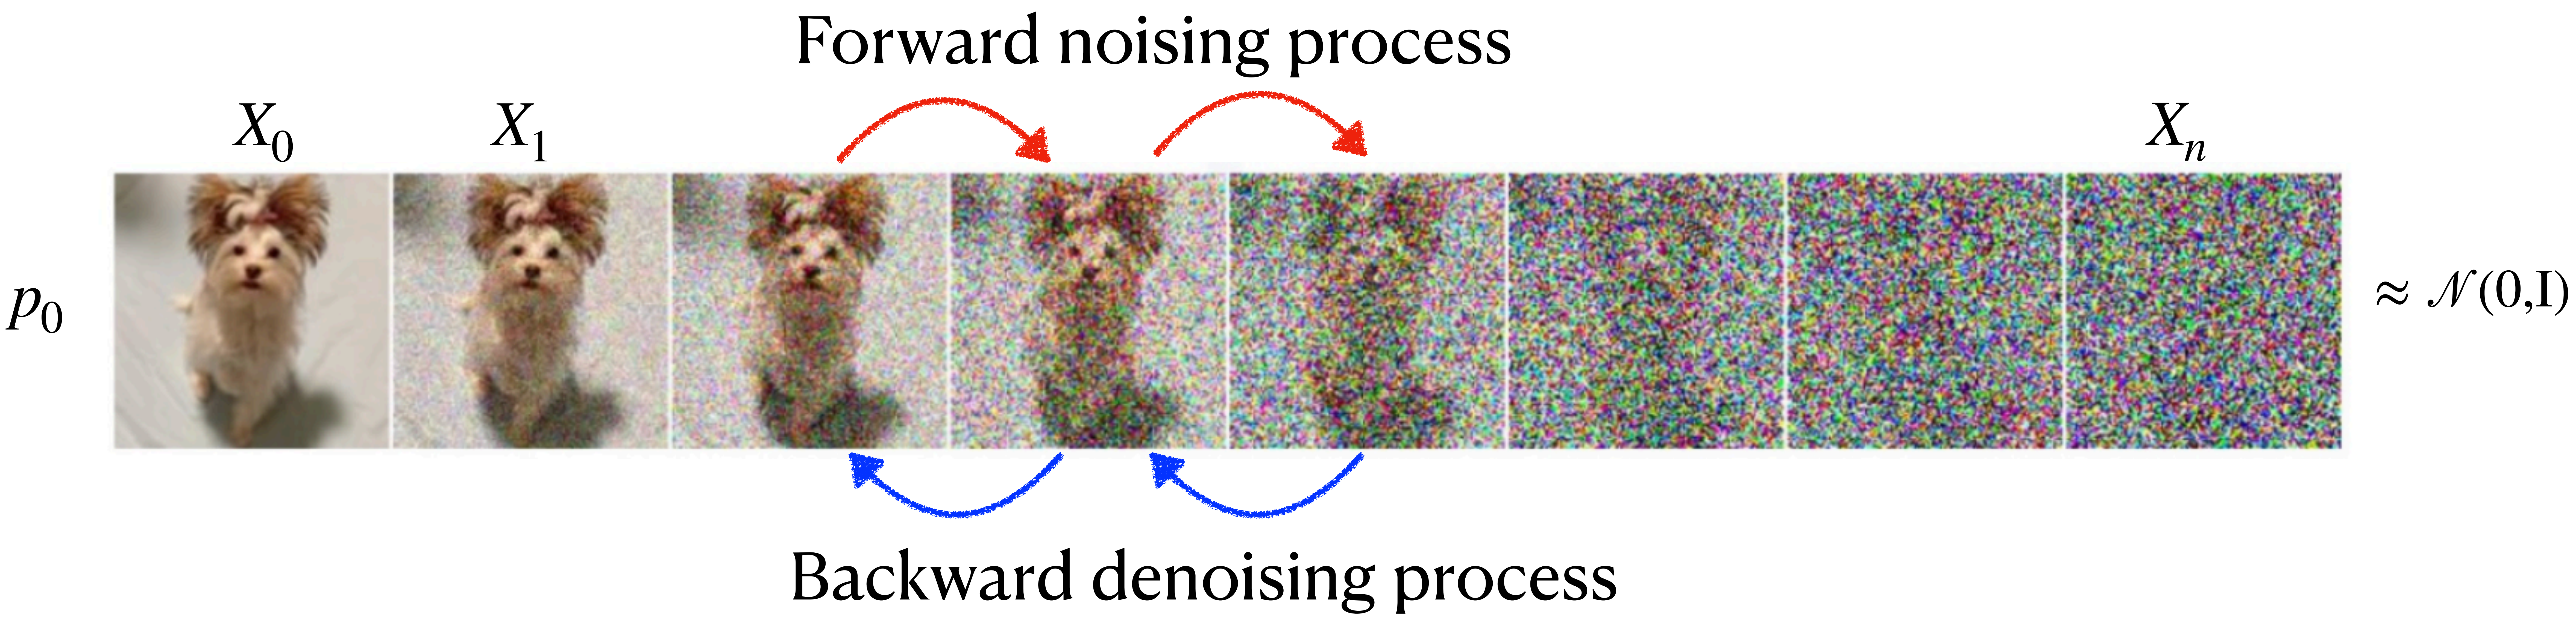

[Ho, J., Jain, A. and Abbeel, P., 2020. Denoising diffusion probabilistic models.]

Diffusion models

$$q_{k+1|k} = \mathrm{N}(x_{k+1}; (\alpha_{k+1}/\alpha_k)x_k, (1 - \alpha_{k+1}^2/\alpha_k^2)\mathrm{I})$$

where $\alpha_0 = 1,\ \alpha_n \approx 0$

X_0 X_1 X_n

p_0

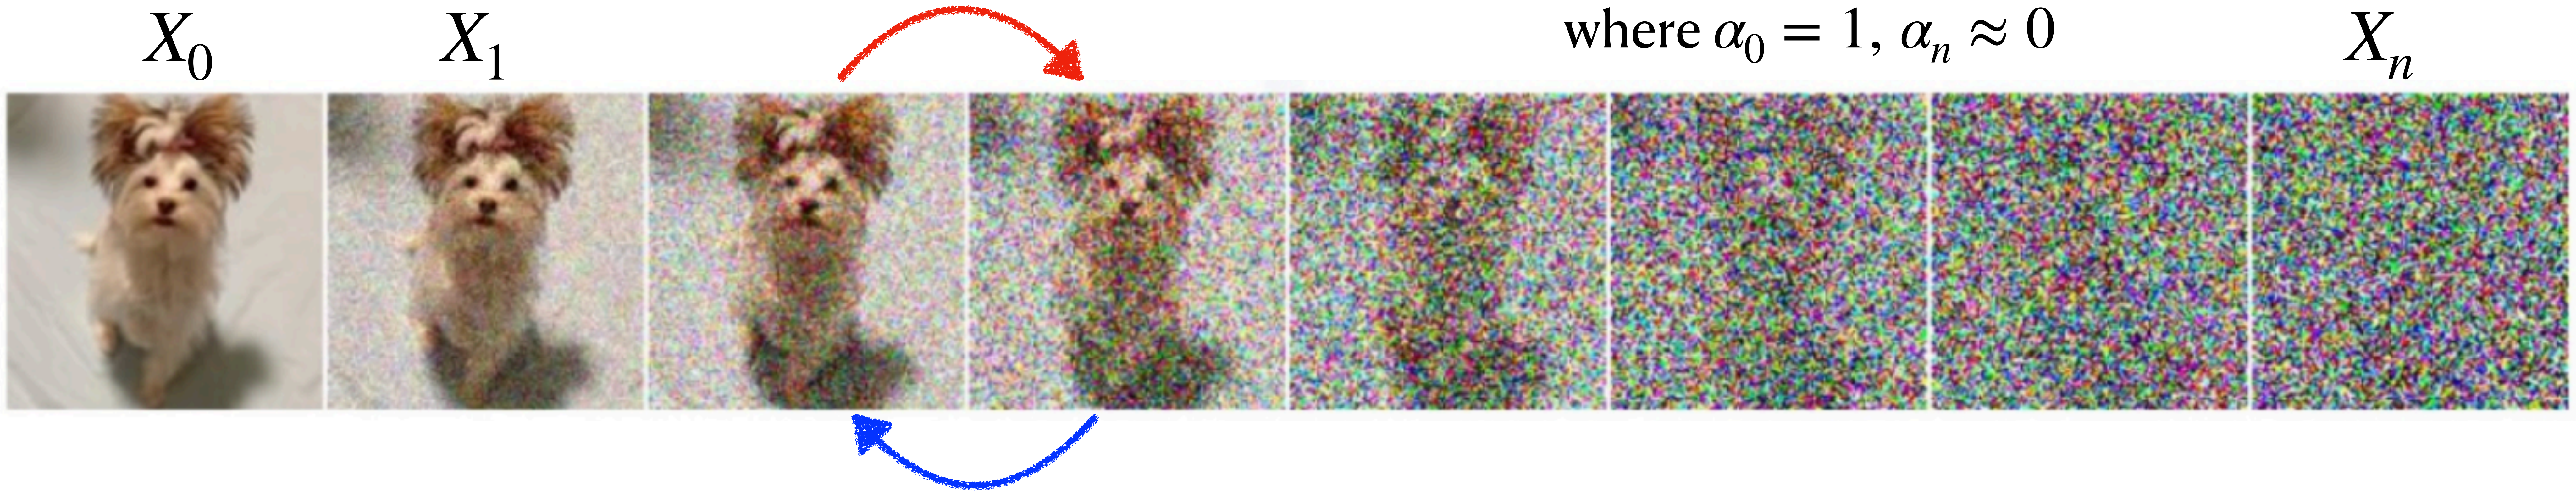

$\approx \mathcal{N}(0,\mathrm{I})$

$$p_{k|k+1} = ?$$

[Ho, J., Jain, A. and Abbeel, P., 2020. Denoising diffusion probabilistic models.]

Diffusion models

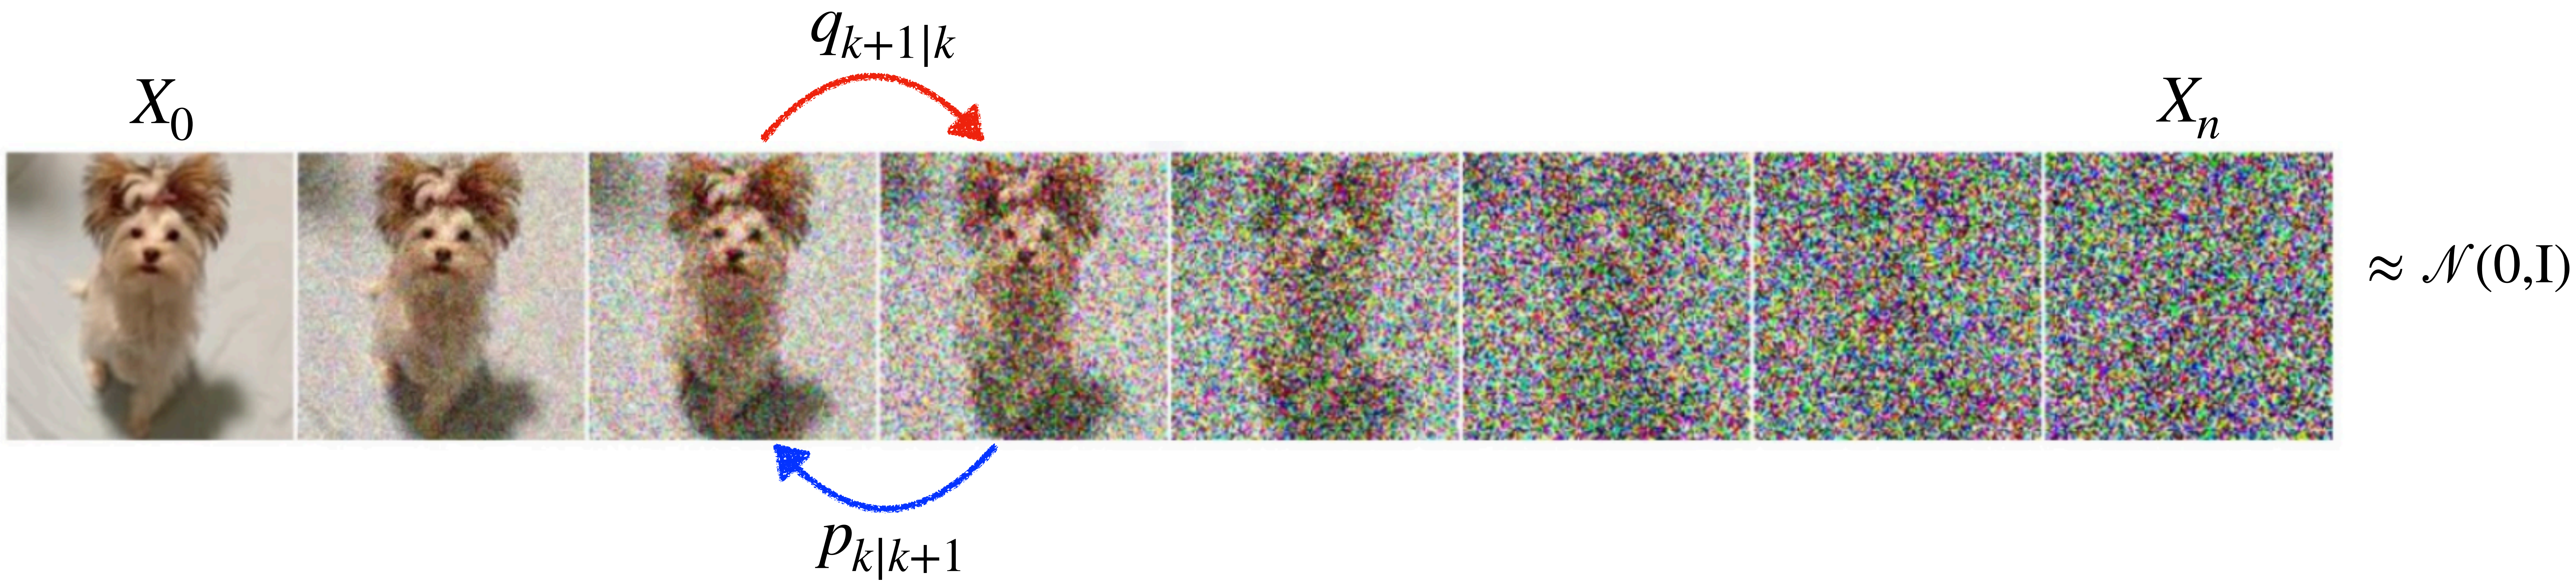

$$p_{0:n}(x_{0:n}) = p_0(x_0)\prod_{k=0}^{n-1} q_{k+1|k}(x_{k+1}\,|\,x_k)$$

$$= \mathrm{N}(x_{k+1}; (\alpha_{k+1}/\alpha_k)x_k, (1-\alpha_{k+1}^2/\alpha_k^2)\mathrm{I})$$

where $\alpha_0 = 1,\ \alpha_n \approx 0$

$$= p_n(x_n)\prod_{k=0}^{n-1} p_{k|k+1}(x_k\,|\,x_{k+1})$$

$p_n(x_n) \approx \mathcal{N}(0,\mathrm{I})$

$p_{k|k+1}(x_k\,|\,x_{k+1})$: intractable

[Ho, J., Jain, A. and Abbeel, P., 2020. Denoising diffusion probabilistic models.]

Diffusion models

$$p_{k|k+1}(x_k \mid x_{k+1}) = \int q_{k|0,k+1}(x_k \mid x_0, x_{k+1}) p_{0|k+1}(x_0 \mid x_{k+1}) \, \mathrm{d}x_0$$

$$\approx q_{k|0,k+1}(x_k \mid D^{\theta}_{k+1}(x_{k+1}), x_{k+1})$$

where D^{θ}_{k+1} is a « denoiser » with parameters θ, minimizing $\mathbb{E}[D^{\theta}_{k+1}(X_{k+1}) - X_0]$

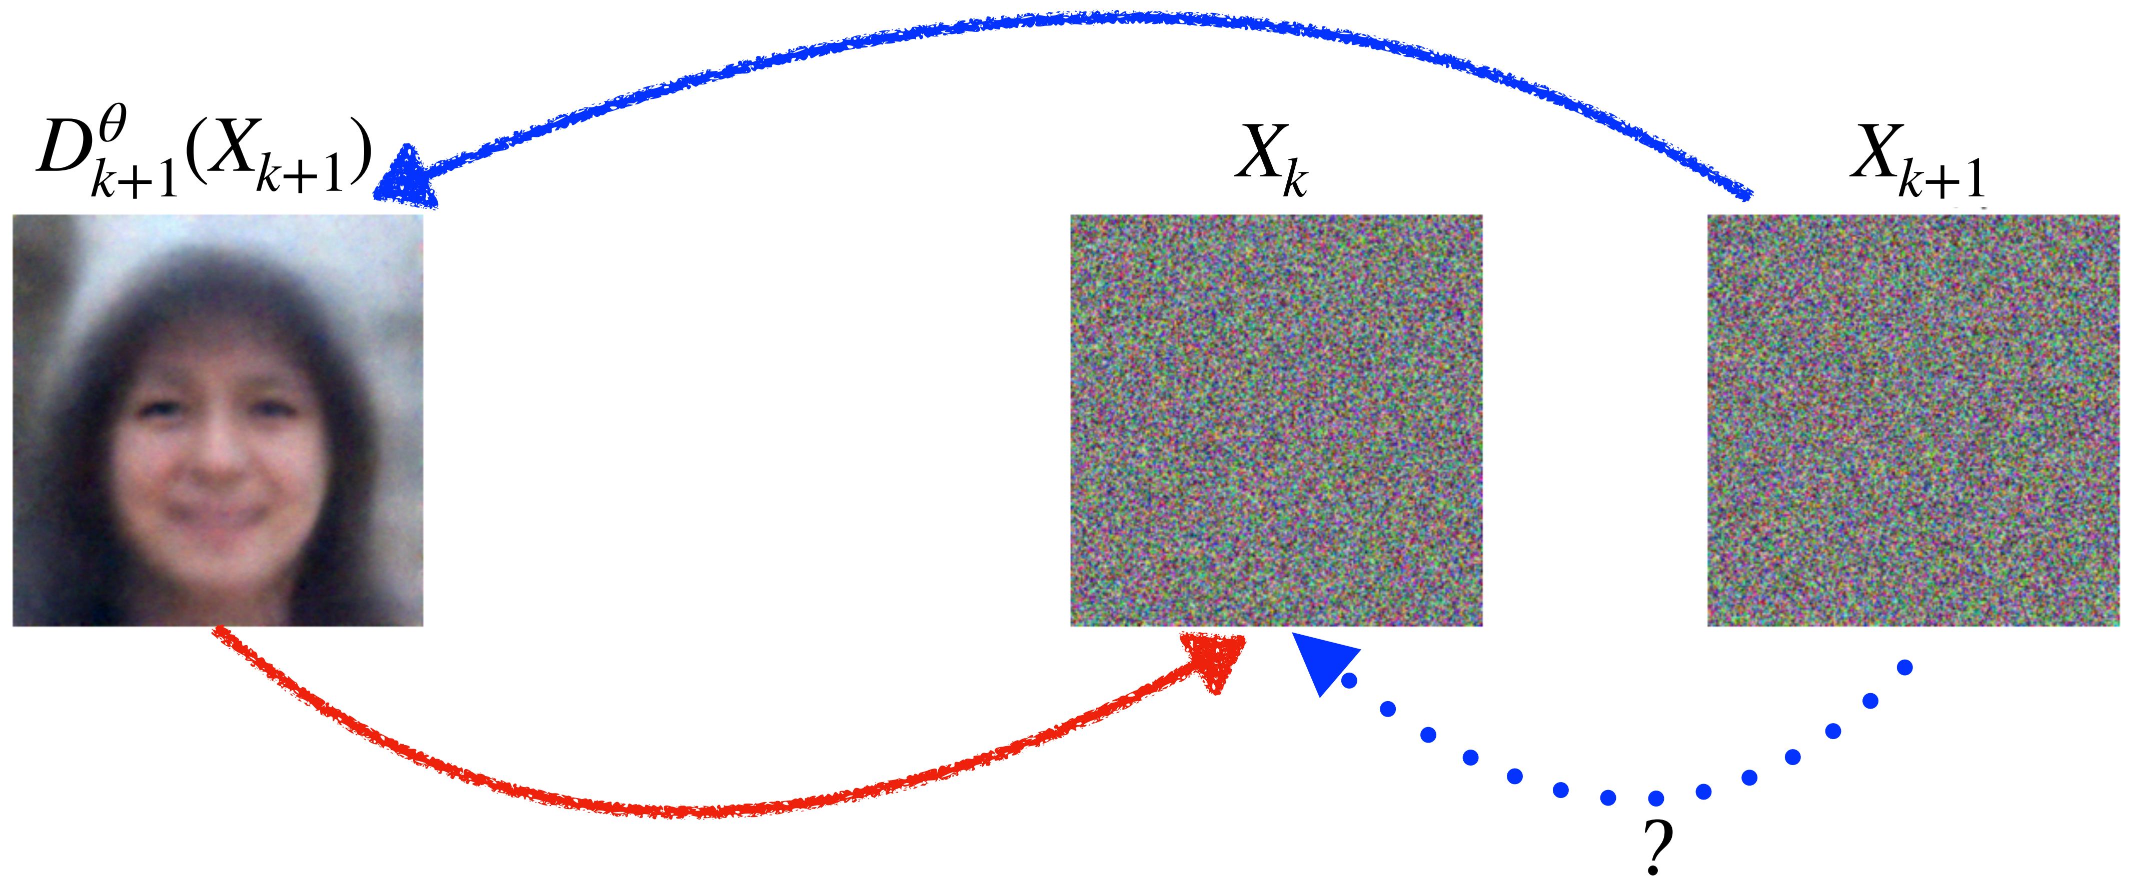

Diffusion models

$$p_{k|k+1}(x_k \mid x_{k+1}) = \int q_{k|0,k+1}(x_k \mid x_0, x_{k+1}) p_{0|k+1}(x_0 \mid x_{k+1})\, \mathrm{d}x_0$$

$$\approx q_{k|0,k+1}(x_k \mid D^\theta_{k+1}(x_{k+1}), x_{k+1})$$

where D^θ_{k+1} is a « denoiser » with parameters θ, minimizing $\mathbb{E}[D^\theta_{k+1}(X_{k+1}) - X_0]$

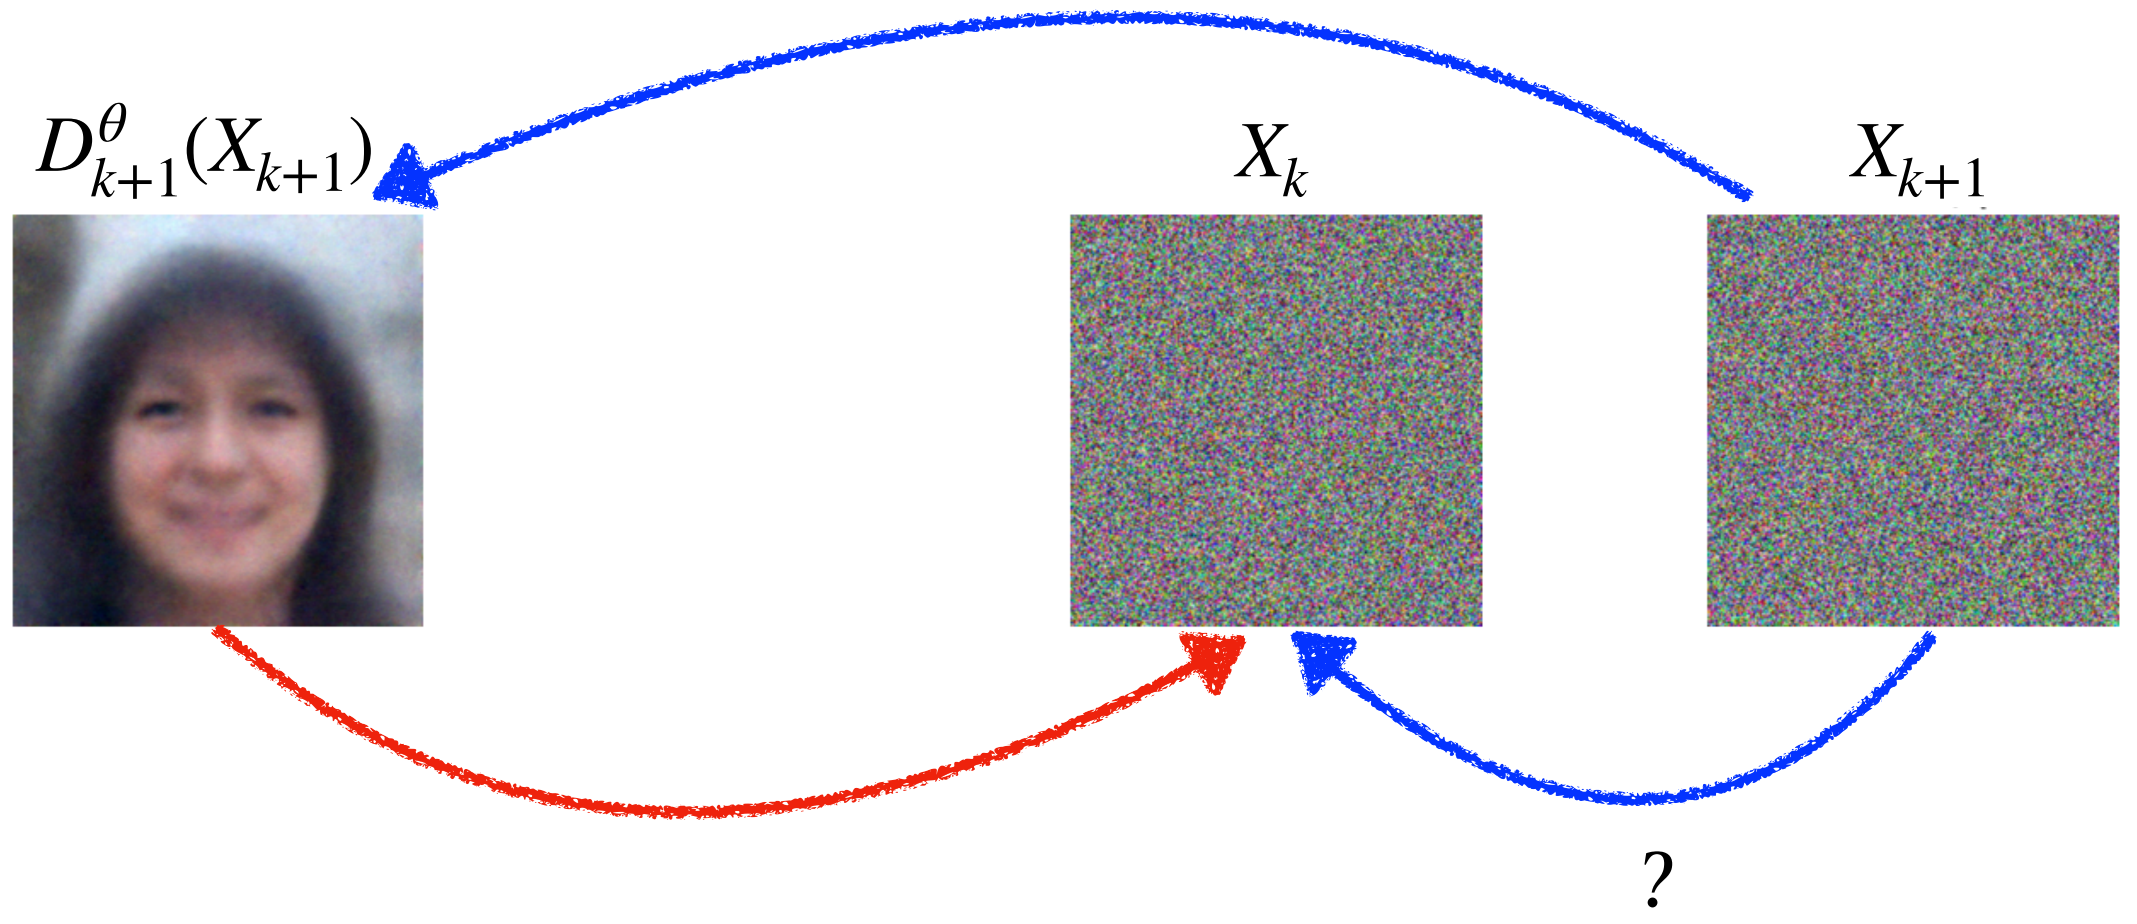

Bayesian Inverse Problems

Bayesian inverse problems

$$Y = A(X) + \sigma_y Z\,, \quad X \sim p_0$$

Given a realisation $Y = y$, sample the most plausible reconstructions X

$X \sim p_0$

$A(X) + \sigma_y Z$

Observation y

sample

Bayesian inverse problems

The reconstructions are encoded in the posterior distribution

$$\pi_0^y(x) \propto g_0(y \,|\, x)\, p_0(x)$$

where $g_0(y \,|\, x) = \mathcal{N}(y; A(x), \Sigma_y^2)$.

Sampling plausible reconstructions

$$\Longleftrightarrow$$

Drawing samples from π_0^y

Diffusion Posterior Sampling

Diffusion Posterior Sampling

$$\pi_0^y(x) \propto g_0(y \mid x)\, p_0(x)$$

Given a **pre-trained** diffusion model $p_0^\theta \approx p_0$, develop an efficient algorithm for sampling from π_0^y with no further model training

Distribution path

$X_0 \sim p_0$ $\cdots$ $X_k \sim p_k$ $\cdots$ $X_n \sim p_n$

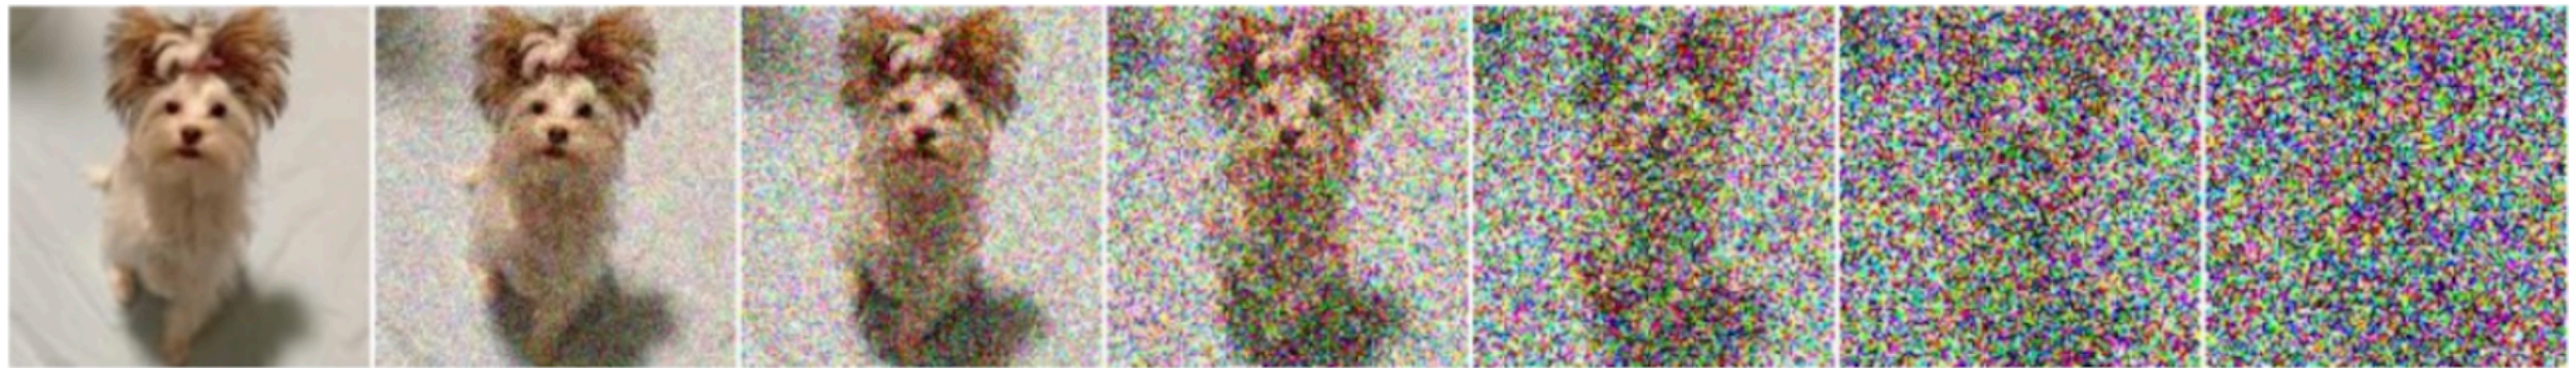

$$p_k(x_k) = \int p_0(x_0) q_{k|0}(x_k \mid x_0)\, \mathrm{d}x_0$$

Diffusion model for $\pi_0^y \iff$ follow path $(\pi_k)_{k=n}^0$ where

$$\pi_k^y(x_k) = \int \pi_0^y(x_0) q_{k|0}(x_k \mid x_0)\, \mathrm{d}x_0$$
$$\propto g_k(y \mid x_k) p_k(x_k)$$

Posterior denoiser

Define the likelihood of y given the noised sample x_k

$$g_k(y \mid x_k) = \int g_0(y \mid x_0) p_{0|k}(x_0 \mid x_k) \, \mathrm{d}x_0$$

We can relate the posterior denoiser D_k^y to the prior denoiser D_k

$$D_k^y(x_k) = D_k(x_k) + \alpha_k^{-1}(1 - \alpha_k^2) \nabla \log g_k(y \mid x_k)$$

$\approx D_k^\theta(x_k)$??

Posterior denoiser

$$D_k^y(x_k) = D_k(x_k) + \alpha_k^{-1}(1-\alpha_k^2)\nabla \log g_k(y\,|\,x_k)$$

$\approx D_k^\theta(x_k)$??

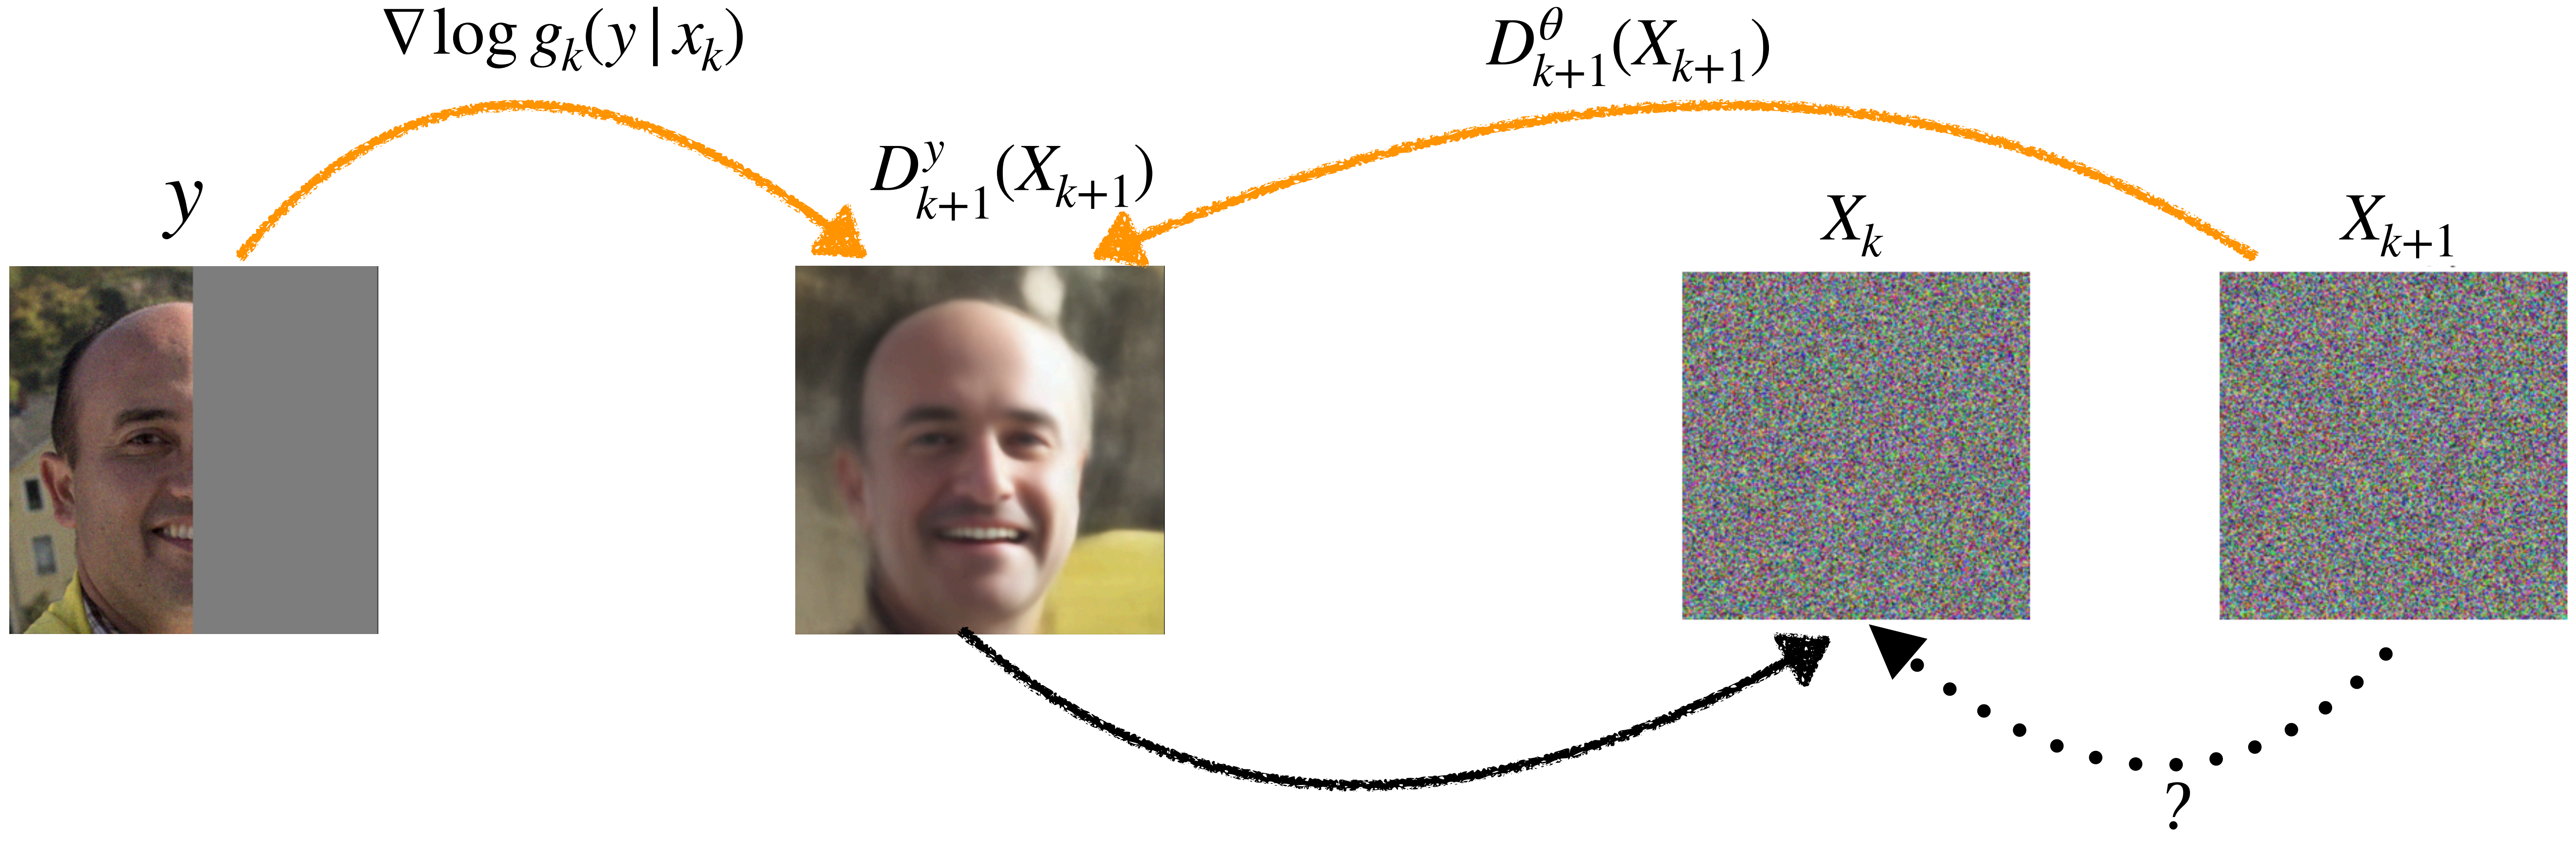

Posterior denoiser

$$D_k^y(x_k) = D_k(x_k) + \alpha_k^{-1}(1 - \alpha_k^2)\nabla \log g_k(y \mid x_k)$$

$$\approx D_k^\theta(x_k) \qquad ??$$

[Ho, J., Salimans, T., Gritsenko, A., Chan, W., Norouzi, M. and Fleet, D.J., 2022. Video diffusion models.]
[Chung, H., Kim, J., Mccann, M.T., Klasky, M.L. and Ye, J.C., 2022. Diffusion posterior sampling for general noisy inverse problems.]

[Song, J., Vahdat, A., Mardani, M. and Kautz, J., 2023, May. Pseudoinverse-guided diffusion models for inverse problems.]
[Boys, B., Girolami, M., Pidstrigach, J., Reich, S., Mosca, A. and Akyildiz, O.D., 2023. Tweedie moment projected diffusions for inverse problems.]

[Moufad, B., Janati, Y., Bedin, L., Durmus, A., Douc, R., Moulines, E., Olsson, J., 2025. Diffusion posterior sampling with midpoint guidance.]

Applications to ECGs

Denoiser for ECGs D_k^θ

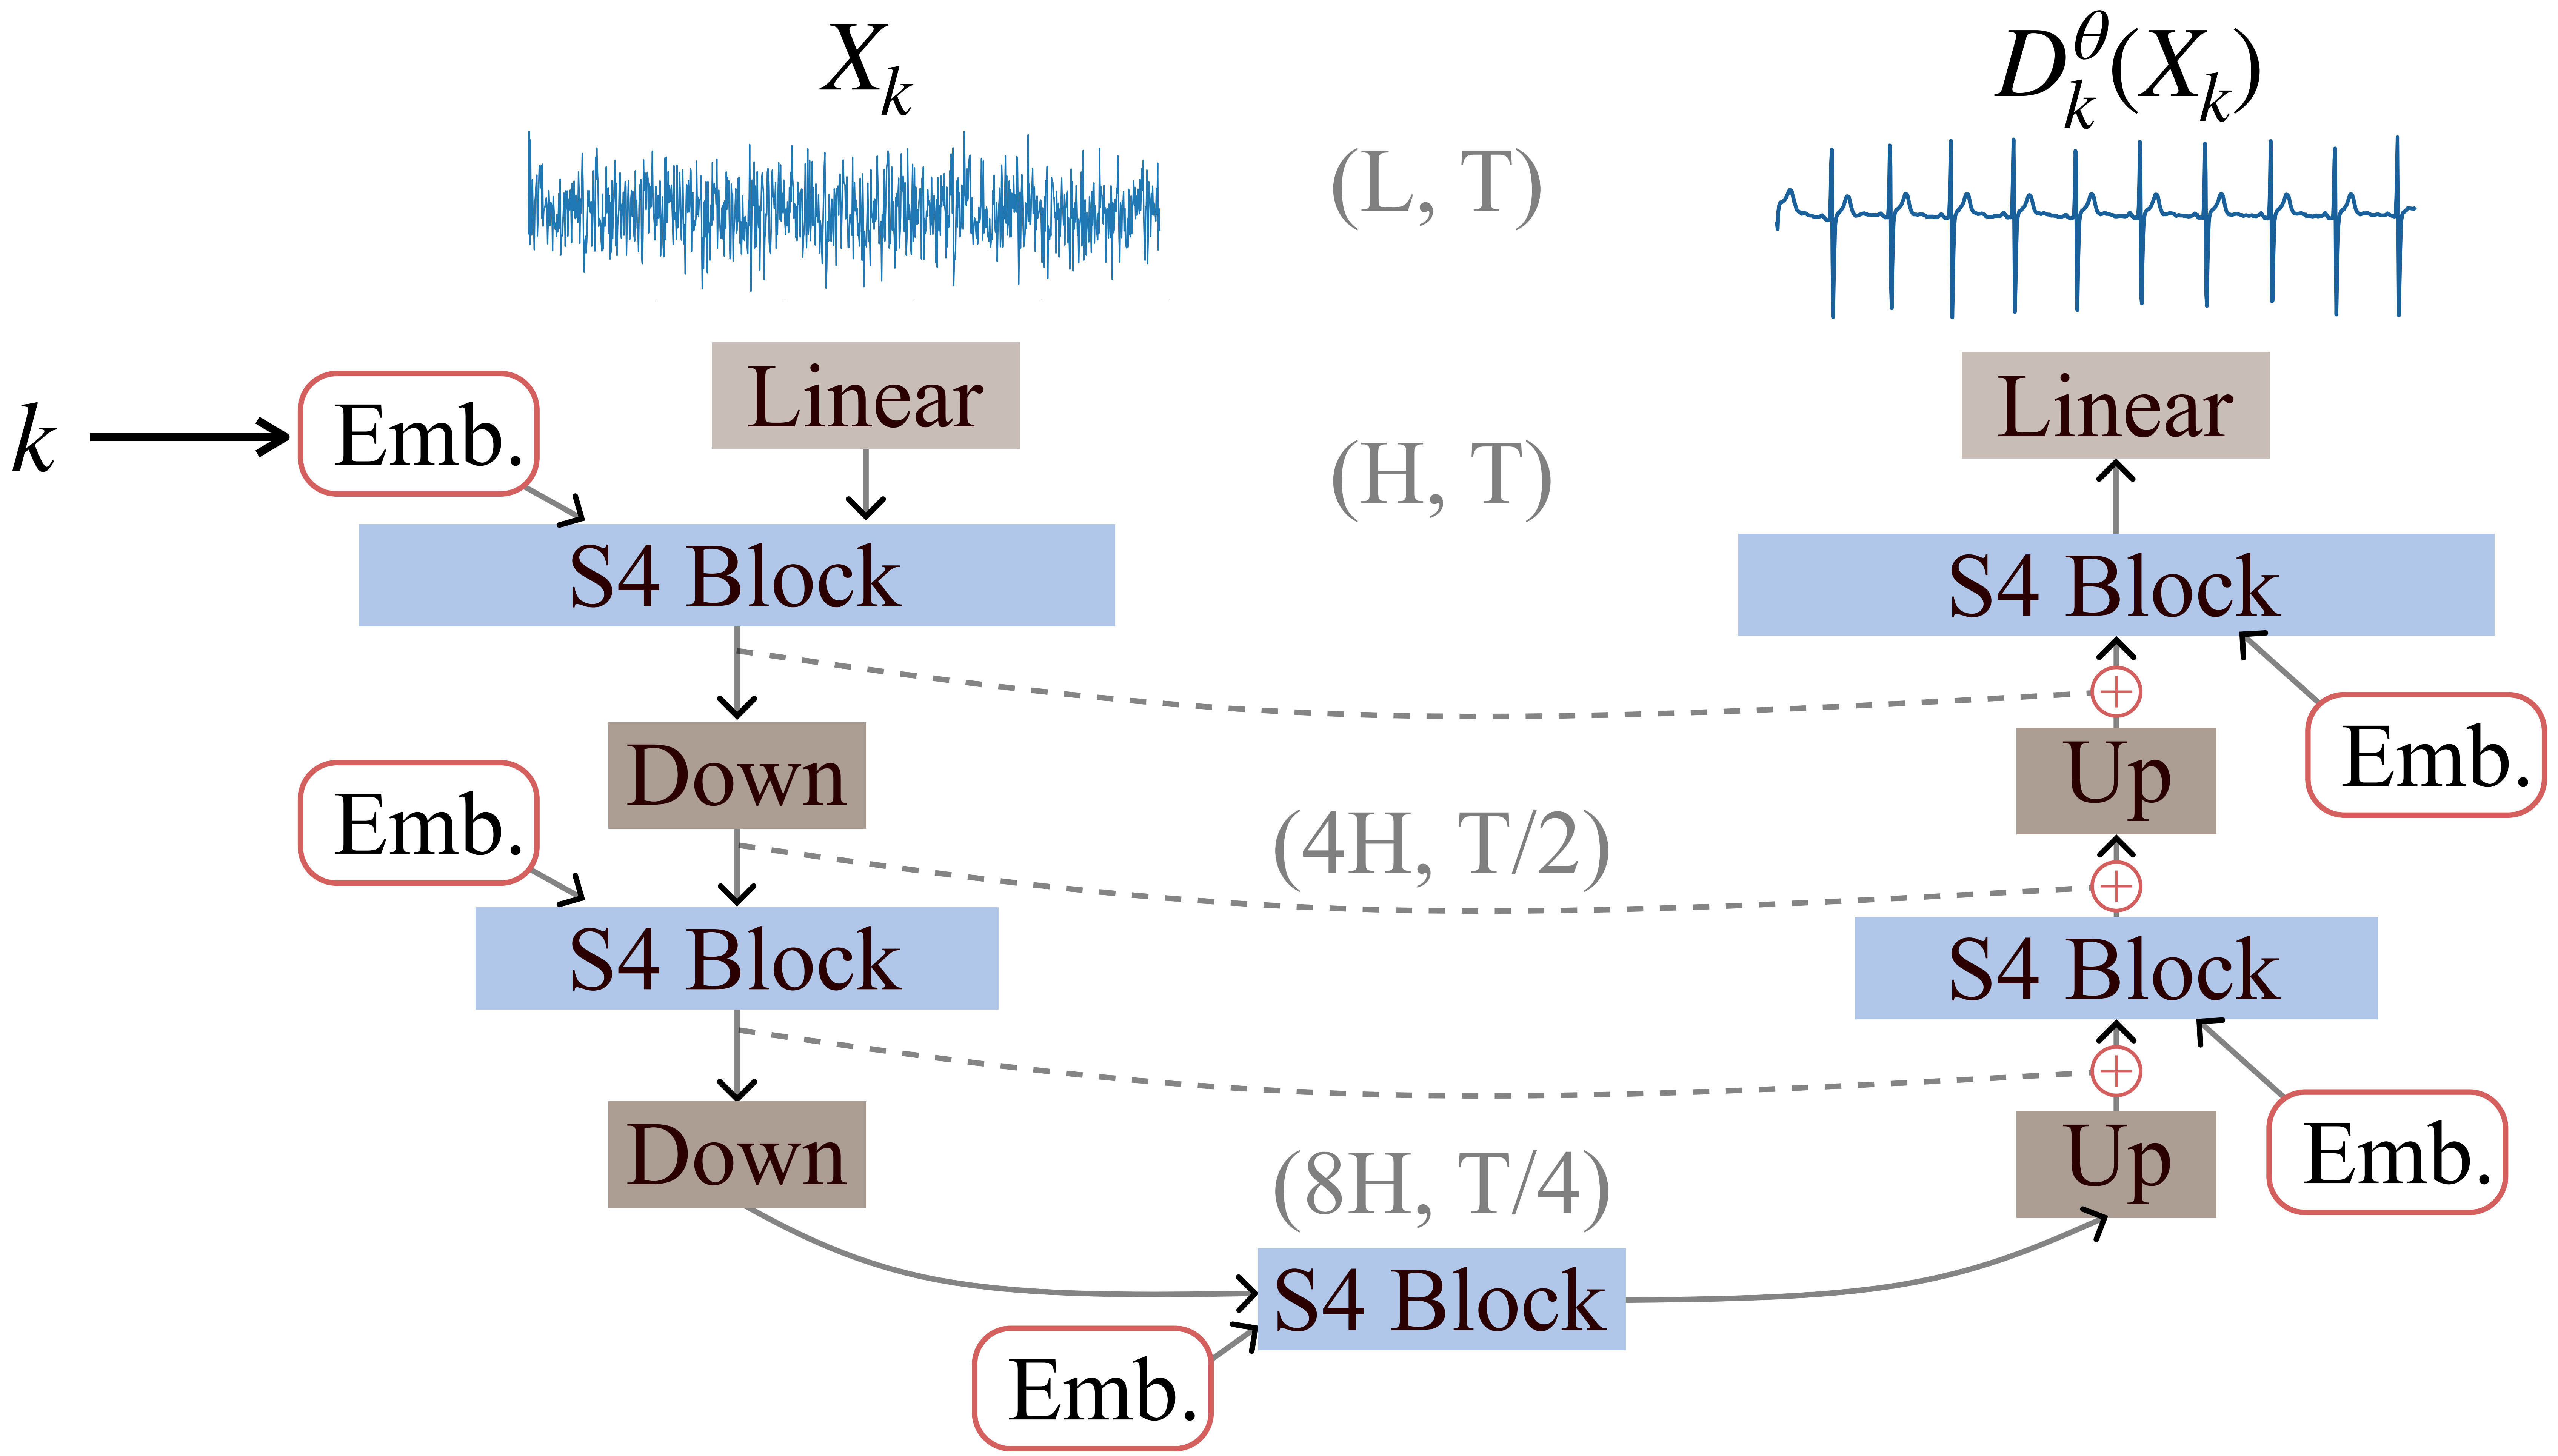

Reconstruction of Missing Leads

Bundle Block Branch

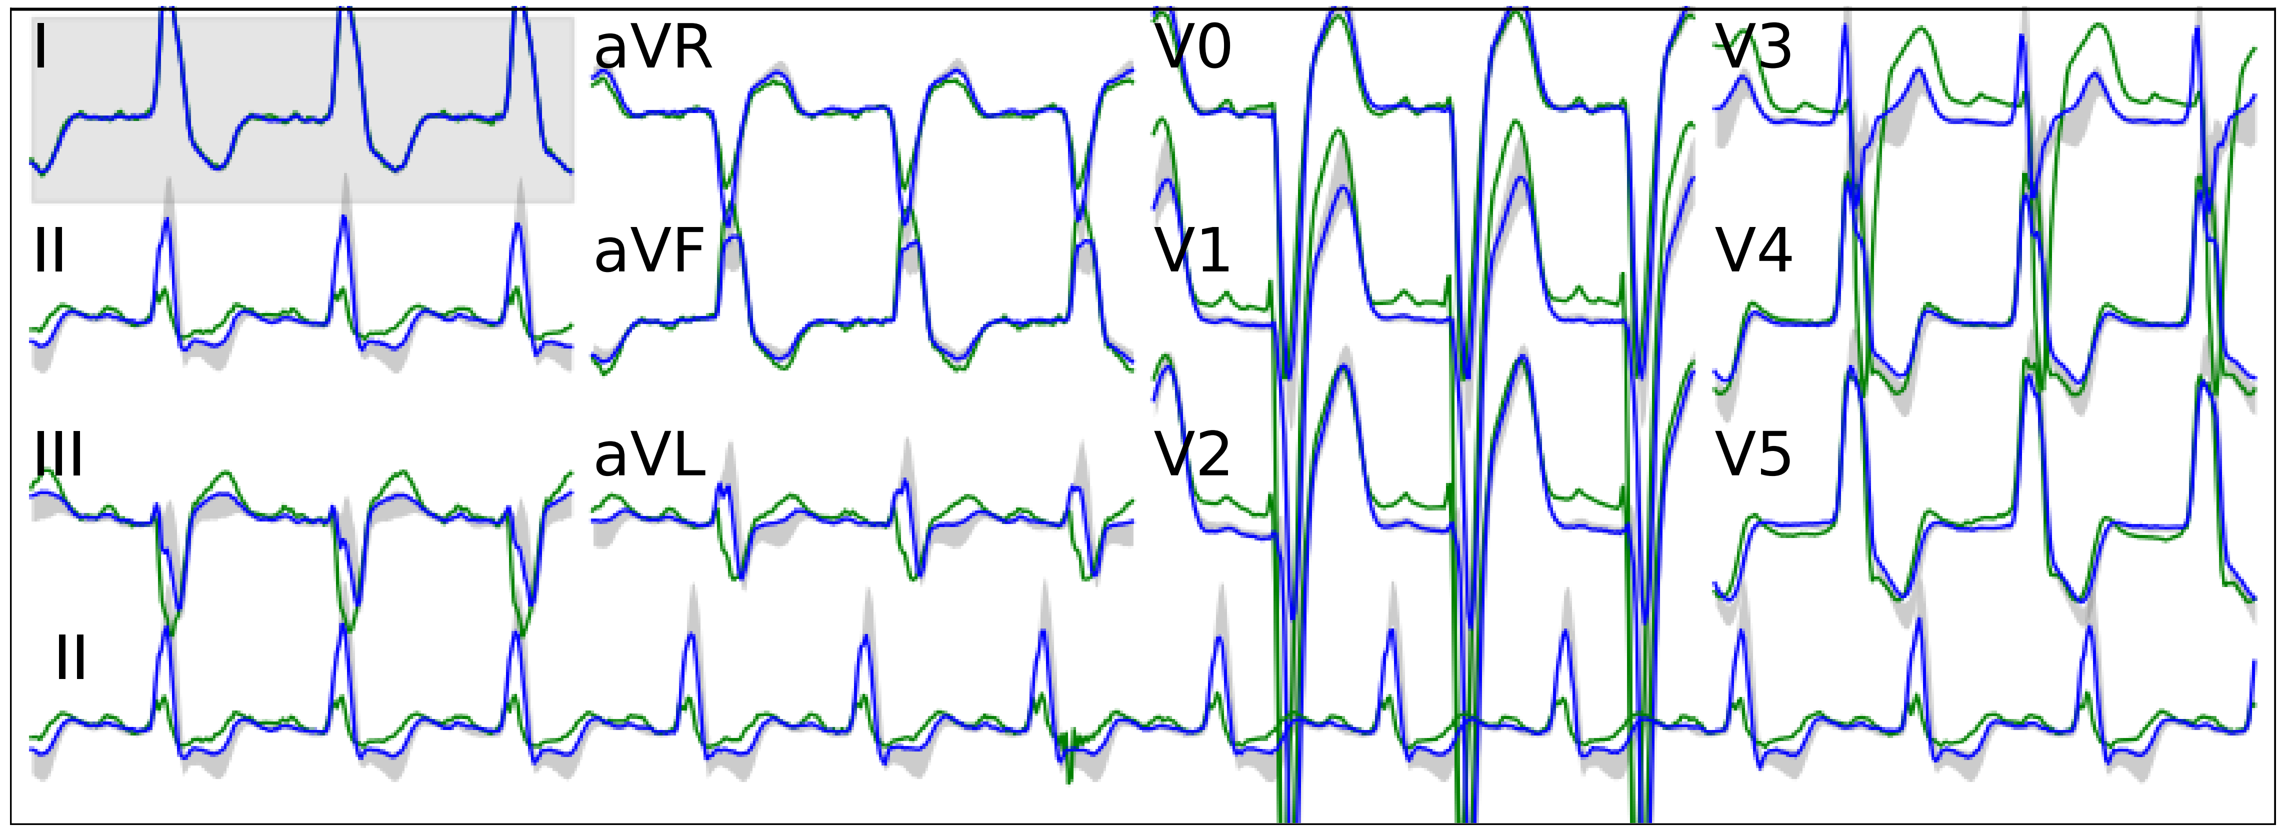

Atrial Fibrillation

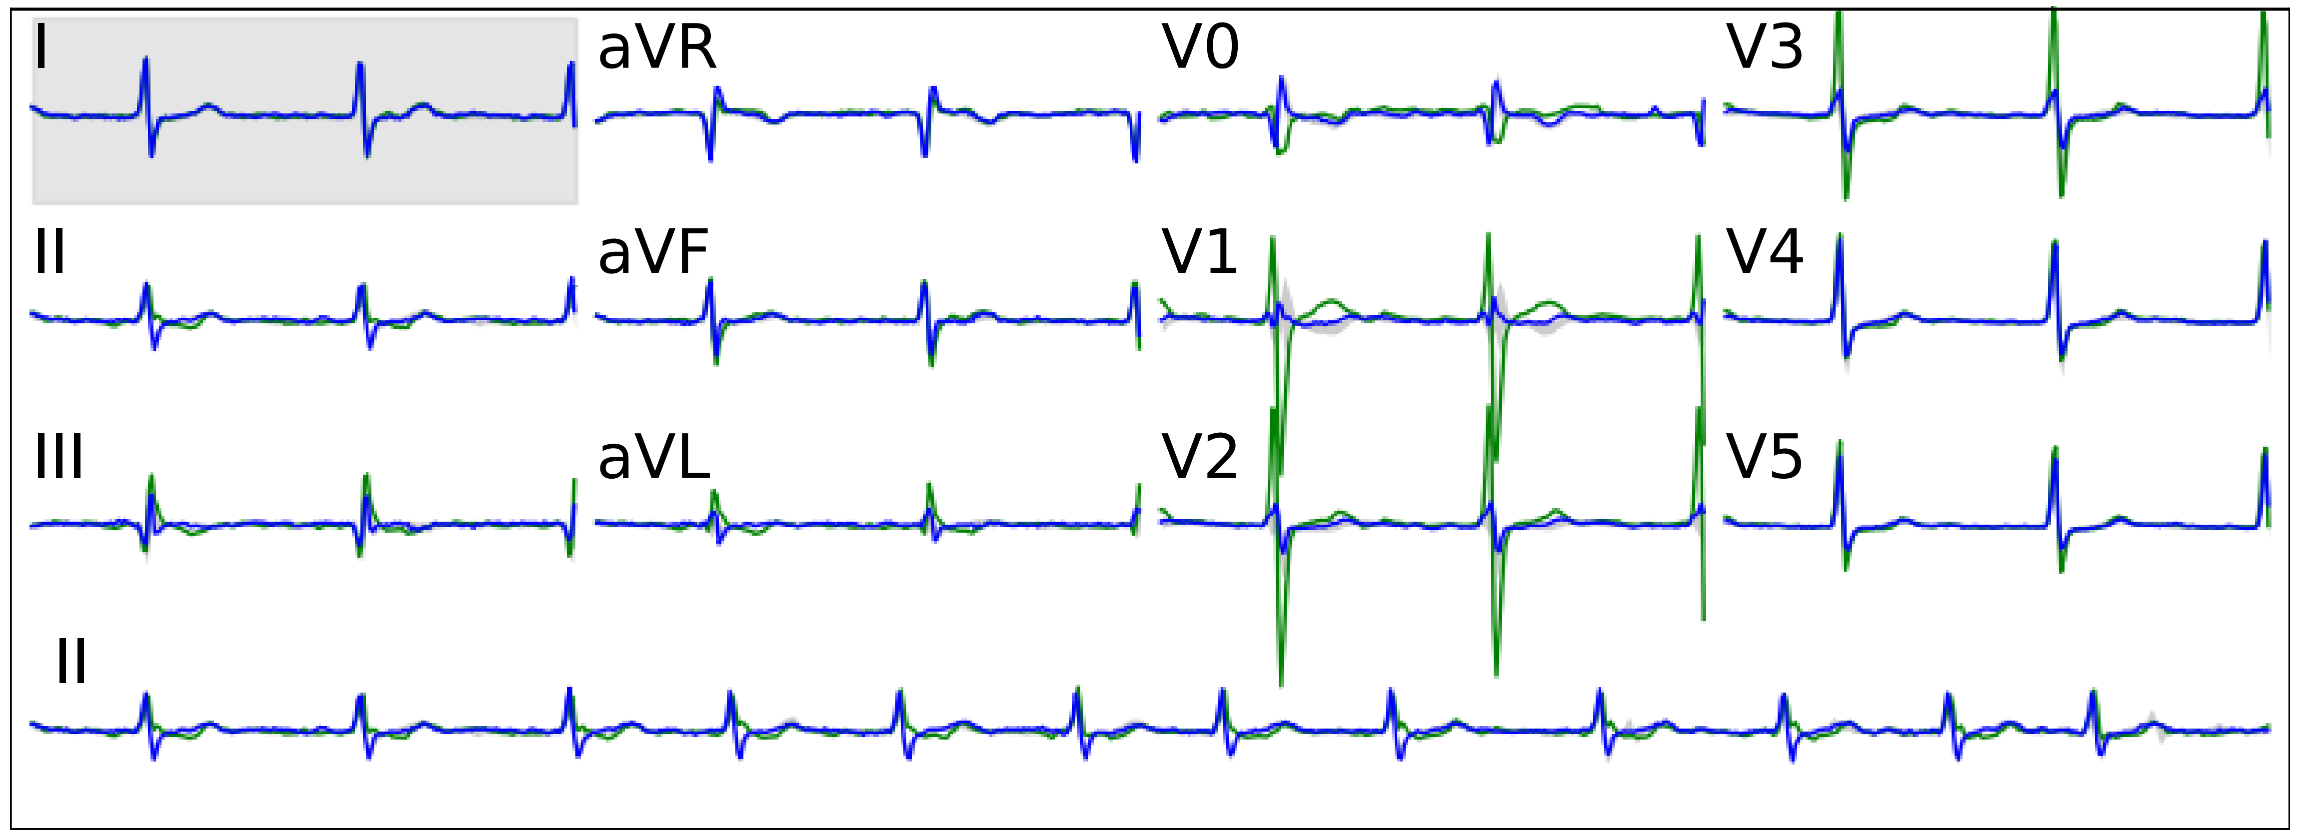

		Errors			XResNet1d50 AUC-PR (↑)		
Method	Task	RMSE (↓)	MAD (↓)	DTW (↓)	NSR	BBB	AF
Ground-truth	-	-	-	-	1.00	1.00	0.98
RhythmDiff-MGPS	MLR	**0.19** $\pm$ 0.01	**0.09** $\pm$ 0.00	**0.29** $\pm$ 0.02	**0.99**	<u>0.91</u>	**0.96**
RhythmDiff-DPS	MLR	1.61 $\pm$ 0.17	0.16 $\pm$ 0.01	77.35 $\pm$ 15.34	0.97	0.65	0.68
RhythmDiff-PGDM	MLR	0.19 $\pm$ 0.01	0.10 $\pm$ 0.003	0.32 $\pm$ 0.03	0.99	0.90	<u>0.89</u>
RhythmDiff-DDNM	MLR	0.19 $\pm$ 0.01	0.10 $\pm$ 0.003	0.36 $\pm$ 0.03	0.99	**0.92**	0.81
RhythmDiff-DiffPIR	MLR	0.21 $\pm$ 0.01	0.11 $\pm$ 0.003	0.39 $\pm$ 0.03	0.94	0.71	0.57
RhythmDiff-RedDiff	MLR	0.20 $\pm$ 0.01	0.11 $\pm$ 0.003	0.36 $\pm$ 0.02	0.97	0.87	0.75
ECGRecover [26]	MLR	0.34 $\pm$ 0.005	0.20 $\pm$ 0.003	1.10 $\pm$ 0.04	0.98	0.73	0.61

Artefact Removal in ECGs

$$Y = X + \sum_{j=1}^{J} \mu_j \phi_j + \sigma_y Z\,, \quad X \sim p_0$$

Given J Fourier harmonics $\phi_{j \in [1,J]}$ modeling additive artifacts, and a realization $Y = y$

1) propose weights $\mu_{j \in [1,J]}$

2) sample most plausible reconstruction X

3) regress new $\mu_{j \in [1,J]}$

... iterate until convergences

Artefact Removal in ECGs

Bundle Block Branch

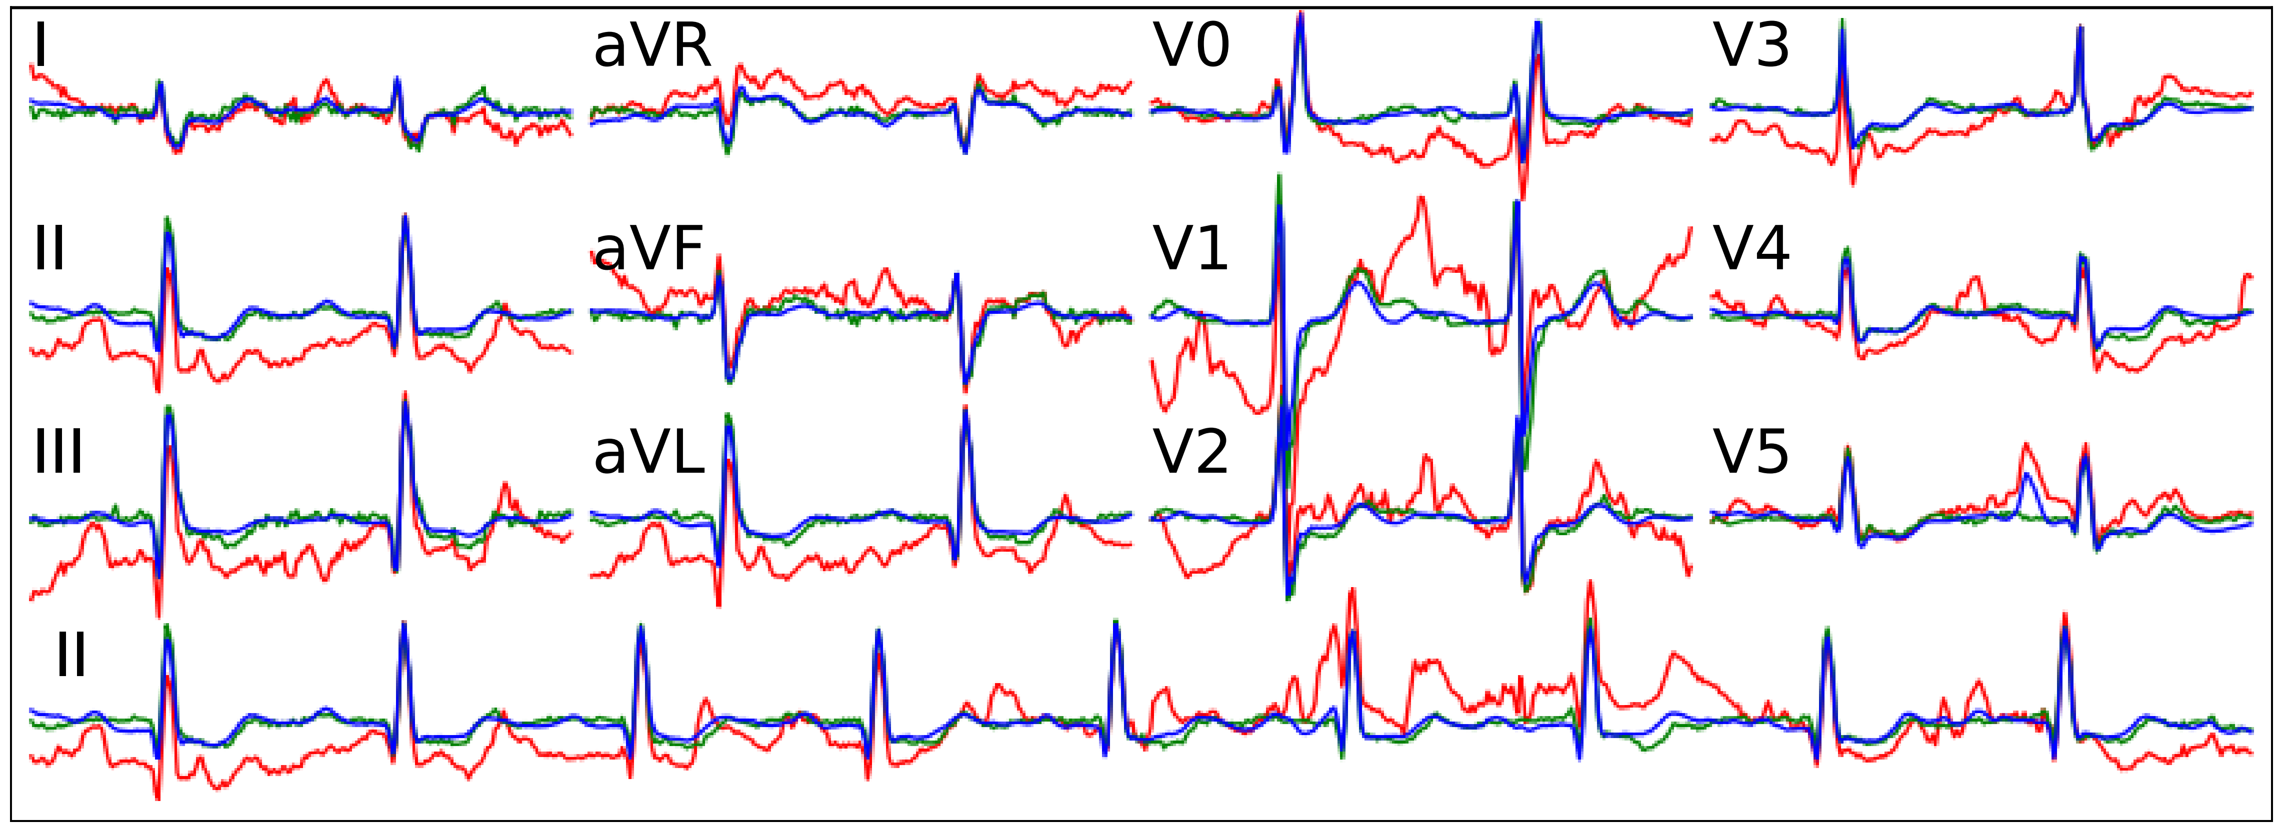

Atrial Fibrillation

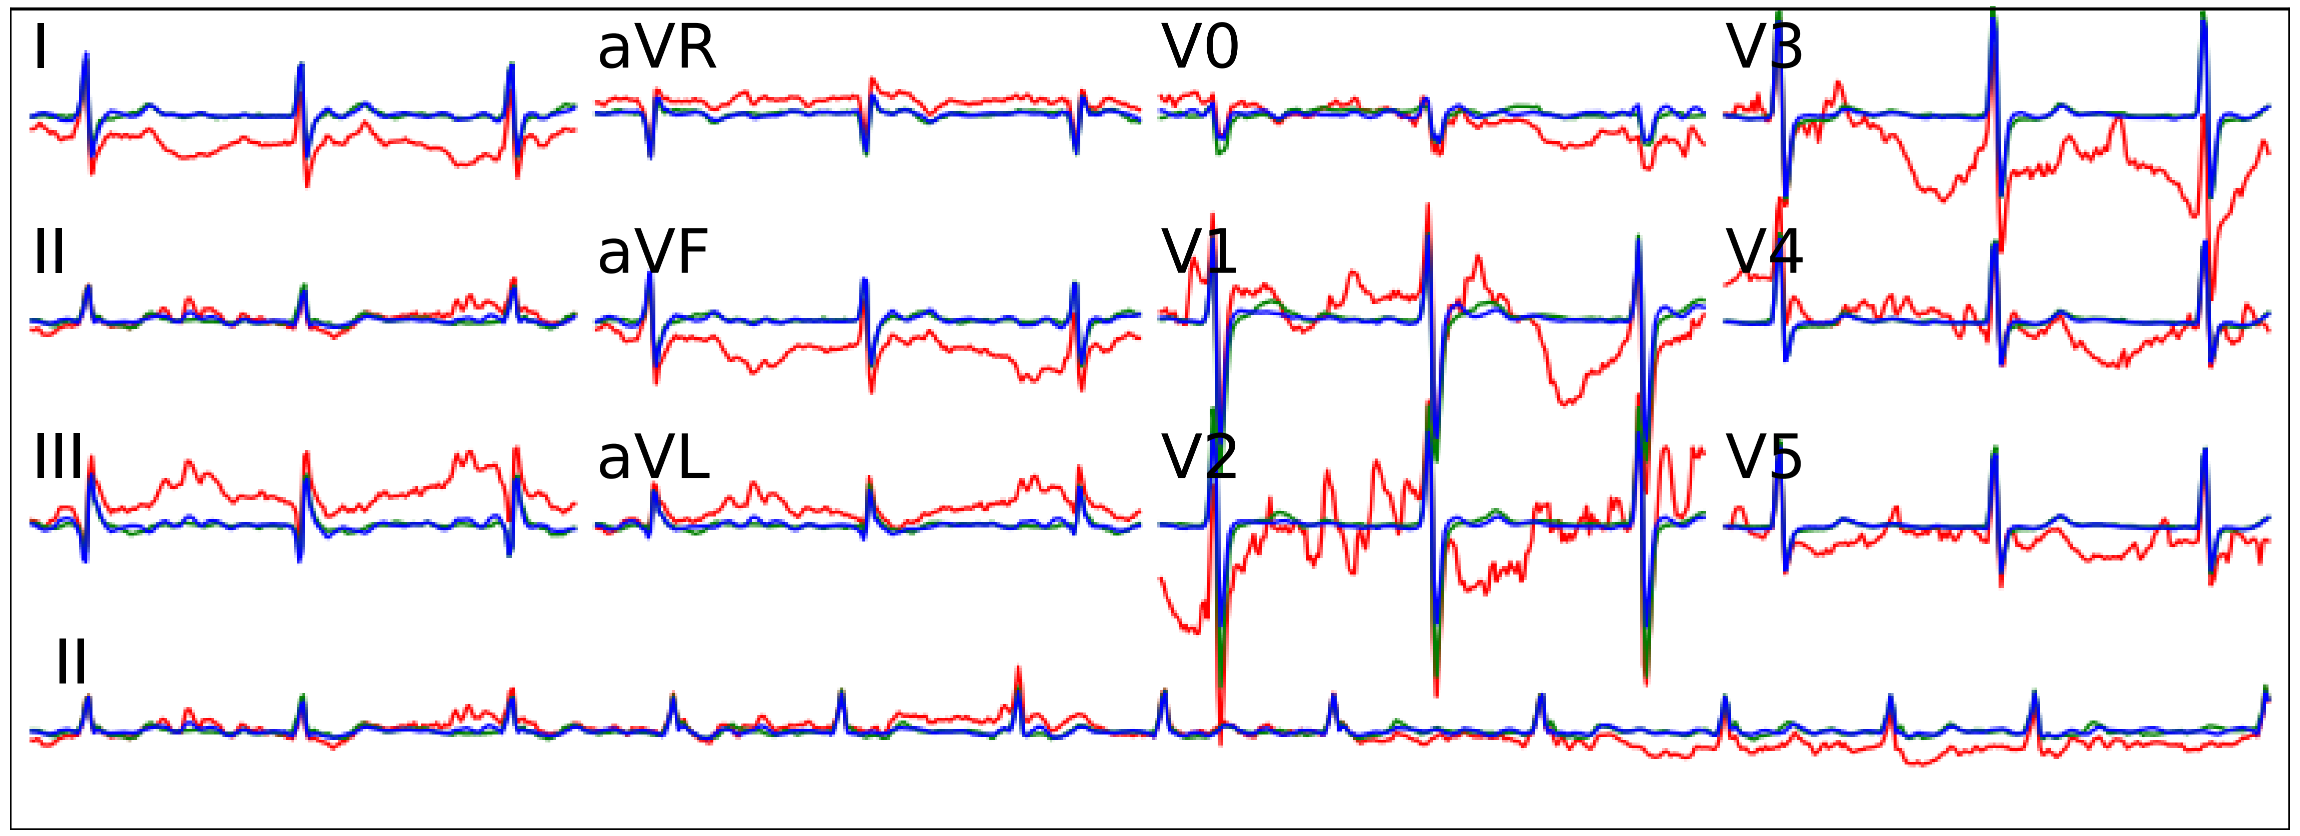

Method	Task	Errors RMSE (↓)	MAD (↓)	DTW (↓)	XResNet1d50 AUC-PR (↑) NSR	BBB	AF
Ground-truth	-	-	-	-	1.00	1.00	0.98
RhythmDiff-MGPS	EMR	**0.097** $\pm$ 0.002	**0.059** $\pm$ 0.001	**0.13** $\pm$ 0.01	**1.00**	**1.00**	**0.86**
DeScoD [29]	EMR	0.100 $\pm$ 0.002	0.074 $\pm$ 0.001	0.13 $\pm$ 0.01	1.00	0.98	0.81
RhythmDiff-MGPS	BWR	0.134 $\pm$ 0.003	0.078 $\pm$ 0.001	0.24 $\pm$ 0.02	0.99	0.97	0.86
DeScoD [29]	BWR	0.088 $\pm$ 0.002	0.062 $\pm$ 0.001	0.11 $\pm$ 0.01	0.99	0.99	0.88

Conclusion

A single trained diffusion model for multiple downstream task

[Bedin, L., Cardoso, G., Duchateau, J., Dubois, R., Moulines, E., 2024. Leveraging an ECG beat diffusion model for morphological reconstruction from indirect signals.]

[Moufad, B., Janati, Y., Bedin, L., Durmus, A., Douc, R., Moulines, E., Olsson, J., 2025. Diffusion posterior sampling with midpoint guidance.]

Midpoint Guidance Posterior Sampling

Posterior denoiser approximation

$$\int g_0(y \,|\, x_0) p_{0|k}(x_0 \,|\, x_k)\, \mathrm{d}x_0 \approx g_0(y \,|\, D_k(x_k))$$

implicitely assumes that $p_{0|k}(\,\cdot\,|\,x_k) \approx \delta_{D_k(x_k)}$

$$D_k^y(x_k) \approx D_k(x_k) + \alpha_k^{-1}(1 - \alpha_k^2)\,\nabla_{x_k} \log g_0\big(y \,|\, D_k(x_k)\big)$$

- Samples diverge after a few iterations
- Instead, [Chung et al. 2023] plugs $\dfrac{C_k}{\|y - A(D_k(x_k))\|} \nabla_{x_k} \log g_0\big(y \,|\, D_k(x_k)\big)$
- $\sim 50\,\%$ more expensive in terms of memory and runtime:

$$\nabla_{x_k} \log g_0(y \,|\, D_k(x_k)) = \nabla_{x_k} D_k(x_k)^\intercal\, \nabla_{x_0} \log g_0(y \,|\, x_0)_{|x_0 = D_k(x_k)}$$

[Ho, J., Salimans, T., Gritsenko, A., Chan, W., Norouzi, M. and Fleet, D.J., 2022. Video diffusion models.]
[Chung, H., Kim, J., Mccann, M.T., Klasky, M.L. and Ye, J.C., 2022. Diffusion posterior sampling for general noisy inverse problems.]

Midpoint guidance

- Usual improvement: Gaussian approximation of $p_{0|k}(\cdot \,|\, x_k)$, but too expensive
- As $\ell \to 0$, $g_\ell(y \,|\, x_\ell) \approx g_0(y \,|\, D_\ell(x_\ell))$ is a more reasonable approximation
- How can it be leveraged at step k of the denoising process?

$$\text{Step } k \implies \text{sample} \quad q_{k|0,k+1}(\cdot \,|\, \hat{D}^y_k(x_{k+1}), x_{k+1})$$

[Song, J., Vahdat, A., Mardani, M. and Kautz, J., 2023, May. Pseudoinverse-guided diffusion models for inverse problems.]
[Boys, B., Girolami, M., Pidstrigach, J., Reich, S., Mosca, A. and Akyildiz, O.D., 2023. Tweedie moment projected diffusions for inverse problems.]

Posterior backward chain

$$\pi^y_{0:n}(x_{0:n}) = \pi^y_0(x_0)\prod_{k=0}^{n-1} q_{k+1|k}(x_{k+1}\,|\,x_k)$$

$$= \pi^y_n(x_n)\prod_{k=0}^{n-1} \pi^y_{k|k+1}(x_k\,|\,x_{k+1})$$

$\approx \mathcal{N}(0,\mathrm{I})$ to be approximated

where $$\pi^y_{k|k+1}(x_k\,|\,x_{k+1}) = \pi^y_k(x_k)q_{k+1|k}(x_{k+1}\,|\,x_k)\big/\pi^y_{k+1}(x_{k+1})$$

$$= g_k(y\,|\,x_k)p_{k|k+1}(x_k\,|\,x_{k+1})\big/g_{k+1}(y\,|\,x_{k+1})$$

not very useful

Midpoint decomposition

For all $\ell \in [0:k]$

$$\pi^y_{k|k+1}(x_k \mid x_{k+1}) = \int q_{k|\ell,k+1}(x_k \mid x_\ell, x_{k+1}) \pi^y_{\ell|k+1}(x_\ell \mid x_{k+1}) \, \mathrm{d}x_\ell$$

Gaussian
+ easy to sample from

where $\quad \pi^y_{\ell|k+1}(x_\ell \mid x_{k+1}) \propto g_\ell(y \mid x_\ell) \, p_{\ell|k+1}(x_\ell \mid x_{k+1})$

prior backward transition

- As $\ell \to 0$, $g_\ell(y \mid x_\ell) \approx g_0(y \mid D_\ell(x_\ell))$ is a more reasonable approximation
- But, $p_{\ell|k+1}(x_\ell \mid x_{k+1}) = \int \prod_{j=\ell}^{k} p_{j|j+1}(x_j \mid x_{j+1}) \, \mathrm{d}x_{\ell+1:k}$: more multi-modal as $\ell \to 0$

There is a tradeoff!

Midpoint decomposition

Conditional densities involved in $\pi^y_{k|k+1}(\cdot \mid x_{k+1})$ for different choices of ℓ

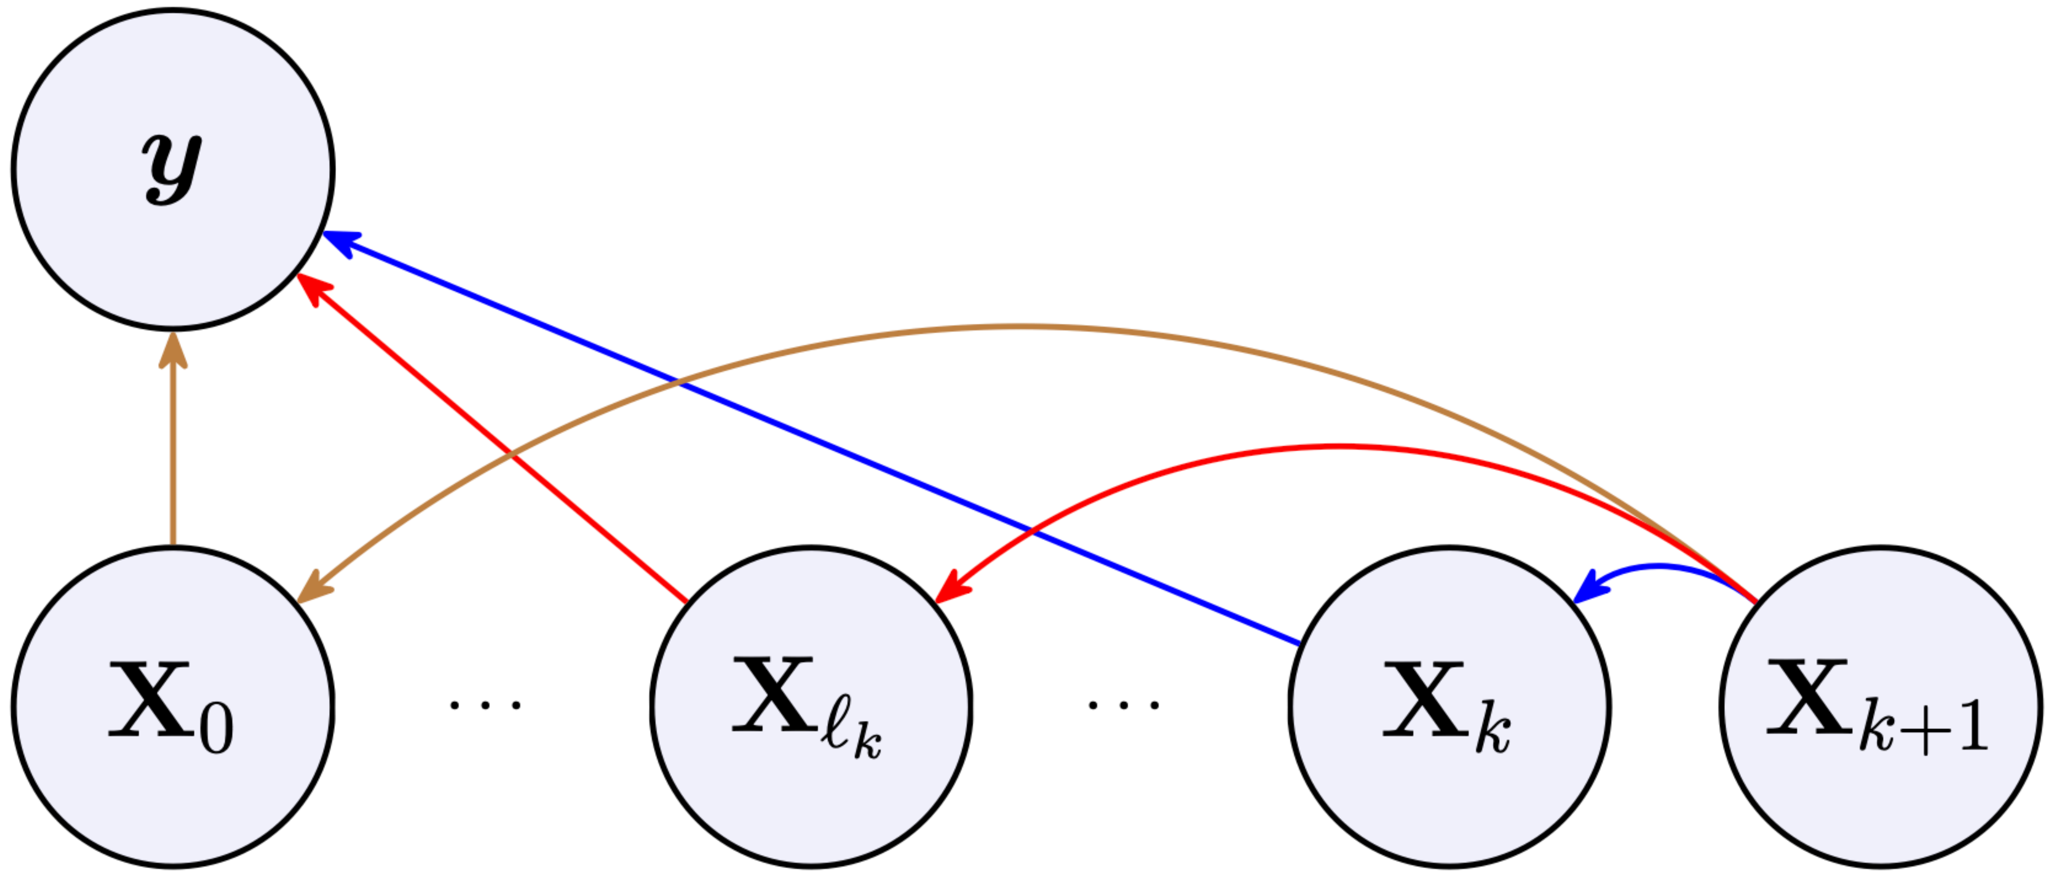

Long arrow $\iff$ more difficult to approximate

Surrogate model

Let $(\ell_k)_{k=1}^n \in [1:n]^n$ with $\ell_k \le k$.

$$\pi^y_{\ell_k|k+1}(x_{\ell_k} | x_{k+1}) \propto g_{\ell_k}(y | x_{\ell_k}) \, p_{\ell_k|k+1}(x_{\ell_k} | x_{k+1})$$

For $p_{\ell_k|k+1}(\cdot | x_{k+1})$ we use the Gaussian approximation:

$$p^\theta_{\ell_k|k+1}(x_{\ell_k} | x_{k+1}) = q_{\ell_k|0,k+1}(x_{\ell_k} | D^\theta_{k+1}(x_{k+1}), x_{k+1})$$

Surrogate backward posterior transition:

$$\hat{\pi}^\theta_{\ell_k|k+1}(x_{\ell_k} | x_{k+1}) \propto g_0\big(y | D^\theta_{\ell_k}(x_{\ell_k})\big) p^\theta_{\ell_k|k+1}(x_{\ell_k} | x_{k+1})$$

Surrogate model

Our approximation of π is

$$\hat{\pi}_0^{\ell}(x_0) = \int \mathrm{N}(x_n; 0_d, I_d) \prod_{k=0}^{n-1} \hat{\pi}^{\ell}_{k|k+1}(x_k \,|\, x_{k+1}) \, \mathrm{d}x_{1:n}$$

where $$\hat{\pi}^{\ell}_{k|k+1}(x_k \,|\, x_{k+1}) = \int \underbrace{q_{k|\ell_k,k+1}(x_k \,|\, x_{\ell_k}, x_{k+1})}_{\text{straightforward to sample}} \underbrace{\hat{\pi}^{\theta}_{\ell_k|k+1}(x_{\ell_k} \,|\, x_{k+1})}_{\text{approximate inference}} \, \mathrm{d}x_{\ell_k}$$

$$\pi^{\theta}_{\ell_k|k+1}(x_{\ell_k} \,|\, x_{k+1}) \propto g_0\big(y \,|\, D^{\theta}_{\ell_k}(x_{\ell_k})\big) p^{\theta}_{\ell_k|k+1}(x_{\ell_k} \,|\, x_{k+1})$$