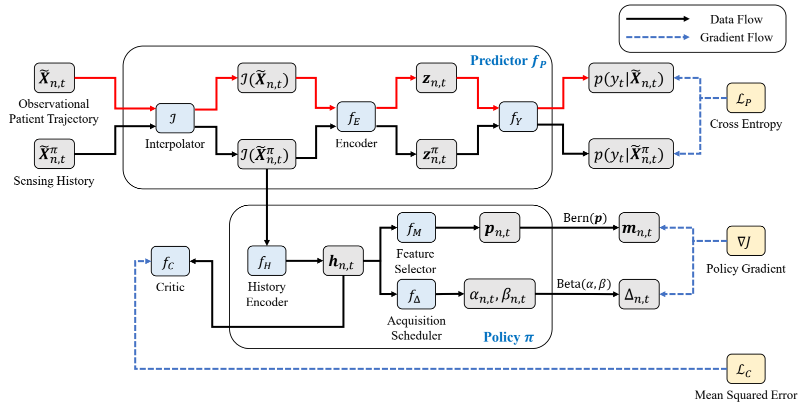

Figure 1: A schematic illustration of RAS.

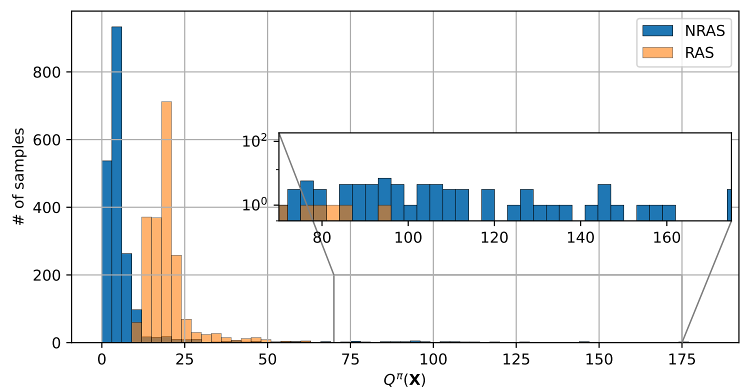

Figure 2: Comparison of cumulative cost distribution on the synthetic dataset. Our proposed sensing approach RAS is trained to optimize sensing performance for tail-risk samples at the last $\alpha = 0.1$ quantile. NRAS is the non-risk-averse version of RAS. RAS can effectively reduce the long tail in the $Q^{\pi}(\mathbf{X})$ distribution observed under NRAS.

Table 1: Benchmark of active sensing approaches on the synthetic dataset.

Method	ROC	PRC	Cost	$d_{\delta=0.3}$	$d_{\delta=0.5}$	$d_{\delta=0.7}$
FO	0.668±0.000	0.634±0.000	39.600±0.000	0.582±0.000	0.229±0.000	0.181±0.000
ASAC	0.582±0.035	0.527±0.023	9.189±1.895	1.052±0.339	1.326±0.063	1.323±0.065
Fixed	0.655±0.006	0.600±0.005	0.907±0.034†	1.384±0.000	1.398±0.000	1.359±0.000
LL	0.547±0.021	0.537±0.009	7.453±0.119	1.314±0.033	1.336±0.055	1.305±0.067
RAS (Ours)	**0.678±0.006**†	0.635±0.005	4.535±0.088	**0.142±0.018**†	**0.132±0.028**†	**0.154±0.032**†

The 95% confidence interval (CI) is evaluated with 5 evaluations with different random seeds. Best performance numbers in each column are highlighted in **bold**. Symbol † indicates p-value < 0.01.