SoQal: Selective Oracle Questioning in Active Learning

Dani Kiyasseh Tingting Zhu* David A. Clifton*

Department of Engineering Science
University of Oxford

CoRR 2020

Motivation

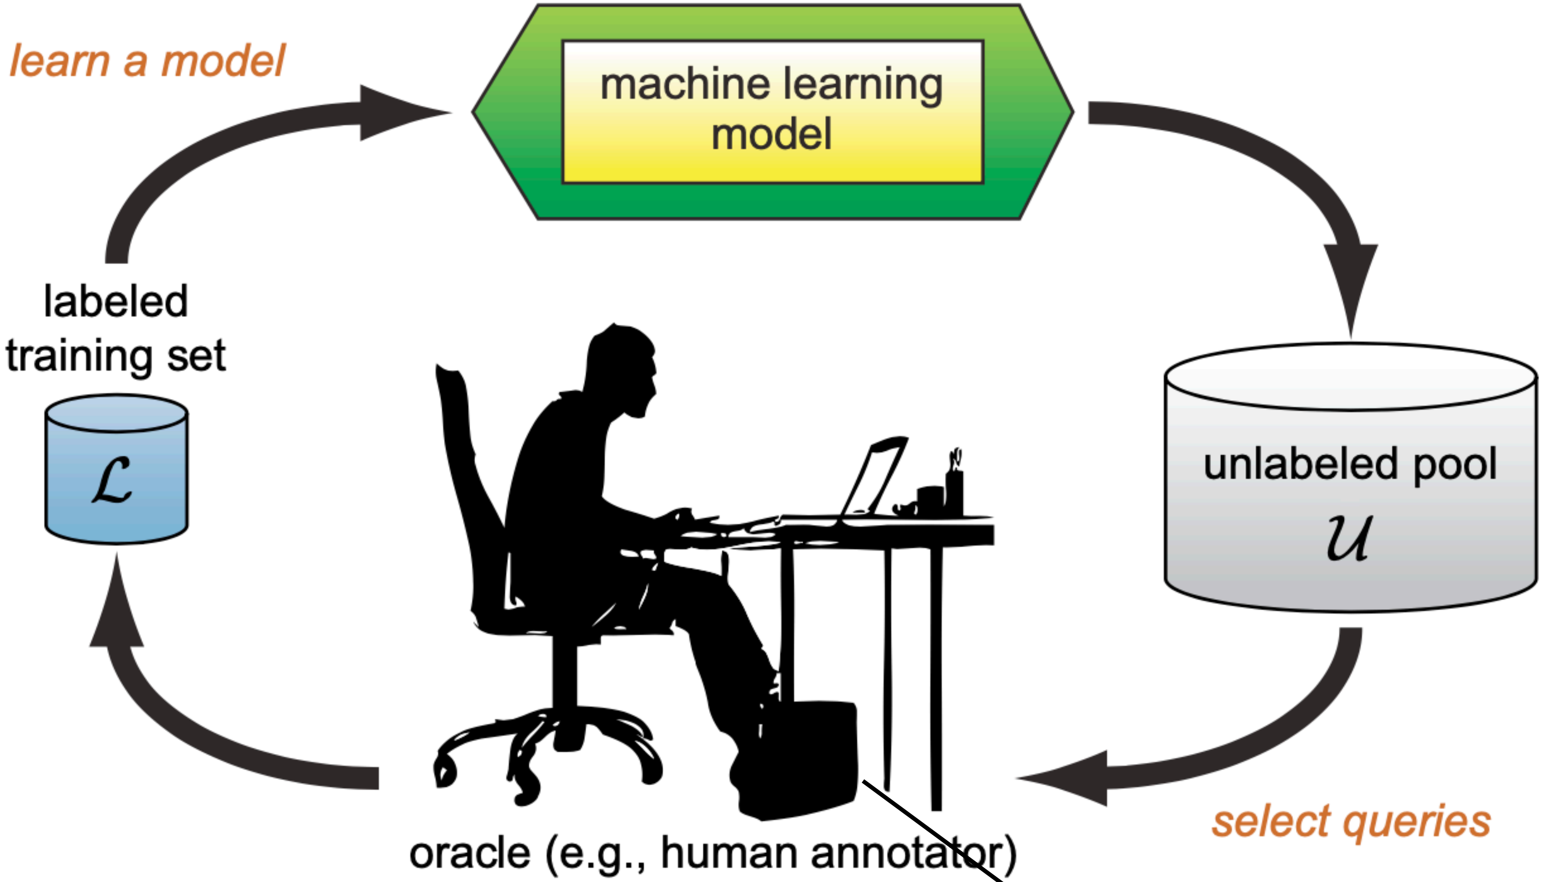

overly-dependent upon an oracle

Motivation

Over-reliance is detrimental:

- It negatively affects the applicability of AL algorithms to scenarios where an oracle is either ***unavailable*** or is ***unreliable***.
- Over-reliance can still ***inundate experts*** with a significant number of label requests, the very goal AL is supposed to minimize.

Selective Oracle Questioning in Active Learning(SoQal):

- a dynamic strategy that learns when to request a label from an oracle during active learning.

Selective Oracle Questioning

Architecture.: We assume the existence of a prediction network, f_ω, which for each instance, x, generates posterior class probabilities, $p(y|x, \omega)$, and an oracle selection network, $g_\theta : X \rightarrow o \in [0, 1]$ parameterized by θ that maps that same instance to a scalar, as shown in Fig. 1.

prediction network

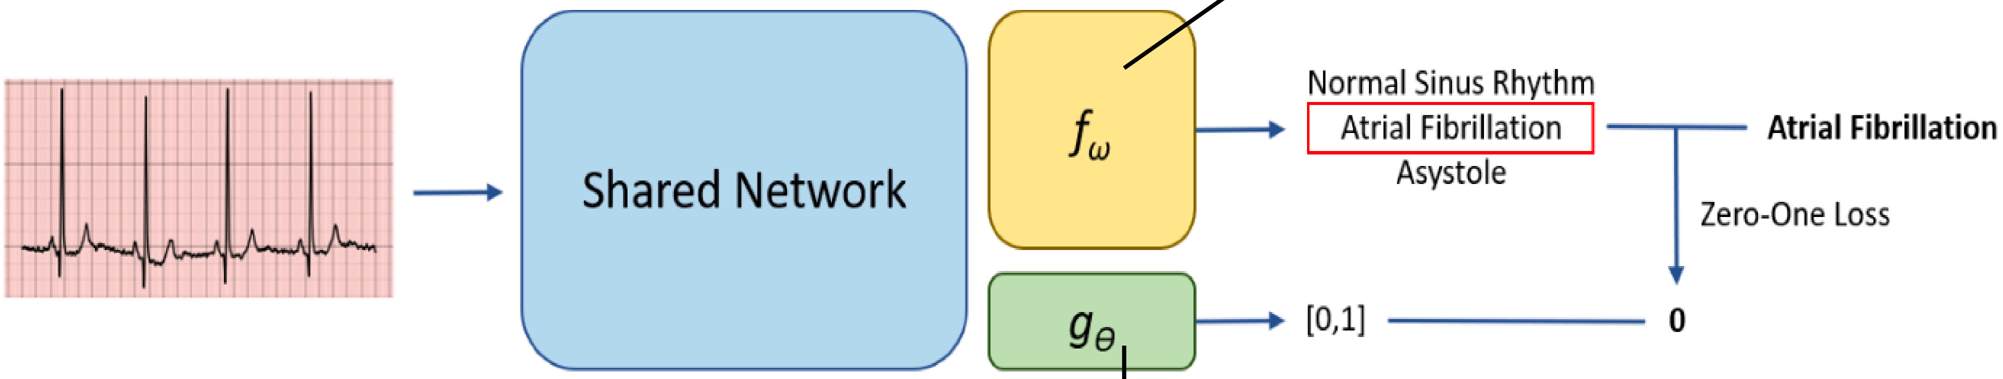

Figure 1: Selective Oracle Questioning Framework.

oracle selection network: approximating the probability that an oracle is requested for a label

Selective Oracle Questioning

Objective Function： Our objective function for a mini-batch of size, B, thus consists of two terms: 1) a cross-entropy class prediction loss for the main task, and 2) a weighted binary cross-entropy loss for the oracle selection network.

$$\mathcal{L} = \sum_{i=1}^{B} \overbrace{-\log\left(p(y_i = c|x_i, \omega)\right)}^{\text{Class Prediction Loss}} - \overbrace{\beta e_i \log\left(g_\theta(o|x_i)\right) - (1 - e_i)\log\left(1 - g_\theta(o|x_i)\right)}^{\text{Oracle Selection Loss}}$$

$\beta = \frac{\sum \delta_{e=0}}{\sum \delta\delta_{e=1}}$,a dynamic hyperparamenter,which changes according to the ratio of correctly classified to misclassified instances within a mini-batch.

Selective Oracle Questioning

Thresholding： As we are dealing with unlabelled instances, we are interested in exploiting the output of gθ as a proxy for whether an instance is correctly classified (e = 0) or not (e = 1). *The separability of these two states determine the reliability of such a proxy.* In Fig. 2b,

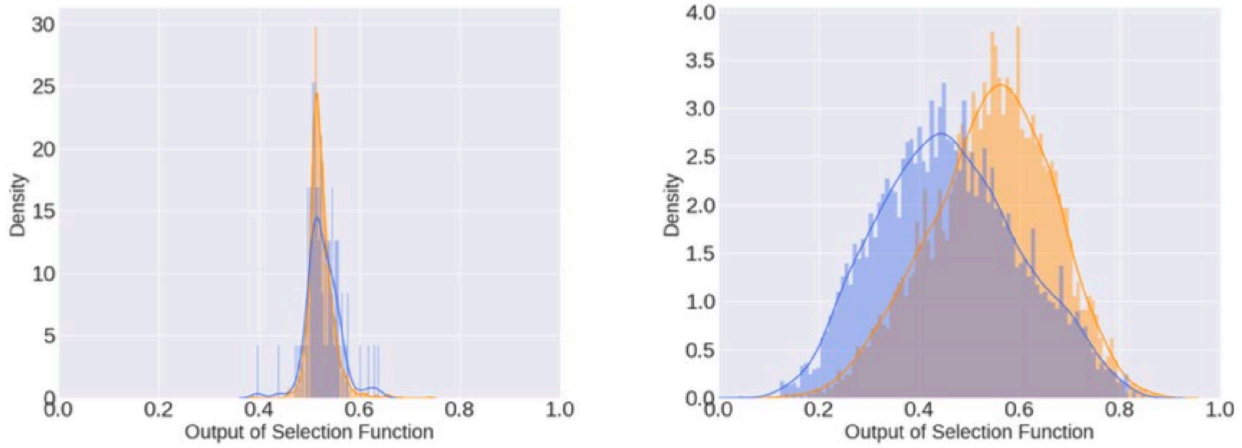

(a) Early in Training

(b) Late in Training

Quantify these parability of these two distributions using the Hellinger distance,DH ∈[0,1].

$$\mathcal{D}_H = \sqrt{1 - \sqrt{\frac{2\sigma_0\sigma_1}{\sigma_0^2\sigma_1^2}} e^{-\frac{1}{4}\frac{(\mu_0-\mu_1)^2}{\sigma_0^2\sigma_1^2}}}$$

Selective Oracle Questioning

We outline the probability of asking an oracle, p(A

$$p(\mathrm{A}) = \begin{cases} 1, & \mathcal{D}_H < S \\ 1, & \mathcal{N}\left(o|\mu_1, \sigma_1^2, e=1\right) > \mathcal{N}\left(o|\mu_0, \sigma_0^2, e=0\right) \text{ and } \mathcal{D}_H \geq S \\ 0, & \text{otherwise} \end{cases}$$

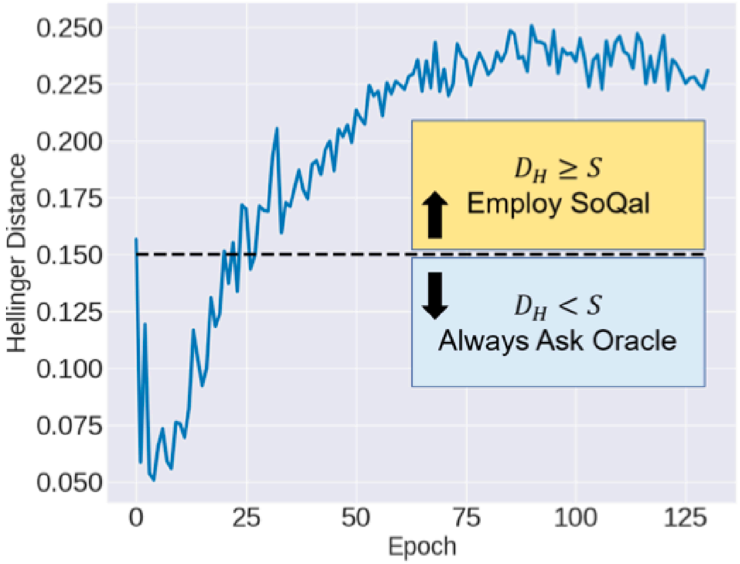

(c) Hellinger Distance

Experiment

1.Selective Oracle Questioning with Noise-Free Oracle

Table 1: Mean test AUC of oracle questioning strategies in the presence of a noise-free oracle. Results are shown for datasets $\mathcal{D}_1 - \mathcal{D}_5$ and all acquisition functions. Mean and standard deviation values are shown across five seeds. 'No AL' is the strategy that does not employ active learning.

Dataset	Ac. Function α	Oracle Questioning Method					
		No Oracle	Entropy Response	Epsilon Greedy	SoQal (ours)	100% Oracle	No AL
$\mathcal{D}_1$	BALD_{MCD}	0.465 ± 0.017	0.496 ± 0.039	0.491 ± 0.028	$\mathbf{0.621 \pm 0.021}$	0.653 ± 0.013	0.577 ± 0.014
	BALD_{MCP}	0.464 ± 0.023	0.517 ± 0.043	0.501 ± 0.043	$\mathbf{0.645 \pm 0.015}$	0.676 ± 0.020	
	BALC_{KLD}	0.500 ± 0.023	0.548 ± 0.034	0.548 ± 0.042	$\mathbf{0.598 \pm 0.055}$	0.634 ± 0.030	
	Temporal BALC_{KLD}	0.496 ± 0.024	0.536 ± 0.040	0.521 ± 0.059	$\mathbf{0.646 \pm 0.067}$	0.659 ± 0.033	
$\mathcal{D}_2$	BALD_{MCD}	0.573 ± 0.063	0.584 ± 0.041	0.609 ± 0.071	$\mathbf{0.707 \pm 0.038}$	0.713 ± 0.053	0.679 ± 0.040
	BALD_{MCP}	0.589 ± 0.045	0.638 ± 0.043	0.637 ± 0.044	$\mathbf{0.677 \pm 0.042}$	0.735 ± 0.028	
	BALC_{KLD}	0.602 ± 0.044	0.582 ± 0.017	0.643 ± 0.033	$\mathbf{0.677 \pm 0.024}$	0.722 ± 0.018	
	Temporal BALC_{KLD}	0.575 ± 0.017	0.612 ± 0.050	0.605 ± 0.019	$\mathbf{0.648 \pm 0.057}$	0.735 ± 0.011	
$\mathcal{D}_3$	BALD_{MCD}	0.581 ± 0.014	0.588 ± 0.013	0.673 ± 0.015	$\mathbf{0.721 \pm 0.025}$	0.802 ± 0.008	0.716 ± 0.012
	BALD_{MCP}	0.623 ± 0.020	0.676 ± 0.058	0.665 ± 0.028	$\mathbf{0.720 \pm 0.044}$	0.798 ± 0.007	
	BALC_{KLD}	0.631 ± 0.010	0.629 ± 0.004	0.643 ± 0.041	$\mathbf{0.731 \pm 0.033}$	0.787 ± 0.008	
	Temporal BALC_{KLD}	0.600 ± 0.005	0.630 ± 0.014	0.654 ± 0.019	$\mathbf{0.730 \pm 0.024}$	0.794 ± 0.002	
$\mathcal{D}_4$	BALD_{MCD}	0.486 ± 0.011	0.489 ± 0.030	0.474 ± 0.037	0.468 ± 0.021	0.585 ± 0.011	0.486 ± 0.023
	BALD_{MCP}	0.493 ± 0.030	0.504 ± 0.026	0.492 ± 0.024	0.499 ± 0.029	0.605 ± 0.024	
	BALC_{KLD}	0.505 ± 0.032	0.504 ± 0.039	0.473 ± 0.010	0.495 ± 0.012	0.588 ± 0.033	
	Temporal BALC_{KLD}	0.511 ± 0.030	0.496 ± 0.023	0.496 ± 0.023	0.503 ± 0.010	0.532 ± 0.027	
$\mathcal{D}_5$	BALD_{MCD}	0.717 ± 0.006	0.715 ± 0.005	$\mathbf{0.718 \pm 0.006}$	0.661 ± 0.105	0.937 ± 0.004	0.710 ± 0.097
	BALD_{MCP}	0.719 ± 0.009	0.678 ± 0.074	$\mathbf{0.774 \pm 0.047}$	0.453 ± 0.136	0.705 ± 0.013	
	BALC_{KLD}	0.679 ± 0.056	0.664 ± 0.064	$\mathbf{0.726 \pm 0.019}$	0.638 ± 0.145	0.900 ± 0.036	
	Temporal BALC_{KLD}	0.720 ± 0.010	0.689 ± 0.061	$\mathbf{0.741 \pm 0.028}$	0.571 ± 0.161	0.708 ± 0.002	

Experiment

1.Selective Oracle Questioning with Noise-Free Oracle

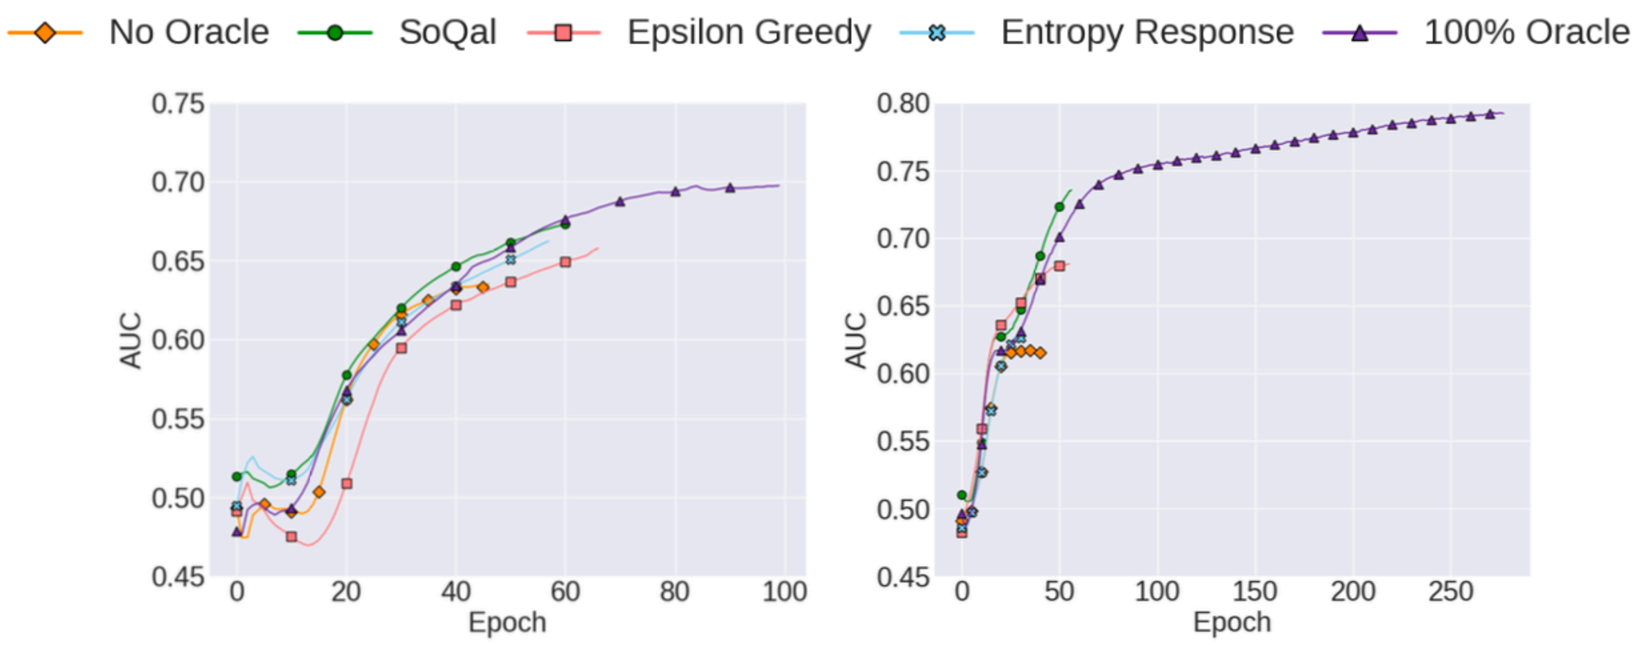

(a) $\mathcal{D}_2$, BALD_{MCP} (b) $\mathcal{D}_3$, BALD_{MCD}

Figure 3: Mean validation AUC as a function of oracle selection strategies on (a) $\mathcal{D}_2$ using BALD_{MCP} and (b) $\mathcal{D}_3$ using BALD_{MCD}. Results are averaged across 5 seeds.

Experiment

2.Selective Oracle Questioning with Noisy Oracle

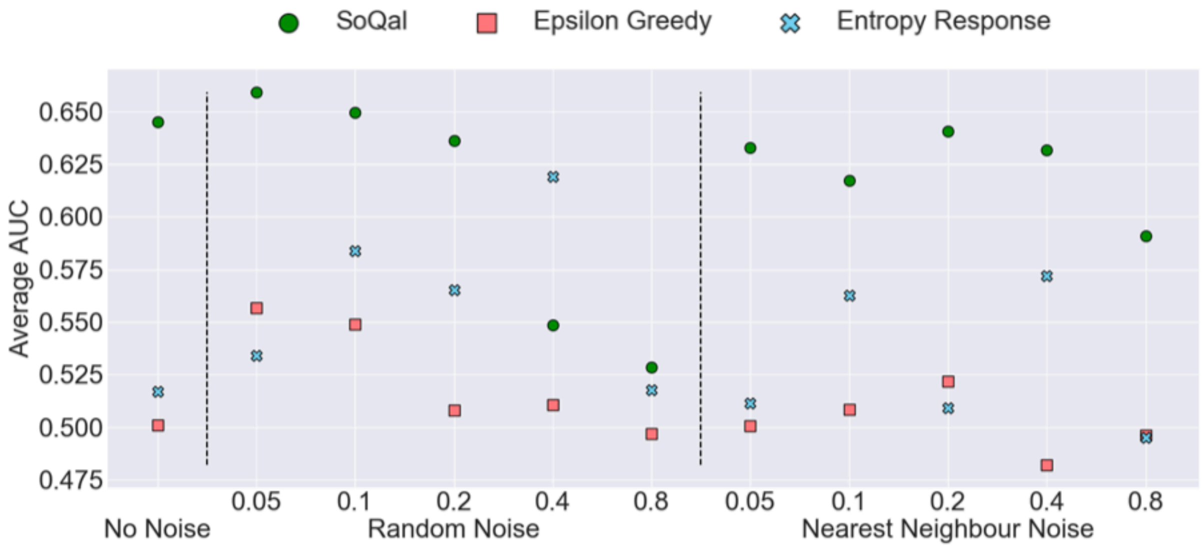

(a) $\mathcal{D}_1$, $\mathrm{BALD}_{\mathrm{MCP}}$

Figure 4: Average AUC of the oracle questioning strategies in the absence and presence of various magnitudes of label noise on $\mathcal{D}_1$ using $\mathrm{BALD}_{\mathrm{MCP}}$. With up to 80% random or nearest neighbour label noise, SoQal still outperforms its counterpart methods that are trained *without* label noise.

Thanks