## A Relation to Other Notions of Error

## A. 1 Connections to Cross-Entropy Loss

We establish that in the classification setting, our notion of error is equivalent to cross-entropy loss up to translations.
Example 1. (cross-entropy loss) Suppose the set $\mathcal{Y}$ of responses is finite. Then,

$$
\mathbf{d}_{\mathrm{KL}}\left(P^{*} \| P\right)=\sum_{y \in \mathcal{Y}} P^{*}(y) \ln \frac{P^{*}(y)}{P(y)}=\sum_{y \in \mathcal{Y}} P^{*}(y) \ln P^{*}(y)-\sum_{y \in \mathcal{Y}} P^{*}(y) \ln P(y)
$$

The first term does not depend on $P$, so minimizing KL-divergence is equivalent to minimizing the final term,

$$
-\sum_{y \in \mathcal{Y}} P^{*}(y) \ln P(y)=-\mathbb{E}\left[\ln \left(P\left(Y_{t+1}\right)\right) \mid \mathcal{E}, P, X\right]
$$

which is exactly the expected cross-entropy loss of $P$, as is commonly used to assess classifiers.

## A. 2 Connections to Mean-Squared Error

In the regression setting we first establish a direct link between KL-divergence and mean-squared error for the case in which $P_{t}^{*}$ and $P_{t}$ are Gaussian.
Example 2. (Gaussian mean-squared error) Fix $\mu \in \Re$ and $\sigma^{2} \in \Re_{++}$. Let $\mathbb{P}\left(Y_{t+1} \in \cdot \mid \mathcal{E}, X_{t}\right) \sim$ $\mathcal{N}\left(\mu, \sigma^{2}\right)$. Consider a point prediction $\hat{\mu}_{t}$ that is determined by $H_{t}$ and a distributional prediction $P_{t} \sim \mathcal{N}\left(\hat{\mu}_{t}, \sigma^{2}\right)$. Then,

$$
\mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| P_{t}\right)=\frac{\mathbb{E}\left[\left(\mu-\hat{\mu}_{t}\right)^{2} \mid \mathcal{E}, H_{t}\right]}{2 \sigma^{2}}=\frac{\mathbb{E}\left[\left(Y_{t+1}-\hat{\mu}_{t}\right)^{2} \mid \mathcal{E}, H_{t}\right]-\mathbb{E}\left[\left(Y_{t+1}-\mu_{t}\right)^{2} \mid \mathcal{E}, H_{t}\right]}{2 \sigma^{2}}
$$

Hence, KL-divergence grows monotonically the squared error $\mathbb{E}\left[\left(Y_{t+1}-\hat{\mu}_{t}\right)^{2} \mid \mathcal{E}, H_{t}\right]$. However, while the minimal squared error $\mathbb{E}\left[\left(Y_{t+1}-\mu\right)^{2} \mid \mathcal{E}, H_{t}\right]=\sigma^{2}$ that is attainable with full knowledge of the environment remains positive, the minimal KL-divergence, which is delivered by $P_{t} \sim \mathcal{N}\left(\mu, \sigma^{2}\right)$, is zero.

Now consider a distributional prediction $P_{t} \sim \mathcal{N}\left(\hat{\mu}_{t}, \hat{\sigma}_{t}^{2}\right)$, based on a variance estimate $\hat{\sigma}_{t}^{2} \neq \sigma^{2}$. Then,

$$
\mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| P_{t}\right)=\frac{\mathbb{E}\left[\left(\mu-\hat{\mu}_{t}\right)^{2} \mid \mathcal{E}, H_{t}\right]}{2 \hat{\sigma}_{t}^{2}}+\frac{1}{2}\left(\frac{\sigma^{2}}{\hat{\sigma}_{t}^{2}}-1-\ln \frac{\sigma^{2}}{\hat{\sigma}_{t}^{2}}\right) .
$$

Consider optimizing the choice of $\hat{\sigma}_{t}^{2}$ given $H_{t}$ :

$$
\min _{\hat{\sigma}_{t}^{2}} \mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| P_{t}\right) \mid H_{t}\right]
$$

The minimum is attained by

$$
\hat{\sigma}_{t}^{2}=\underbrace{\sigma^{2}}_{\text {aleatoric }}+\underbrace{\mathbb{E}\left[\left(\mu-\mathbb{E}\left[\mu \mid H_{t}\right]\right)^{2} \mid H_{t}\right]}_{\text {epistemic }}+\underbrace{\mathbb{E}\left[\left(\mathbb{E}\left[\mu \mid H_{t}\right]-\hat{\mu}_{t}\right)^{2} \mid H_{t}\right]}_{\text {bias }},
$$

which differs from $\sigma^{2}$. While $\sigma^{2}$ characterizes aleatoric uncertainty, the incremental variance $\hat{\sigma}_{t}^{2}-$ $\sigma^{2}$ accounts for epistemic uncertainty and bias.

Now, for $P_{t}^{*}$ and $P_{t}$ that are not Gaussian, we have the following upper bound:
Lemma 13. For all $t \in \mathbb{Z}_{+}$, if $\hat{\mu}_{t}=\int_{y \in \mathcal{Y}} y d P_{t}(y)$, then

$$
\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| P_{t}\right)\right] \leq \frac{1}{2} \ln \left(1+\frac{\mathbb{E}\left[\left(Y_{t+1}-\hat{\mu}_{t}\right)^{2}\right]}{\sigma^{2}}\right)
$$

Therefore, decreasing the mean squared error will always decrease the expected KL-divergence. A corresponding lower bound holds for data generating processes for which $Y_{t+1}$ satisfies a certain subgaussian condition:
Lemma 14. For all $t \in \mathbb{Z}_{+}$, let $\hat{\mu}_{t}=\int_{y \in \mathcal{Y}} y d P_{t}(y)$. If $P_{t}\left(Y_{t+1} \in \cdot\right)$ is $\delta_{t}^{2}$-subgaussian conditioned on $H_{t}$ w.p 1, then

$$
\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| P_{t}\right)\right] \geq \frac{\mathbb{E}\left[\left(Y_{t+1}-\hat{\mu}_{t}\right)^{2}\right]}{\delta_{t}^{2}}
$$

Therefore, for data generating processes that obey this subgaussian condition, we have both upper and lower bounds for expected KL divergence in terms of mean-squared error.

## B Proof of Information and Error Results

Here, we provide a proof of Theorem 3.
Theorem 3. (optimal prediction) For all $t \geq 0$,

$$
\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| \hat{P}_{t}\right) \mid H_{t}\right]=\inf _{\pi} \mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| P_{t}\right) \mid H_{t}\right]
$$

where $P_{t}=\pi\left(H_{t}, Z_{t}\right)$.
Proof. Let $\hat{P}_{t}=\mathbb{P}\left(Y_{t+1} \in \cdot \mid H_{t}\right)$. By Gibbs' inequality,

$$
\inf _{P_{t}} \mathbf{d}_{\mathrm{KL}}\left(\hat{P}_{t} \| P_{t}\right)=\mathbf{d}_{\mathrm{KL}}\left(\hat{P}_{t} \| \hat{P}_{t}\right)=0
$$

Let $P_{t}^{*}=\mathbb{P}\left(Y_{t+1} \in \cdot \mid \mathcal{E}, X_{t}\right)$. Then, for all $P_{t}$,

$$
\begin{aligned}
\mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| P_{t}\right)= & \mathbb{E}\left[\left.\ln \frac{d P_{t}^{*}}{d P_{t}}\left(Y_{t+1}\right) \right\rvert\, \mathcal{E}, H_{t}\right] \\
= & \mathbb{E}\left[\ln d P_{t}^{*}\left(Y_{t+1}\right) \mid \mathcal{E}, H_{t}\right]-\mathbb{E}\left[\ln d P_{t}\left(Y_{t+1}\right) \mid \mathcal{E}, H_{t}\right] \\
= & \mathbb{E}\left[\ln d P_{t}^{*}\left(Y_{t+1}\right) \mid \mathcal{E}, H_{t}\right]-\mathbb{E}\left[\ln d \hat{P}_{t}\left(Y_{t+1}\right) \mid \mathcal{E}, H_{t}\right] \\
& +\mathbb{E}\left[\ln d \hat{P}_{t}\left(Y_{t+1}\right) \mid \mathcal{E}, H_{t}\right]-\mathbb{E}\left[\ln d P_{t}\left(Y_{t+1}\right) \mid \mathcal{E}, H_{t}\right] \\
= & \mathbb{E}\left[\left.\ln \frac{d P_{t}^{*}}{d \hat{P}_{t}}\left(Y_{t+1}\right) \right\rvert\, \mathcal{E}, H_{t}\right]+\mathbb{E}\left[\left.\ln \frac{d \hat{P}_{t}}{d P_{t}}\left(Y_{t+1}\right) \right\rvert\, \mathcal{E}, H_{t}\right] \\
= & \mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| \hat{P}_{t}\right)+\mathbb{E}\left[\left.\ln \frac{d \hat{P}_{t}}{d P_{t}}\left(Y_{t+1}\right) \right\rvert\, \mathcal{E}, H_{t}\right] .
\end{aligned}
$$

It follows that

$$
\begin{aligned}
\inf _{\pi} \mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| P_{t}\right) \mid H_{t}\right] & =\inf _{\pi} \mathbb{E}\left[\left.\mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| \hat{P}_{t}\right)+\mathbb{E}\left[\left.\ln \frac{d \hat{P}_{t}}{d P_{t}}\left(Y_{t+1}\right) \right\rvert\, \mathcal{E}, H_{t}\right] \right\rvert\, H_{t}\right] \\
& =\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| \hat{P}_{t}\right) \mid H_{t}\right]+\inf _{\pi} \mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(\hat{P}_{t} \| P_{t}\right) \mid H_{t}\right] \\
& =\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| \hat{P}_{t}\right) \mid H_{t}\right]+\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(\hat{P}_{t} \| \hat{P}_{t}\right) \mid H_{t}\right] \\
& =\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| \hat{P}_{t}\right) \mid H_{t}\right] .
\end{aligned}
$$

A direct corollary of this result is the following shortfall quantification for a misspecified/suboptimal prediction.
Corollary 15. (misspecified/suboptimal prediction) For all $t \geq 0$ and $P_{t}=\pi\left(H_{t}, Z_{t}\right)$,

$$
\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| P_{t}\right) \mid H_{t}\right]=\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| \hat{P}_{t}\right) \mid H_{t}\right]+\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(\hat{P}_{t} \| P_{t}\right) \mid H_{t}\right]
$$

An interpretation of the above result is that the error of a suboptimal agent with prediction $P_{t}$ is the sum of the error of an optimal agent $\hat{P}_{t}$ and the expected KL divergence between $\hat{P}_{t}$ and $P_{t}$. Since all the results in this paper bound the first term, in order to understand the behavior of a suboptimal agent, techniques to bound the second term would be required.

Next, we provide a proof for Lemma 4.
Lemma 4. (expected prediction error equals information gain) For all $t \in \mathbb{Z}_{+}$,

$$
\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| \hat{P}_{t}\right)\right]=\mathbb{I}\left(\mathcal{E} ; Y_{t+1} \mid H_{t}\right),
$$

and $\mathcal{R}(t)=\mathbb{I}\left(\mathcal{E} ; H_{t}\right)$.

Proof. It is well known that the mutual information $\mathbb{I}(A ; B)$ between random variables $A$ and $B$ can be expressed in terms of the expected KL-divergence $\mathbb{I}(A ; B)=\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}(\mathbb{P}(A \in \cdot \mid B) \| \mathbb{P}(A \in \cdot))\right]$. It follows that

$$
\begin{aligned}
\mathbb{I}\left(Y_{t+1} ; \mathcal{E} \mid H_{t}\right) & =\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(\mathbb{P}\left(Y_{t+1} \in \cdot \mid \mathcal{E}, H_{t}\right) \| \mathbb{P}\left(Y_{t+1} \in \cdot \mid H_{t}\right)\right)\right] \\
& \stackrel{(a)}{=} \mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(\mathbb{P}\left(Y_{t+1} \in \cdot \mid \mathcal{E}, X_{t}\right) \| \mathbb{P}\left(Y_{t+1} \in \cdot \mid H_{t}\right)\right)\right] \\
& =\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| \hat{P}_{t}\right)\right]
\end{aligned}
$$

where $(a)$ follows from the fact that $Y_{t+1} \perp H_{t} \mid\left(\mathcal{E}, X_{t}\right)$. We then have

$$
\mathcal{R}(T)=\mathbb{E}\left[\sum_{t=0}^{T-1} \mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| \hat{P}_{t}\right)\right]=\sum_{t=0}^{T-1} \mathbb{I}\left(Y_{t+1} ; \mathcal{E} \mid H_{t}\right) \stackrel{a}{=} \mathbb{I}\left(\mathcal{E} ; H_{T}\right)
$$

where $(a)$ follows from the chain rule of mutual information.

Lastly, we provide a proof for Lemma 5.
Lemma 5. (expected prediction error is monotonically nonincreasing) For all $t \in \mathbb{Z}_{+}$,

$$
\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| \hat{P}_{t}\right)\right] \geq \mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P_{t+1}^{*} \| \hat{P}_{t+1}\right)\right] .
$$

Proof. We have

$$
\begin{aligned}
\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P_{t+1}^{*} \| \hat{P}_{t+1}\right)\right] & \stackrel{(a)}{=} \mathbb{I}\left(\mathcal{E} ; Y_{t+2} \mid H_{t+1}\right) \\
& =\mathbf{h}\left(Y_{t+2} \mid H_{t+1}\right)-\mathbf{h}\left(Y_{t+2} \mid \mathcal{E}, H_{t+1}\right) \\
& \stackrel{(b)}{=} \mathbf{h}\left(Y_{t+2} \mid H_{t+1}\right)-\mathbf{h}\left(Y_{t+2} \mid \mathcal{E}, H_{t-1}, Y_{t}, X_{t+1}\right) \\
& \stackrel{(c)}{\leq} \mathbf{h}\left(Y_{t+2} \mid H_{t-1}, Y_{t}, X_{t+1}\right)-\mathbf{h}\left(Y_{t+2} \mid \mathcal{E}, H_{t-1}, Y_{t}, X_{t+1}\right) \\
& =\mathbb{I}\left(\mathcal{E} ; Y_{t+2} \mid H_{t-1}, Y_{t}, X_{t+1}\right) \\
& \stackrel{(d)}{=} \mathbb{I}\left(\mathcal{E} ; Y_{t+1} \mid H_{t-1}, Y_{t}, X_{t}\right) \\
& =\mathbb{I}\left(\mathcal{E} ; Y_{t+1} \mid H_{t}\right) \\
& \stackrel{(e)}{=} \mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P_{t}^{*} \| \hat{P}_{t}\right)\right],
\end{aligned}
$$

where (a) follows from Lemma 4, (b) follows from the the fact that $Y_{t+2} \perp\left(X_{t}, Y_{t+1}\right) \mid\left(\mathcal{E}, X_{t+1}\right)$, (c) follows from the fact that conditioning reduces differential entropy, $(d)$ follows from the fact that $\left(X_{t}, Y_{t+1}\right)$ and $\left(X_{t+1}, Y_{t+2}\right)$ are independent and identically distributed conditioned on $\left(H_{t-1}, Y_{t}\right)$, and $(e)$ follows from the equivalence between mutual information and expected KL-divergence.

## C Proofs of Regret and Sample Complexity Results

In this section, we prove results pertaining to bounding optimal regret and sample complexity in terms of rate-distortion.

Lemma 6. (proxy error equals information gain) For all $\tilde{\mathcal{E}} \in \tilde{\Theta}$,

$$
\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P^{*} \| \tilde{P}\right)\right]=\mathbb{I}(\mathcal{E} ; Y \mid \tilde{\mathcal{E}}, X)
$$

Proof. It is well known that the mutual information $\mathbb{I}(A ; B)$ between random variables $A$ and $B$ can be expressed in terms of the expected KL-divergence $\mathbb{I}(A ; B)=\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}(\mathbb{P}(A \in \cdot \mid B) \| \mathbb{P}(A \in \cdot))\right]$. We therefore have

$$
\begin{aligned}
\mathbb{I}(\mathcal{E} ; Y \mid \tilde{\mathcal{E}}, X) & =\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}(\mathbb{P}(Y \in \cdot \mid \mathcal{E}, \tilde{\mathcal{E}}, X) \| \mathbb{P}(Y \in \cdot \mid \tilde{\mathcal{E}}, X))\right] \\
& =\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}(\mathbb{P}(Y \in \cdot \mid \mathcal{E}, X) \| \mathbb{P}(Y \in \cdot \mid \tilde{\mathcal{E}}, X))\right] \\
& =\mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P^{*} \| \tilde{P}\right)\right]
\end{aligned}
$$

where the second equation follows from the fact that $(X, Y) \perp \tilde{\mathcal{E}} \mid \mathcal{E}$.
Theorem 8. (rate-distortion regret bounds) For all $T$,

$$
\sup _{\epsilon \geq 0} \min \left\{\mathbb{H}_{\epsilon}(\mathcal{E}), \epsilon T\right\} \leq \mathcal{R}(T) \leq \inf _{\epsilon \geq 0}\left(\mathbb{H}_{\epsilon}(\mathcal{E})+\epsilon T\right)
$$

Proof. We begin by establishing the upper bound.

$$
\begin{aligned}
\mathcal{R}(T) & =\sum_{t=0}^{T-1} \mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P_{T}^{*} \| \hat{P}_{t}\right)\right] \\
& \stackrel{(a)}{=} \sum_{t=0}^{T-1} \mathbb{I}\left(Y_{t+1} ; \mathcal{E} \mid H_{t}\right) \\
& =\sum_{t=0}^{T-1} \mathbb{I}\left(Y_{t+1} ; \mathcal{E}, \tilde{\mathcal{E}} \mid H_{t}\right) \\
& \stackrel{(b b}{=} \sum_{t=0}^{T-1} \mathbb{I}\left(Y_{t+1} ; \tilde{\mathcal{E}} \mid H_{t}\right)+\mathbb{I}\left(Y_{t+1} ; \mathcal{E} \mid \tilde{\mathcal{E}}, H_{t}\right) \\
& \stackrel{(c)}{\leq} \sum_{t=0}^{T-1} \mathbb{I}\left(Y_{t+1} ; \tilde{\mathcal{E}} \mid H_{t}\right)+\mathbb{I}\left(Y_{t+1} ; \mathcal{E} \mid \tilde{\mathcal{E}}, H_{t}\right) \\
& \stackrel{(d)}{=} \mathbb{I}\left(H_{T} ; \tilde{\mathcal{E}}\right)+\sum_{t=0}^{T-1} \mathbb{I}\left(Y_{t+1} ; \mathcal{E} \mid \tilde{\mathcal{E}}, H_{t}\right) \\
& \stackrel{(e)}{\leq} \mathbb{I}\left(H_{T} ; \tilde{\mathcal{E}}\right)+\sum_{t=0}^{T-1} \mathbb{I}\left(Y_{t+1} ; \mathcal{E} \mid \tilde{\mathcal{E}}, X_{t}\right) \\
& \stackrel{(f)}{\leq} \mathbb{I}\left(H_{T} ; \tilde{\mathcal{E}}\right)+\epsilon T \\
& \stackrel{(g)}{\leq} \mathbb{I}(\mathcal{E} ; \tilde{\mathcal{E}})+\epsilon T
\end{aligned}
$$

where (a) follows from Lemma 4, (b) follows from the chain rule of mutual information, $(c)$ follows from the chain rule of mutual information, $(d)$ follows from the facts that $\mathbf{h}\left(Y_{t+1} \mid \tilde{\mathcal{E}}, H_{t}\right) \leq$ $\mathbf{h}\left(Y_{t+1} \mid \tilde{\mathcal{E}}, X_{t}\right)$ and $\mathbf{h}\left(Y_{t+1} \mid \mathcal{E}, H_{t}\right)=\mathbf{h}\left(Y_{t+1} \mid \mathcal{E}, X_{t}\right),(e)$ holds for any $\tilde{\mathcal{E}} \in \tilde{\Theta}_{\epsilon}$, and $(f)$ follows from the data processing inequality. Since the above inequality holds for all $\epsilon \geq 0$ and $\tilde{\mathcal{E}} \in \tilde{\Theta}_{\epsilon}$, the result follows.
Next, we establish the lower bound. Fix $T \in \mathbb{Z}_{+}$. Let $\tilde{\mathcal{E}}=\left(\tilde{H}_{T-2}, \tilde{Y}_{T-1}\right)$ be independent from but distributed identically with $\left(H_{T-2}, Y_{T-1}\right)$, conditioned on $\mathcal{E}$. In other words, $\tilde{\mathcal{E}} \perp\left(H_{T-2}, Y_{T-1}\right) \mid \mathcal{E}$ and $\mathbb{P}(\tilde{\mathcal{E}} \in \cdot \mid \mathcal{E})=\mathbb{P}\left(\left(H_{T-2}, Y_{T-1}\right) \in \cdot \mid \mathcal{E}\right)$. This implies that $\mathbb{P}\left(\left(\mathcal{E}, \tilde{H}_{T-2}, \tilde{Y}_{T-1}, X_{T-1}, Y_{T}\right) \in\right.$ $\cdot)=\mathbb{P}\left(\left(\mathcal{E}, H_{T-2}, Y_{T-1}, X_{T-1}, Y_{T}\right) \in \cdot\right)$, and therefore, $\mathbb{I}\left(\mathcal{E} ; Y_{T} \mid H_{T-1}\right)=\mathbb{I}\left(\mathcal{E} ; Y_{T} \mid \mathcal{E}, X_{T-1}\right)$.

Fix $\epsilon \geq 0$. If $\mathcal{R}(T)<\mathbb{H}_{\epsilon}(\mathcal{E})$ then,

$$
\begin{aligned}
\mathcal{R}(T) & \stackrel{(a)}{=} \mathbb{I}\left(\mathcal{E} ; H_{T}\right) \\
& \stackrel{(b)}{=} \sum_{t=0}^{T-1} \mathbb{I}\left(\mathcal{E} ; Y_{t+1} \mid H_{t}\right) \\
& \stackrel{(c)}{\geq} \mathbb{I}\left(\mathcal{E} ; Y_{T} \mid H_{T-1}\right) T \\
& =\mathbb{I}\left(\mathcal{E} ; Y_{T} \mid \tilde{\mathcal{E}}, X_{T-1}\right) T \\
& \stackrel{(d)}{=} \mathbb{E}\left[\mathbf{d}_{\mathrm{KL}}\left(P_{T-1}^{*} \| \tilde{P}_{T-1}\right)\right] T \\
& \stackrel{(e)}{>} \epsilon T
\end{aligned}
$$

where (a) follows from Lemma 4, (b) follows from the chain rule of mutual information, (c) follows from Lemma 5, (d) follows from Lemma 6, and (e) follows from the fact that $\tilde{\mathcal{E}} \notin \tilde{\Theta}_{\epsilon}$ because $\mathcal{R}(T)=\mathbb{I}\left(\mathcal{E} ; H_{T}\right)=\mathbb{I}(\mathcal{E} ; \tilde{\mathcal{E}})<H_{\epsilon}(\mathcal{E})$. Therefore,

$$
\mathcal{R}(T) \geq \min \left\{\mathbb{H}_{\epsilon}(\mathcal{E}), \epsilon T\right\}
$$

Since this holds for any $\epsilon \geq 0$, the result follows.
Theorem 9. (rate-distortion sample complexity bounds) For all $\epsilon \geq 0$,

$$
\frac{\mathbb{H}_{\epsilon}(\mathcal{E})}{\epsilon} \leq T_{\epsilon} \leq \inf _{\delta \in[0, \epsilon]}\left\lceil\frac{\mathbb{H}_{\epsilon-\delta}(\mathcal{E})}{\delta}\right\rceil \leq\left\lceil\frac{2 \mathbb{H}_{\epsilon / 2}(\mathcal{E})}{\epsilon}\right\rceil
$$

Proof. We begin by showing the upper bound. Fix $\epsilon \geq 0$ and $\delta \in[0, \epsilon]$. Let

$$
T=\left\lceil\frac{\mathbb{H}_{\epsilon-\delta}(\mathcal{E})}{\delta}\right\rceil \text {, }
$$

so that $\mathbb{H}_{\epsilon-\delta}(\mathcal{E}) \leq \delta T$. We have that:

$$
\begin{aligned}
\mathcal{R}(T) & \stackrel{(a)}{\leq} \mathbb{H}_{\epsilon-\delta}(\mathcal{E})+(\epsilon-\delta) T \\
& \stackrel{(b)}{\leq} \delta T+(\epsilon-\delta) T \\
& =\epsilon T
\end{aligned}
$$

where (a) follows from the upper bound of Theorem 8 and $(b)$ follows from our choice of $T$. Since $T_{\epsilon}=\min \{T: \mathcal{R}(T) \leq \epsilon \mathcal{R}(1) T\}$, it follows that $T \geq T_{\epsilon}$. Since the above holds for arbitrary $\delta \in[0, \epsilon]$, the result follows.
We now show the lower bound. Fix $\epsilon \geq 0$. By the definition of $T_{\epsilon}$, we have

$$
\mathcal{R}\left(T_{\epsilon}\right) \leq \epsilon T_{\epsilon}
$$

In the proof of the lower bound in Theorem 8, we show that for all $\epsilon \geq 0, \mathcal{R}(T)<\mathbb{H}_{\epsilon}(\mathcal{E}) \Longrightarrow$ $\mathcal{R}(T)>\epsilon T$. Therefore, using the contrapositive and the above definition of $T_{\epsilon}$, we have that $\mathbb{H}_{\epsilon}(\mathcal{E}) \leq \mathcal{R}\left(T_{\epsilon}\right)$ and therefore

$$
\mathbb{H}_{\epsilon}(\mathcal{E}) \leq \mathcal{R}\left(T_{\epsilon}\right) \leq \epsilon T_{\epsilon}
$$

The result follows.

## D Proofs of Single-Layer Rate-Distortion Bounds

We have established in Theorem 10 that the rate-distortion function for a multilayer environment is a sum of the rate-distortion functions for each single-layer environment under tolerance $\frac{\epsilon}{K}$. In this section, we show the following rate-distortion bounds for a single layer of the deep ReLU network and the deep nonparametric network.

We begin by stating a well known result from information theory:

Lemma 16. (maximum differential entropy) For all random vectors $X: \Omega \mapsto \Re^{d}$ with covariance $K$,

$$
\mathbf{h}(X) \leq \frac{1}{2} \ln \left((2 \pi e)^{d}|K|\right)
$$

with equality iff $\operatorname{Pr}(X \in \cdot) \sim \mathcal{N}(\mu, K)$ for some $\mu \in \Re^{d}$.

Proof. Follows from Theorems 8.6.3 and 8.6.5 of Cover and Thomas [2006].

Theorem 17. (single-layer relu neural network rate-distortion bound) For all $N, k \in \mathbb{Z}_{++}$and $\sigma^{2}, \epsilon \geq 0$, if $\mathcal{E}_{k}$ is identified by $(A, b)$ where $A: \Omega \mapsto \Re^{N \times N}, b: \Omega \mapsto \Re^{N}$ both consist of independent elements each with variance $\frac{1}{N}, X: \Omega \mapsto \Re^{N}$ is a random vector with covariance $I_{N}$, and $Y \sim \mathcal{N}\left(\operatorname{ReLU}(A X+b), \sigma^{2} I_{N}\right)$, then

$$
\mathbb{H}_{\epsilon}\left(\mathcal{E}_{k}\right)=\mathcal{O}\left(N^{2} \ln \left(\frac{N}{2 \sigma^{2} \epsilon}\right)\right)
$$

Proof. We use $A_{i}, b_{i}$ to denote the $i$ th rows of $A, b$ respectively. Let $\tilde{\mathcal{E}}_{k}=(\tilde{A}, \tilde{b})$ where $\tilde{A}=A+V$ where $V \perp A$ and $V \sim \mathcal{N}\left(0, \delta^{2} I_{d}\right)$ and likewise $\tilde{b}=b+Z$ where $Z \perp b$ and $Z \sim \mathcal{N}\left(0, \delta^{2}\right)$. $\delta^{2}=\frac{\sigma^{2}\left(e^{\frac{2 \epsilon}{N}}-1\right)}{N+1}$. We have that

$$
\begin{aligned}
\mathbb{I}\left(\mathcal{E}_{k} ; \tilde{\mathcal{E}}_{k}\right) & =\mathbb{I}\left(\mathcal{E}_{k} ; \tilde{\mathcal{E}}_{k}\right) \\
& =\mathbb{I}(A ; \tilde{A})+\mathbb{I}(b ; \tilde{b}) \\
& \leq \mathbf{h}(\tilde{A})-\mathbf{h}(\tilde{A} \mid A)+\mathbb{I}(b ; \tilde{b}) \\
& =\frac{N^{2}}{2} \ln \left(2 \pi e\left(\delta^{2}+\frac{1}{N}\right)\right)-\mathbf{h}(V)+\mathbb{I}(b ; \tilde{b}) \\
& =\frac{N^{2}}{2} \ln \left(1+\frac{1}{N \delta^{2}}\right)+\frac{N}{2} \ln \left(1+\frac{1}{N \delta^{2}}\right) \\
& \leq \frac{N(N+1)}{2} \ln \left(1+\frac{1}{N \delta^{2}}\right) \\
& =\mathcal{O}\left(N^{2} \ln \left(1+\frac{1}{N \delta^{2}}\right)\right) \\
& =\mathcal{O}\left(N^{2} \ln \left(1+\frac{1}{\sigma^{2}\left(e^{\frac{2 \epsilon}{N}}-1\right)}\right)\right) \\
& =\mathcal{O}\left(N^{2} \ln \left(\frac{N}{2 \sigma^{2} \epsilon}\right)\right) .
\end{aligned}
$$

We now verify that our choice of $\tilde{\mathcal{E}}_{k}$ satisfies the distortion constraint:

$$
\begin{aligned}
\mathbb{I}\left(Y ; \mathcal{E}_{k} \mid \tilde{\mathcal{E}}_{k}, X\right) & =\mathbb{I}(Y ; A, b \mid \tilde{A}, \tilde{b}, S, X) \\
& =\mathbf{h}(Y \mid \tilde{A}, \tilde{b}, S, X)-\mathbf{h}(Y \mid A, b, S, X) \\
& =N\left(\mathbf{h}\left(Y_{i} \mid \tilde{A}, \tilde{b}, S, X\right)-\mathbf{h}\left(Y_{i} \mid A, b, S, X\right)\right) \\
& \stackrel{(a)}{=} N\left(\mathbf{h}\left(Y_{i} \mid \tilde{A}_{i}, \tilde{b}_{i}, S_{i}, X\right)-\mathbf{h}\left(W_{i}\right)\right) \\
& =N\left(\mathbf{h}\left(Y_{i}-\operatorname{ReLU}\left(\tilde{A}_{i}^{\top} X+\tilde{b}_{i}\right) \mid \tilde{\theta}_{i}, \tilde{b}_{i}, S_{i}, X\right)-\mathbf{h}\left(W_{i}\right)\right) \\
& \stackrel{(b)}{\leq} N\left(\mathbf{h}\left(Y_{i}-\operatorname{ReLU}\left(\tilde{A}_{i}^{\top} X+\tilde{b}_{i}\right) \mid X\right)-\mathbf{h}\left(W_{i}\right)\right) \\
& \stackrel{(c)}{\leq} \mathbb{E}\left[N\left(\frac{1}{2} \ln \left(2 \pi e\left(\sigma^{2}+\mathbb{V}\left[\operatorname{ReLU}\left(A_{i}^{\top} X+b_{i}\right)-\operatorname{ReLU}\left(\tilde{A}_{i}^{\top} X+\tilde{b}_{i}\right) \mid X\right]\right)\right)-\mathbf{h}\left(W_{i}\right)\right)\right] \\
& \stackrel{(d)}{\leq} \frac{N}{2} \ln \left(1+\frac{\mathbb{V}\left[\operatorname{ReLU}\left(A_{i}^{\top} X+b_{i}\right)-\operatorname{ReLU}\left(\tilde{A}_{i}^{\top} X+\tilde{b}_{i}\right)\right]}{\sigma^{2}}\right) \\
& (e) N \\
& \frac{N}{2} \ln \left(1+\frac{\mathbb{V}\left[A_{i}^{\top} X+b_{i}-\tilde{A}_{i}^{\top} X-\tilde{b}_{i}\right]}{\sigma^{2}}\right) \\
& =\frac{N}{2} \ln \left(1+\frac{(N+1) \delta^{2}}{\sigma^{2}}\right) \\
& =\epsilon
\end{aligned}
$$

where in $(a), W_{i} \sim \mathcal{N}\left(0, \sigma^{2}\right),(b)$ follows from the fact that conditioning reduces entropy, $(c)$ follows from Lemma 16, (d) follows from Jensen's Inequality, and $(e)$ follows from the fact that for all $x, y \in \Re,(\operatorname{ReLU}(x)-\operatorname{ReLU}(y))^{2} \leq(x-y)^{2}$.

Theorem 18. (single-layer nonparametric rate-distortion bound) For all $N, M, c \in \mathbb{Z}_{++}$and $\sigma^{2}, \epsilon \geq 0$, if $\mathcal{E}$ is identified by $\alpha$ where $\alpha \sim \operatorname{Dir}\left(N,\left[\frac{c}{N}, \ldots, \frac{c}{N}\right]\right)$ and

$$
f(X)=\sqrt{c} \cdot \sum_{n=1}^{N} \alpha_{n} g_{n}(X)
$$

for deterministic basis functions $\left(g_{1}, \ldots, g_{N}\right)$ mapping $\Re^{M} \mapsto \Re^{M}$ that satisfy $\mathbb{E}\left[g_{n}(X)^{\top} g_{n}(X)\right] \leq$ 1 for random vector $X: \Omega \mapsto \Re^{N}$, and $Y \sim \mathcal{N}\left(f(X), \sigma^{2} I_{M}\right)$, then

$$
\mathbb{H}_{\epsilon}(\mathcal{E}) \leq \frac{c \ln (2 N)}{2 \sigma^{2} \epsilon}
$$

Proof. Let $\tilde{\mathcal{E}}=\frac{\sqrt{c}}{r} \sum_{i=1}^{r} h_{i}$ for $r=\frac{c}{2 \sigma^{2} \epsilon}$ and where for all $i \in\{1, \ldots, r\}$,

$$
h_{i} \stackrel{i i d}{\sim}\left\{\delta_{n} \cdot g_{n} \quad \text { w.p } \alpha_{n} .\right.
$$

The distortion of $\tilde{\mathcal{E}}$ is

$$
\begin{aligned}
\mathbb{I}(f(X)+W ; \mathcal{E} \mid \tilde{\mathcal{E}}, X) & =\mathbf{h}(f(X)+W \mid \tilde{\mathcal{E}}, X)-\mathbf{h}(f(X)+W \mid \mathcal{E}, X) \\
& =\mathbf{h}\left(\left.f(X)-\frac{\sqrt{c}}{r} \sum_{i=1}^{r} h_{i}(X)+W \right\rvert\, \tilde{\mathcal{E}}, X\right)-\mathbf{h}(W) \\
& \stackrel{(a)}{\leq} \mathbf{h}\left(f(X)-\frac{\sqrt{c}}{r} \sum_{i=1}^{r} h_{i}(X)+W\right)-\mathbf{h}(W) \\
& \stackrel{(b)}{\leq} \frac{M}{2} \ln \left(2 \pi e\left(\sigma^{2}+\mathbb{E}\left[\frac{\left\|f(X)-\frac{\sqrt{c}}{r} \sum_{i=1}^{r} h_{i}(X)\right\|^{2}}{M}\right]\right)\right)-\mathbf{h}(W) \\
& =\frac{M}{2} \ln \left(1+\frac{\left.\mathbb{E}\left[\left\|f(X)-\frac{\sqrt{c}}{r} \sum_{i=1}^{r} h_{i}(X)\right\|^{2}\right]\right)}{\sigma^{2} M^{\prime}}\right) \\
& \stackrel{(c)}{\leq} \frac{\mathbb{E}\left[\left\|f(X)-\frac{\sqrt{c}}{r} \sum_{i=1}^{r} h_{i}(X)\right\|^{2}\right]}{\sigma^{2}} \\
& =\frac{\mathbb{E}\left[\left(\frac{\sqrt{c}}{r} \sum_{i=1}^{r} h_{i}(X)\right)^{\top}\left(\frac{\sqrt{c}}{r} \sum_{i=1}^{r} h_{i}(X)\right)-f(X)^{\top} f(X)\right]}{2 \sigma^{2}} \\
& =\frac{\mathbb{E}\left[\frac{c}{r^{2}} \sum_{i=1}^{r} h_{i}(X)^{\top} h_{i}(X)+\frac{c}{r^{2}} \sum_{i \neq j} h_{i}(X)^{\top} h_{j}(X)-f(X)^{\top} f(X)\right]}{2 \sigma^{2}} \\
& =\frac{\mathbb{E}\left[\frac{c}{r} \sum_{n=1}^{N} \alpha_{i} g_{n}(X)^{\top} g_{n}(X)+\frac{c(r-1)}{r} f(X)^{\top} f(X)-f(X)^{\top} f(X)\right]}{2 \sigma^{2}} \\
& \leq \frac{\mathbb{E}\left[c-f(X)^{\top} f(X)\right]}{2 r \sigma^{2}} \\
& \leq \frac{c}{2 r \sigma^{2}} \\
& =\epsilon
\end{aligned}
$$

where (a) follows from conditioning reducing differential entropy, (b) follows from Lemma 16, and (c) follows from the fact that for all $m, x \in \Re_{++}, m \ln \left(1+\frac{x}{m}\right) \leq x$.

We now upper bound the rate of $\tilde{\mathcal{E}}$.

$$
\begin{aligned}
\mathbb{I}(\alpha ; \tilde{\mathcal{E}}) & \leq \mathbb{H}(\tilde{\mathcal{E}}) \\
& \leq r \cdot \mathbb{H}\left(h_{i}\right) \\
& \leq r \ln (2 N) \\
& =\frac{c \ln (2 N)}{2 \sigma^{2} \epsilon}
\end{aligned}
$$

The result follows.

## E Proof of Multilayer Results

Lemma 19. (more is learned with the true input) Let $\tilde{\mathcal{E}}_{K: 1}$ be a multilayer proxy. Then,

$$
\mathbb{I}\left(Y ; \mathcal{E}_{k} \mid \mathcal{E}_{K: k+1}, \tilde{\mathcal{E}}_{k: 1}, X\right) \leq \mathbb{I}\left(Y ; \mathcal{E}_{k} \mid \mathcal{E}_{K: k+1}, \tilde{\mathcal{E}}_{k}, U_{k-1}\right) .
$$

Proof.

$$
\begin{aligned}
& \mathbb{I}\left(\mathcal{E}_{k} ; Y \mid \mathcal{E}_{K: k+1}, \tilde{\mathcal{E}}_{k: 1}, X\right) \\
& =\mathbb{I}\left(\mathcal{E}_{k} ; \tilde{\mathcal{E}}_{k-1: 1}, X, Y \mid \mathcal{E}_{K: k+1}, \tilde{\mathcal{E}}_{k}\right)-\mathbb{I}\left(\mathcal{E}_{k} ; \tilde{\mathcal{E}}_{k-1: 1}, X \mid \mathcal{E}_{K: k+1}, \tilde{\mathcal{E}}_{k}\right) \\
& =\mathbb{I}\left(\mathcal{E}_{k} ; \tilde{\mathcal{E}}_{j-1: 1}, X, Y \mid \mathcal{E}_{K: k+1}, \tilde{\mathcal{E}}_{k}\right) \\
& =\mathbb{I}\left(\mathcal{E}_{k} ; Y \mid \mathcal{E}_{K: k+1}, \tilde{\mathcal{E}}_{k}\right)+\mathbb{I}\left(\mathcal{E}_{k} ; \tilde{\mathcal{E}}_{k-1: 1}, X \mid \mathcal{E}_{K: k+1}, \tilde{\mathcal{E}}_{k}, Y\right) \\
& \stackrel{(a)}{\leq} \mathbb{I}\left(\mathcal{E}_{k} ; Y \mid \mathcal{E}_{K: k+1}, \tilde{\mathcal{E}}_{k}\right)+\mathbb{I}\left(\mathcal{E}_{k} ; \mathcal{E}_{k-1: 1}, X \mid \mathcal{E}_{K: k+1}, \tilde{\mathcal{E}}_{k}, Y\right) \\
& =\mathbb{I}\left(\mathcal{E}_{k} ; \mathcal{E}_{k-1,1}, X, Y \mid \mathcal{E}_{K: k+1}, \tilde{\mathcal{E}}_{k}\right) \\
& \stackrel{(b)}{=} \mathbb{I}\left(\mathcal{E}_{k} ; Y \mid \mathcal{E}_{K: k+1}, \tilde{\mathcal{E}}_{k}, \mathcal{E}_{k-1: 1}, X\right) \\
& \stackrel{(c)}{=} \mathbb{I}\left(\mathcal{E}_{k} ; Y \mid \mathcal{E}_{K: k+1}, \tilde{\mathcal{E}}_{k}, U_{k-1}\right),
\end{aligned}
$$

where $(a)$ follows from the fact that $\mathcal{E}_{k} \perp \tilde{\mathcal{E}}_{k-1: 1} \mid\left(X, Y, \mathcal{E}_{k-1: 1}\right)$ and the data processing inequality, (b) follows from the fact that $\mathbb{I}\left(\mathcal{E}_{k} ; \mathcal{E}_{k-1: 1}, X \mid \mathcal{E}_{K: k+1}, \tilde{\mathcal{E}}_{k}\right)=0$, and (c) follows from the fact that $Y \perp\left(\mathcal{E}_{k-1: 1}, X\right) \mid U_{k-1}$.

Lemma 19 states that we learn more information about $\mathcal{E}_{k}$ when we are given the true input $U_{k-1}$ than when we are given $\left(X, \tilde{\mathcal{E}}_{k-1: 1}\right)$ and have to infer $U_{k-1}$. This is intuitive as we should be able to recover more about $\mathcal{E}_{k}$ when we observe its input exactly.

To prove 2), that the immediate output $U_{k}$ provides more information about $\mathcal{E}_{k}$ than does $Y$, each prototypical environment will involve slightly different analysis techniques. In essence, we expect a result like 2) to hold as a result of the data-processing inequality. No amount of post-processing of $U_{k}$ into $(Y)$ can increase the amount of available information conveyed about $\mathcal{E}_{k}$.

Lemma 20. For all $K \in \mathbb{Z}_{++},\left(N_{0}, N_{1}, \ldots, N_{K}\right) \in \mathbb{Z}_{++}^{K+1}$, and $\sigma^{2} \geq 0$, if $N_{K}=1$, and multilayer environment $\mathcal{E}_{K: 1}$ consists of single-layer environments $\mathcal{E}_{k}$ that each identify a random function $f^{(k)}$ from $\Re^{N_{k-1}} \mapsto \Re^{N_{k}}$ for which

$$
L^{(k)}=\sup _{x, y \in \Re^{N_{k-1}}} \mathbb{E}\left[\left.\frac{\left\|f^{(k)}(x)-f^{(k)}(y)\right\|_{2}^{2}}{\|x-y\|_{2}^{2}} \right\rvert\, x=x, y=y\right] \leq 1
$$

and

$$
U_{k}= \begin{cases}X & k=0 \\ f^{(k)}\left(U_{k-1}\right) & 1 \leq k \leq K\end{cases}
$$

with $Y \sim \mathcal{N}\left(U_{K}, \sigma^{2}\right)$, then for any proxy $\tilde{\mathcal{E}}_{k}$,

$$
\mathbb{I}\left(Y ; \mathcal{E}_{k} \mid \tilde{\mathcal{E}}_{k}, \mathcal{E}_{K: k+1}, U_{k-1}\right) \leq \frac{1}{2} \ln \left(1+\frac{\mathbb{E}\left[\left\|U_{k}-\mathbb{E}\left[U_{k} \mid \tilde{\mathcal{E}}_{k}, U_{k-1}\right]\right\|_{2}^{2}\right]}{\sigma^{2}}\right)
$$

Proof.

$$
\begin{aligned}
\mathbb{I}\left(Y ; \mathcal{E}_{k} \mid \tilde{\mathcal{E}}_{k}, \mathcal{E}_{K: k+1}, U_{k-1}\right) & =\mathbf{h}\left(Y \mid \tilde{\mathcal{E}}_{k}, \mathcal{E}_{K: k+1}, U_{k-1}\right)-\mathbf{h}\left(Y \mid \mathcal{E}_{K}: k, U_{k-1}\right) \\
& =\mathbf{h}\left(Y \mid \tilde{\mathcal{E}}_{k}, \mathcal{E}_{K: k+1}, U_{k-1}\right)-\mathbf{h}(W) \\
& =\mathbf{h}\left(Y-\left(f_{K} \circ \ldots \circ f_{k+1}\right)\left(\mathbb{E}\left[U_{k} \mid \tilde{\mathcal{E}}, U_{k-1}\right]\right) \mid \tilde{\mathcal{E}}_{k}, \mathcal{E}_{K: k+1}, U_{k-1}\right)-\mathbf{h}(W) \\
& \leq \mathbf{h}\left(Y-\left(f_{K} \circ \ldots \circ f_{k+1}\right)\left(\mathbb{E}\left[U_{k} \mid \tilde{\mathcal{E}}, U_{k-1}\right]\right)\right)-\mathbf{h}(W) \\
& \leq \frac{1}{2} \ln \left(1+\frac{\mathbb{V}\left[U_{K}-\left(f_{K} \circ \ldots \circ f_{k+1}\right)\left(\mathbb{E}\left[U_{k} \mid \tilde{\mathcal{E}}, U_{k-1}\right]\right)\right]}{\sigma^{2}}\right) \\
& \leq \frac{1}{2} \ln \left(1+\frac{\mathbb{E}\left[\left(U_{K}-\left(f_{K} \circ \ldots \circ f_{k+1}\right)\left(\mathbb{E}\left[U_{k} \mid \tilde{\mathcal{E}}, U_{k-1}\right]\right)\right)^{2}\right]}{\sigma^{2}}\right) \\
& \leq \frac{1}{2} \ln \left(1+\frac{\left.\mathbb{E}\left[\prod_{i=k+1}^{K} L^{(i)}\left\|U_{k}-\mathbb{E}\left[U_{k} \mid \tilde{\mathcal{E}}_{k}, U_{k-1}\right]\right\|^{2}\right]\right)}{\sigma^{2}}\right) \\
& \leq \frac{1}{2} \ln \left(1+\frac{\mathbb{E}\left[\left\|U_{k}-\mathbb{E}\left[U_{k} \mid \tilde{\mathcal{E}}_{k}, U_{k-1}\right]\right\|^{2}\right]}{\sigma^{2}}\right)
\end{aligned}
$$

We now provide a proof of Theorem 10.
Theorem 10. multilayer rate-distortion bound For all $K \in \mathbb{Z}_{++}, \sigma^{2}, \epsilon \geq 0$, if $\mathcal{E}_{K: 1}$ is a multilayer environment such that for all $k \in\{1, \ldots, K\}$ and $\delta \geq 0$, there exist $\tilde{\mathcal{E}}_{k}$ s.t.

$$
\mathbb{I}\left(Y ; \mathcal{E}_{k} \mid \mathcal{E}_{K: k+1}, \tilde{\mathcal{E}}_{k: 1}, X\right) \leq \mathbb{I}\left(U_{k}+W ; \mathcal{E}_{k} \mid \tilde{\mathcal{E}}_{k}, U_{k-1}\right) \leq \delta
$$

where $W \sim \mathcal{N}\left(0, \sigma^{2} I\right)$, then

$$
\mathbb{H}_{\epsilon}\left(\mathcal{E}_{K: 1}\right) \leq \sum_{k=1}^{K} \mathbb{H}_{\frac{\epsilon}{K}}\left(\mathcal{E}_{k}\right)
$$

Proof. Let

$$
\tilde{\Theta}_{\epsilon}^{K: 1}=\left\{\tilde{\mathcal{E}} \in \tilde{\Theta}_{\epsilon}: \tilde{\mathcal{E}}=\left(\tilde{\mathcal{E}}_{1}, \ldots, \tilde{\mathcal{E}}_{K}\right) ; \tilde{\mathcal{E}}_{i} \perp \tilde{\mathcal{E}}_{j} \wedge \tilde{\mathcal{E}}_{i} \perp \mathcal{E}_{j} \text { for } i \neq j\right\}
$$

We have that

$$
\begin{aligned}
\inf _{\tilde{\mathcal{E}} \in \tilde{\Theta}_{\epsilon}} \mathbb{I}(\mathcal{E} ; \tilde{\mathcal{E}}) & =\inf _{\tilde{\mathcal{E}} \in \tilde{\Theta}_{\epsilon}} \sum_{k=1}^{K} \mathbb{I}\left(\mathcal{E}_{k} ; \tilde{\mathcal{E}} \mid \mathcal{E}_{K: k+1}\right) \\
& \stackrel{(a)}{\leq} \inf _{\tilde{\mathcal{E}} \in \tilde{\Theta}_{\epsilon}^{K: 1}} \sum_{k=1}^{K} \mathbb{I}\left(\mathcal{E}_{k} ; \tilde{\mathcal{E}} \mid \mathcal{E}_{K: k+1}\right) \\
& \stackrel{(b)}{=} \inf _{\tilde{\mathcal{E}} \in \tilde{\Theta}_{\epsilon}^{K: 1}} \sum_{k=1}^{K} \mathbb{I}\left(\mathcal{E}_{k} ; \tilde{\mathcal{E}}_{k}\right) \\
& \stackrel{(c)}{\leq} \sum_{k=1}^{K} \inf _{\tilde{\mathcal{E}}_{k} \in \tilde{\Theta}_{\frac{e}{K}}^{(k)}} \mathbb{I}\left(\mathcal{E}_{k}, \tilde{\mathcal{E}}_{k}\right) \\
& =\sum_{k=1}^{K} \mathbb{H}_{\frac{\epsilon}{K}}\left(\mathcal{E}_{k}\right)
\end{aligned}
$$

where (a) follows from the fact that $\tilde{\Theta}_{\epsilon}^{K: 1} \subset \tilde{\Theta}_{\epsilon},(b)$ follows from the fact that $\tilde{\mathcal{E}}_{i} \perp \mathcal{E}_{j}$ for $i \neq j$, and $(c)$ follows from the fact that for $\tilde{\Theta}_{\epsilon / K}^{(k)}:=\left\{\tilde{\mathcal{E}}_{k} \in \tilde{\Theta}: \mathbb{I}\left(U_{k}+W ; \mathcal{E}_{k} \mid \tilde{\mathcal{E}}_{k}, U_{k-1}\right) \leq \frac{\epsilon}{K}\right\}, \tilde{\Theta}_{\epsilon / K}^{(1)} \times$ $\ldots \times \tilde{\Theta}_{\epsilon / K}^{(K)} \subset \tilde{\Theta}_{\epsilon}^{K: 1}$ since we assumed that $\mathbb{I}\left(Y ; \mathcal{E}_{k} \mid \mathcal{E}_{K: k+1}, \tilde{\mathcal{E}}_{k: 1}, X\right) \leq \mathbb{I}\left(U_{k}+W ; \mathcal{E}_{k} \mid \tilde{\mathcal{E}}_{k}, U_{k-1}\right)$ for all $k$.

## F Proof of Main Results

Theorem 11. (relu neural network rate-distortion and sample complexity bound) For all $d, N, K \in \mathbb{Z}_{++}$and $\sigma^{2}, \epsilon \geq 0$, if multilayer environment $\mathcal{E}_{K: 1}$ is the deep ReLU network with input $X: \Omega \mapsto \Re^{d}$ s.t. $\mathbb{V}[X]=I_{d}$ and output $Y \sim \mathcal{N}\left(U_{K}, \sigma^{2}\right)$, then

$$
\mathbb{H}_{\epsilon}\left(\mathcal{E}_{K: 1}\right)=\tilde{\mathcal{O}}\left(K N^{2}+d N\right), \quad T_{\epsilon}=\tilde{\mathcal{O}}\left(\frac{K N^{2}+d N}{\epsilon}\right)
$$

Proof. If we have a distortion function $d\left(\mathcal{E}_{k}, \tilde{\mathcal{E}}_{k}\right)$ for which

$$
\mathbb{I}\left(Y ; \mathcal{E}_{k} \mid \mathcal{E}_{K: k+1}, \tilde{\mathcal{E}}_{k: 1}, X\right) \leq d\left(\mathcal{E}_{k}, \tilde{\mathcal{E}}_{k}\right)
$$

then by the same proof as in Theorem 10,

$$
\mathbb{H}_{\epsilon}\left(\mathcal{E}_{K: 1}\right) \leq \sum_{k=1}^{K} \mathbb{H}_{\frac{e}{K}}\left(\mathcal{E}_{k}, d\right),
$$

where $\mathbb{H}_{\frac{\epsilon}{K}}\left(\mathcal{E}_{k}, d\right)$ denotes the rate-distortion function for random variable $\mathcal{E}_{k}$ under distortion function $d\left(\mathcal{E}_{k}, \tilde{\mathcal{E}}_{k}\right)$. By Lemma 20, we have the above condition for $d\left(\mathcal{E}_{k}, \tilde{\mathcal{E}}_{k}\right)=$ $\frac{1}{2} \ln \left(1+\frac{\mathbb{E}\left[\left\|U_{k}-\mathbb{E}\left[U_{k} \mid \tilde{\mathcal{E}}_{k}, U_{k-1}\right]\right\|_{2}^{2}\right]}{\sigma^{2}}\right)$. Furthermore, distortion used in the proof of Theorem 17 is exactly $d\left(\mathcal{E}_{k} ; \tilde{\mathcal{E}}_{k}\right)$. As a result,

$$
\begin{aligned}
\mathbb{H}_{\epsilon}\left(\mathcal{E}_{K: 1}\right) & \leq \sum_{k=1}^{K} \mathbb{H}_{\frac{e}{K}}\left(\mathcal{E}_{k}, d\right) \\
& \stackrel{(a)}{=} \mathcal{O}\left((K-2) N^{2} \ln \left(\frac{N K}{2 \sigma^{2} \epsilon}\right)+d N \ln \left(\frac{d K}{2 \sigma^{2} \epsilon}\right)+N \ln \left(\frac{N K}{2 \sigma^{2} \epsilon}\right)\right) \\
& =\tilde{O}\left(K N^{2}+d N\right)
\end{aligned}
$$

where (a) follows from Theorem 17. The sample complexity result follows from applying Theorem 9.

