## Appendix

## A Probabilistic Specifications: Examples

Below we provide further examples of specifications that can be captured by our framework.
Uncertainty calibration. Another desirable specification towards ensuring reliable uncertainty calibration for NNs is that the expected uncertainty in the predictions increases monotonically with an increase in the variance of the input-noise distribution. Formally, consider a set of zeromean distributions $\mathcal{P}_{\text {noise }}$ with diagonal covariance matrices. For any two such noise distributions $p_{\omega_{1}}, p_{\omega_{2}} \in \mathcal{P}_{\text {noise }}$ such that $\operatorname{Var}\left(p_{\omega_{1}}\right) \succcurlyeq \operatorname{Var}\left(p_{\omega_{2}}\right)$ (where $\succcurlyeq$ is the element-wise inequality and $\operatorname{Var}$ denotes the diagonal of the covariance matrix) and a given image $x$ from the training distribution, we wish to guarantee that the expected entropy of the resulting predictions corresponding to $p_{\omega_{1}}$ is greater than that of $p_{\omega_{2}}$, i.e., $\mathbb{E}_{x_{1} \sim x+p_{\omega_{1}}}\left[H\left(\operatorname{softmax}\left(\phi\left(x_{1}\right)\right] \geq \mathbb{E}_{x_{2} \sim x+p_{\omega_{2}}}\left[H\left(\operatorname{sof} \operatorname{tmax}\left(\phi\left(x_{2}\right)\right]\right.\right.\right.\right.$, where $H$ is the entropy functional: $H(p)=-\sum_{i=1}^{|\mathcal{Y}|} p_{i} \log p_{i}$. Intuitively, this captures the desired behaviour that as the strength of the noise grows, the expected uncertainty in the network predictions increases. We can capture this specification within the formulation described by equation 1 , by letting:

1. $\mathcal{P}_{0}=\mathcal{P}_{\text {noise }} \times \mathcal{P}_{\text {noise }}$,
2. For a given $\mathrm{NN} \phi$ and an input $x$, we define another network $\bar{\phi}: \mathcal{X} \times \mathcal{X} \rightarrow \mathcal{P}(\mathcal{Y}) \times \mathcal{P}(\mathcal{Y})$, where $\bar{\phi}$ is such that: $\bar{\phi}(a, b)=(\phi(x+a), \phi(x+b))$.

We can then define the verification problem as certifying that the following holds :

$$
\begin{aligned}
\boldsymbol{\psi}\left(\bar{\phi}\left(\bar{p}_{0}, \overline{\bar{p}}_{0}\right)\right):=-\left(\underset{\left(y_{1}, y_{2}\right) \sim \bar{\phi}\left(\bar{p}_{0}, \overline{\bar{p}}_{0}\right)}{\mathbb{E}}[ \right. & \left.\left.H\left(\operatorname{softmax}\left(y_{1}\right)\right)-H\left(\operatorname{softmax}\left(y_{2}\right)\right)\right]\right) \leq 0 \\
& \forall\left(\bar{p}_{0}, \overline{\bar{p}}_{0}\right) \in \mathcal{P}_{0} \text { such that } \operatorname{Var}\left(\bar{p}_{0}\right) \succcurlyeq \operatorname{Var}\left(\overline{\bar{p}}_{0}\right)
\end{aligned}
$$

Robust VAE In Dathathri et al. [2020], the authors consider a specification that corresponds to certifying low reconstruction losses for a VAE decoder over a set of inputs in the neighborhood of the latent-variable mean predicted by the encoder for a given image. A natural generalization of this specification is one where low reconstruction error is guaranteed in expectation, since in practice the latent-representations that are fed into the decoder are drawn from a normal distribution whose mean and variance are predicted by the encoder. A more general specification is one where we wish to verify that for a set of norm-bounded points around a given input, the expected reconstruction error from the VAE is small. Formally, for a VAE $\phi$ (note that a VAE directly fits within our framework, where the distribution for the latent-variable can be obtained as the output of a stochastic layer), a given threshold $\tau \in \mathbb{R}_{+}$and for a set of inputs $\mathcal{S}$, we wish to certify that the following holds:

$$
\begin{equation*}
\forall p_{0} \in \mathcal{P}_{0}, \quad \boldsymbol{\psi}\left(\phi\left(p_{0}\right)\right):=\underset{\mu_{\phi}^{s^{\prime}} \sim \phi\left(p_{0}\right)}{\mathbb{E}}\left[\left\|\underset{s \sim p_{0}}{\mathbb{E}}[s]-\mu_{\phi}^{s^{\prime}}\right\|_{2}^{2}\right]-\tau \leq 0 \tag{8}
\end{equation*}
$$

where $\mathcal{P}_{0}=\left\{\delta_{s}: s \in \mathcal{S}\right\}$.

## B Proof of Functional Lagrangian Theorem

## Proof of Theorem 1

Proof. The optimization problem (2) is equivalent to the following optimization problem:

$$
\begin{array}{ll}
\max _{p_{0}, p_{1}, \ldots, p_{K}} & \boldsymbol{\psi}\left(p_{K}\right) \\
\text { Subject to } & p_{k+1}(y)=\int \pi_{k}(y \mid x) p_{k}(x) d x \quad \forall y \in \mathcal{X}_{k+1}, \forall k \in\{0, \ldots, K-1\} \\
& p_{0} \in \mathcal{P}_{0} \tag{9b}
\end{array}
$$

where $\psi$ is a functional of the output distribution $p_{K}$. We refer to the space of probability measures on $\mathcal{X}_{k}$ as $\mathcal{P}_{k}$ (not that for $k=0$, this may not be the whole space of probability measures but a
constrained set of measures depending on the specification we would like to verify). The dual of this optimization problem can be written as follows:

$$
\begin{aligned}
\max _{p_{0} \in \mathcal{P}_{0}, p_{1} \in \mathcal{P}_{1}, \ldots} & \boldsymbol{\psi}\left(p_{K}\right)-\sum_{k=0}^{K-1} \int \lambda_{k+1}(z) p_{k+1}(z) d z \\
& +\sum_{k=0}^{K-1} \int_{\mathcal{X}_{k+1}, \mathcal{X}_{k}} \lambda_{k+1}(z) \pi_{k}(z \mid x) p_{k}(x) d x d z
\end{aligned}
$$

where we assigned Lagrange multipliers $\lambda_{k+1}(y)$ for every $y \in \mathcal{X}_{k+1}$. The above optimization problem can be decomposed as follows:

$$
\begin{aligned}
& \boldsymbol{\psi}\left(p_{K}\right)-\int_{\mathcal{X}_{K}} \lambda_{K}(x) p_{K}(x) d x \\
& +\sum_{k=0}^{K-1} \int_{\mathcal{X}_{k}}\left(\int_{\mathcal{X}_{k+1}} \lambda_{k+1}(y) \pi_{k}(y \mid x) d y-\lambda_{k}(x)\right) p_{k}(x) d x
\end{aligned}
$$

This is a separable problem in each $p_{k}$ and since $p_{k}$ is constrained to be a probability measure, the optimal choice of $p_{k}$ (for $k=1, \ldots, K-1$ ) is a $\delta$ measure with probability 1 assigned to the the $x \in \mathcal{X}_{k}$ that maximizes:

$$
\int_{\mathcal{X}_{k+1}} \lambda_{k+1}(y) \pi_{k}(y \mid x) d y-\lambda_{k}(x)
$$

The optimization over $p_{0}$ can be rewritten as follows:

$$
\max _{p_{0} \in \mathcal{P}_{0}} \int_{\mathcal{X}_{0}}\left(\int_{\mathcal{X}_{1}} \lambda_{1}(y) \pi_{0}(y \mid x) d y\right) p_{0}(x) d x
$$

The optimization over $p_{K}$ can be rewritten as follows:

$$
\boldsymbol{\psi}^{\star}\left(\lambda_{K}\right)=\max _{p_{K} \in \mathcal{P}_{K}} \boldsymbol{\psi}\left(p_{K}\right)-\int \lambda_{K}(x) p_{K}(x) d x
$$

The overall dual problem can be rewritten as:

$$
\begin{aligned}
& \boldsymbol{\psi}^{\star}\left(\lambda_{K}\right)+\sum_{k=1}^{K-1} \max _{x \in \mathcal{X}_{k}}\left(\int_{\mathcal{X}_{k+1}} \lambda_{k+1}(y) \pi_{k}(y \mid x) d y-\lambda_{k}(x)\right) \\
& \quad+\max _{p_{0} \in \mathcal{P}_{0}} \int_{\mathcal{X}_{0}}\left(\int_{\mathcal{X}_{1}} \lambda_{1}(y) \pi_{0}(y \mid x) d y\right) p_{0}(x) d x
\end{aligned}
$$

Writing this in terms of expected values, we obtain $g(\lambda)$. Plugging in $\lambda^{\star}$ into $g(\lambda)$, all the terms cancel except the first term which evaluates to:

$$
\max _{p_{0} \in \mathcal{P}_{0}} \underset{x \sim p_{0}}{\mathbb{E}}\left[\underset{y \in \pi_{0}(x)}{\mathbb{E}}\left[\lambda_{1}^{\star}(x)\right]\right]=\max _{p_{0} \in \mathcal{P}_{0}} \underset{x \sim p_{0}}{\mathbb{E}}\left[\underset{y \in \pi_{0}(x)}{\mathbb{E}}\left[\underset{z \sim \pi_{1}(x)}{\mathbb{E}}\left[\lambda_{2}^{\star}(z)\right]\right]\right] \ldots=\max _{p_{0} \in \mathcal{P}_{0}} \boldsymbol{\psi}\left(p_{0}\right)
$$

## C Additional Theoretical Results

## C. 1 Computation of Expected Values

Since $\pi_{k}(x)$ is typically a distribution that one can sample from easily (as it is required to perform forward inference through the neural network), estimating this expectation via sampling is a viable option. However, in order to turn this into verified bounds on the specification, one needs to appeal to concentration inequalities and the final guarantees would only be probabilistically valid. We leave this direction for future work.

Instead, we focus on situations where the expectations can be computed in closed form. In particular, we consider layers of the form $\pi_{k}(x)=w s(x)+b,(w, b) \sim p_{k}^{w}$, where $s$ is an element-wise function like ReLU, sigmoid or tanh and $(w, b)$ represents a fully connected or convolutional layer. We consider a general form of Lagrange multipliers as a sum of quadratic and exponential terms as follows:

$$
\lambda(x)=q^{T} x+\frac{1}{2} x^{T} Q x+\kappa \exp \left(\gamma^{T} x\right)
$$

Let:

$$
\tilde{s}(x)=\binom{1}{s(x)}, \tilde{Q}=\left(\begin{array}{cc}
0 & q^{T} \\
q & Q
\end{array}\right), \tilde{w}=\left(\begin{array}{ll}
b & w
\end{array}\right) .
$$

Then:

$$
\lambda(W s(x)+b)=\frac{1}{2}(\tilde{s}(x))^{T} \tilde{w}^{T} \tilde{Q} \tilde{w} \tilde{s}(x)+\kappa \exp \left(\gamma^{T} \tilde{w} \tilde{s}(x)\right)
$$

Taking expected values with respect to $\tilde{W} \sim p_{k}^{w}$, we obtain:

$$
\begin{equation*}
\mathbb{E}[\lambda(w s(x)+b)]=\frac{1}{2}(\tilde{s}(x))^{T} \underset{\tilde{s}(x)}{\mathbb{E}}\left[\tilde{w}^{T} \tilde{Q} \tilde{w}\right]+\kappa \prod_{i, j} \mathbb{E}\left[\exp \left(\gamma_{i} \tilde{w}_{i j} \tilde{s}\left(x_{j}\right)\right)\right] \tag{10}
\end{equation*}
$$

where we have assumed that each element of $\tilde{W}_{i j}$ is independently distributed. The first term in (10) evaluates to:

$$
\frac{1}{2} \operatorname{Trace}(\operatorname{Cov}(\tilde{w} \tilde{s}(x)) \tilde{Q})+\frac{1}{2}(\mathbb{E}[\tilde{w} \tilde{s}(x)])^{\top} \tilde{Q}(\mathbb{E}[\tilde{w} \tilde{s}(x)])
$$

and the second one to:

$$
\kappa \prod_{i, j} \operatorname{mgf}_{i j}\left(\gamma_{i} \tilde{s}_{j}(x)\right)
$$

where $\operatorname{Cov}(X)$ refers to the covariance matrix of the random vector, and $\mathrm{mgf}_{i j}$ refers to the moment generating function of the random variable $\tilde{w}_{i j}$ :

$$
\operatorname{mgf}_{i j}(\theta)=\mathbb{E}\left[\exp \left(\tilde{w}_{i j} \theta\right)\right]
$$

The details of this computation for various distributions on $\tilde{w}$ (Gaussian posterior, MC-dropout) are worked out below.

Diagonal Gaussian posterior. Consider a BNN with a Gaussian posterior, $\tilde{w} \sim \mathcal{N}\left(\mu, \operatorname{diag}\left(\sigma^{2}\right)\right)$, where $\mu, \sigma \in \mathbb{R}^{m n}$, let $\operatorname{Mat}(\mu) \in \mathbb{R}^{m \times n}$ denote a reshaped version of $\mu$. Then, we have:

$$
\begin{aligned}
\mathbb{E}[\lambda(w s(x)+b)] & =\frac{1}{2} \operatorname{Trace}\left(\operatorname{diag}\left(\operatorname{Mat}\left(\sigma^{2}\right) \tilde{s}(x)\right) \tilde{Q}\right)+\frac{1}{2}(\operatorname{Mat}(\mu) \tilde{s}(x))^{T} \tilde{Q}(\operatorname{Mat}(\mu) \tilde{s}(x)) \\
& +\kappa \prod_{i, j} \exp \left(\operatorname{Mat}(\mu)_{i j} \tilde{s}\left(x_{j}\right) \gamma_{i}+\frac{1}{2} \operatorname{Mat}\left(\sigma^{2}\right)_{i j}\left(\tilde{s}\left(x_{j}\right) \gamma_{i}\right)^{2}\right)
\end{aligned}
$$

MC dropout. Now assume a neural network with dropout: $\tilde{w}=\mu \odot \operatorname{Bernoulli}\left(p_{\text {dropout }}\right)$, where $\mu \in \mathbb{R}^{m n}$ denotes the weight in the absence of dropout and $p_{\text {dropout }} \in \mathbb{R}^{m n}$ denotes the probability of dropout. Then, we have:

$$
\begin{aligned}
\mathbb{E}[\lambda(w s(x)+b)] & =\frac{1}{2} \operatorname{Trace}\left(\operatorname{diag}\left(\operatorname{Mat}\left(\mu \odot p_{\text {dropout }} \odot\left(1-p_{\text {dropout }}\right)\right) \tilde{s}(x)\right) \tilde{Q}\right) \\
& +\frac{1}{2}\left(\operatorname{Mat}\left(\mu \odot p_{\text {dropout }}\right) \tilde{s}(x)\right)^{T} \tilde{Q}\left(\operatorname{Mat}\left(\mu \odot p_{\text {dropout }}\right) \tilde{s}(x)\right) \\
& +\kappa \prod_{i, j}\left(\operatorname{Mat}\left(p_{\text {dropout }}\right)_{i j} \exp \left(\operatorname{Mat}(\mu)_{i j} \tilde{s}\left(x_{j}\right) \gamma_{i}\right)+1-\operatorname{Mat}\left(p_{\text {dropout }}\right)_{i j}\right) .
\end{aligned}
$$

## C. 2 Expected-Softmax Optimization

We describe an algorithm to solve optimization problems of the form

$$
\max _{\ell \leq x \leq u} \mu^{T} \operatorname{softmax}(x)-\lambda^{T} x
$$

Our results will rely on the following lemma:

## Proposition 3. Consider the function

$$
f(x)=\frac{\sum_{i} \mu_{i} \exp \left(x_{i}\right)+D}{\sum_{j} \exp \left(x_{j}\right)+B}-\lambda^{T} x
$$

where $B \geq 0$ and $D=0$ if $B=0$. Let $r=\frac{D}{B}$ if $B>0$ and 0 otherwise. Define the set

$$
\Delta=\left\{\kappa \in \mathbb{R}:(\kappa-r)\left(\prod_{i=1}^{n}\left(\mu_{i}-\kappa\right)\right)-\sum_{i=1}^{n} \mu_{i}(1-r) \lambda_{i}\left(\prod_{j \neq i}\left(\mu_{j}-\kappa\right)\right)=0\right\}
$$

which is a set with at most $n+1$ elements. Define further

$$
\Delta_{f}= \begin{cases}\left\{\kappa \in \Delta: 0<\frac{\lambda}{\mu-\kappa}<1, \sum_{i=1}^{n} \frac{\lambda_{i}}{\mu_{i}-\kappa} \leq 1\right\} & \text { if } B>0 \\ \left.\kappa \kappa \Delta: 0<\frac{\lambda}{\mu-\kappa}<1, \sum_{i=1}^{n} \frac{\lambda_{i}}{\mu_{i}-\kappa}=1\right\} & \text { if } B=0\end{cases}
$$

Then, the set of stationary points of $f$ are given by

$$
\left\{\log \left(h\left(\frac{\lambda}{\mu-\kappa}\right)\right): \kappa \in \Delta_{f}\right\}
$$

where

$$
h(v)=\left\{\begin{array}{l}
\frac{B v}{1-\mathbf{1}^{T} v} \text { if } B>0 \\
\{\theta v: \theta>0\} \text { if } B=0
\end{array}\right.
$$

Proof. Differentiating with respect to $x_{i}$, we obtain

$$
\frac{\exp \left(x_{i}\right)}{\sum_{j} \exp \left(x_{j}\right)+B}\left(\mu_{i}-\left(\frac{\sum_{j} \mu_{j} \exp \left(x_{j}\right)+D}{\sum_{j} \exp \left(x_{j}\right)+B}\right)\right)-\lambda_{i}=p_{i}\left(\mu_{i}-\mu^{T} p-q\right)-\lambda_{i}
$$

where

$$
p_{i}=\frac{\exp \left(x_{i}\right)}{\sum_{j} \exp \left(x_{j}\right)+B}, q=\frac{D}{\sum_{j} \exp \left(x_{j}\right)+B} .
$$

If we set the derivative to 0 (to obtain a stationary point) we obtain the following coupled set of equations in $p, q, \kappa$ :

$$
\begin{aligned}
p_{i} & =\frac{\lambda_{i}}{\mu_{i}-\kappa} \quad i=1, \ldots, n \\
q & =r\left(1-\sum_{i} p_{i}\right) \\
\kappa & =\sum_{i} \mu_{i} p_{i}+q
\end{aligned}
$$

where $r=\frac{D}{B}$. We can solve this by first solving the scalar equation

$$
\kappa-r=\sum_{i} \frac{\mu_{i}(1-r) \lambda_{i}}{\mu_{i}-\kappa}
$$

for $\kappa$ (this is derived by adding up the first $n$ equations above weighted by $\mu_{i}$ and plugging in the value of $q$ ). This equation can be converted into a polynomial equation in $\kappa$

$$
(\kappa-r)\left(\prod_{i}\left(\mu_{i}-\kappa\right)\right)-\sum_{i} \mu_{i}(1-r) \lambda_{i}\left(\prod_{j \neq i}\left(\mu_{j}-\kappa\right)\right)=0
$$

which we can solve for all possible real solutions, denote this set $\Delta$. Note that this set has at most $n+1$ elements since it is the set of real solutions to a degree $n+1$ polynomial.

In order to recover $x$ from this, we first recall:

$$
p_{i}=\frac{\lambda_{i}}{\mu_{i}-\kappa}
$$

Since $p_{i}=\frac{\exp x_{i}}{\sum_{j} \exp \left(x_{j}\right)+B}$, we require that $p_{i} \in(0,1)$ and $\sum_{i} p_{i} \leq 1$ (with equality when $B=0$ and strict inequality when $B=1$. We thus filter $\Delta$ to the set of $\kappa$ that lead to $p$ satisfying these properties to obtain $\Delta_{f}$.
Once we have these, we are guaranteed that for each $\kappa \in \Delta_{f}$, we can define $p_{i}$ as above and solve for $x_{i}$ by solving the linear system of equations

$$
u_{i}=p_{i}\left(\sum_{j} u_{j}+B\right) \quad i=1,2, \ldots, n
$$

which can be solved as:

$$
u=B\left(I-p \mathbf{1}^{T}\right)^{-1} p=\frac{B p}{1-\mathbf{1}^{T} p}, x=\log (u)
$$

if $B>0$ since the matrix $I-p \mathbf{1}^{T}$ is strictly diagonally dominant and hence invertible, and we applied the Woodbury identity to compute the explicit inverse.
If $B=0$, we have $p=\operatorname{softmax}(x)$ and can recover $x$ as

$$
x=\log (p \theta)
$$

for any $\theta>0$.
The above lemma allows us to characterize all stationary points of the function

$$
\mu^{T} \operatorname{softmax}(x)-\lambda^{T} x
$$

when a subset of entries of $x$ are fixed to their upper or lower bounds, and we search for stationary points wrt the remaining free variables. Given this ability, we can develop an algorithm to globally optimize $\mu^{T}$ softmax $(x)-\lambda^{T} x$ subject to bound constraints by iterating over all possible $3^{n}$ configurations of binding constraints (each variable could be at its lower bound, upper bound or strictly between them). In this way, we are guaranteed to loop over all local optima, and by picking the one achieving the best objective value, we can guarantee that we have obtained the global optimum. The overall algorithm is presented in Algorithm 2 .
Proposition 4. Algorithm 2 finds the global optimum of the optimization problem

$$
\min _{x: \ell \leq x \leq u} \mu^{T} \operatorname{softmax}(x)-\lambda^{T} x
$$

and runs in time $O\left(n 3^{n}\right)$ where $n$ is the dimension of $x$.

```
Algorithm 2 Solving softmax layer problem via exhaustive enumeration
    Inputs: \(\lambda, \mu, \ell, u \in \mathbb{R}^{n}\)
    \(x^{\star} \leftarrow \ell\)
    \(f_{\text {opt }}(x) \leftarrow \mu^{T}\) softmax \((x)-\lambda^{T} x, f^{\star} \leftarrow f_{\text {opt }}\left(x^{\star}\right)\)
    for \(v \in[\text { Lower, Upper, Interior }]^{n}\) do
        nonbinding \([i] \leftarrow(v[i]==\) Interior \(), x_{i} \leftarrow\left\{\begin{array}{l}l[i] \text { if } v[i]=\text { Lower } \\ u[i] \text { if } v[i]=\text { Upper }\end{array} \quad\right.\) for \(i=1, \ldots, n\)
        \(B \leftarrow \sum_{i \text { such that nonbinding }[i]==\text { False }} \exp (x[i])\)
        \(D \leftarrow \sum_{i \text { such that nonbinding }[i]==\text { False }} \mu[i] \exp (x[i])\)
        Use proposition 3 to find the set of stationary points \(\mathcal{S}_{f}\) of the function
\[
f(x)=\frac{\sum_{i \text { such that nonbinding }[i]} \mu_{i} \exp \left(x_{i}\right)+D}{\sum_{j \text { such that nonbinding }[j]} \exp \left(x_{j}\right)+B}-\sum_{j \text { such that nonbinding }[j]} \lambda[j] x_{j}
\]
        for \(x^{s} \in \mathcal{S}_{f}\) do
            if \(x_{i}^{s} \in[\ell[i], u[i]] \quad \forall i\) s.t nonbinding \([i]\) then
            \(x_{i} \leftarrow x_{i}^{s} \quad \forall i\) s.t nonbinding \([i]\).
            if \(f_{\text {opt }}(x)>f^{\star}\) then
                \(x^{\star} \leftarrow x\)
                \(f^{\star} \leftarrow f_{\text {opt }}\left(x^{\star}\right)\)
            end if
            end if
        end for
    end for
    Return \(x^{\star}, f^{\star}\)
```


## C. 3 Input Layer with Linear-Exponential Multipliers

We recall the setting from Proposition 1 Let $\lambda_{1}(x)=\alpha^{T} x+\exp \left(\gamma^{T} x+\kappa\right), \lambda_{2}(x)=\beta^{T} x$, $g_{0}^{\star}=\max _{p_{0} \in \mathcal{P}_{0}} g_{0}\left(p_{0}, \lambda_{1}\right)$, and $g_{1}^{\star}=\max _{x \in \mathcal{X}_{1}} g_{0}\left(x, \lambda_{1}, \lambda_{2}\right)$.

Proposition 5. In the setting described above, and with s as the element-wise activation function:

$$
\begin{aligned}
& g_{0}^{\star} \leq \alpha^{T}(w \mu+b)+\exp \left(\frac{\left\|w^{T} \gamma\right\|^{2} \sigma^{2}}{2}+\gamma^{T} b+\kappa\right), \\
& g_{1}^{\star} \leq \max _{\substack{x \in \mathcal{X}_{2} \\
z=s(x)}} \beta^{T}\left(w_{2} z+b_{2}\right)-\alpha^{T} x-\exp \left(\gamma^{T} x+\kappa\right) .
\end{aligned}
$$

The maximization in the second equation can be bounded by solving the following convex optimization problem:

$$
\begin{aligned}
\min _{\eta \in \mathbb{R}^{n}, \zeta \in \mathbb{R}_{+}} & \zeta(\log (\zeta)-1-\kappa)+\mathbf{1}^{T} \max \left(\left(\eta+w_{2}^{T} \beta\right) \odot s\left(l_{2}\right),\left(\eta+w_{2}^{T} \beta\right) \odot s\left(u_{2}\right)\right) \\
& +\sum_{i} s^{\star}\left(\alpha_{i}+\zeta \gamma_{i}, \eta_{i}, l_{2 i}, u_{2 i}\right)
\end{aligned}
$$

where $s^{\star}(a, b, c, d)=\max _{z \in[c, d]}-a z-b s(z)$.

Proof.

$$
\begin{aligned}
\max _{x \in \mathcal{X}}^{z=s(x)} & \beta^{T}\left(w_{2} z+b_{2}\right)-\alpha^{T} x-\exp \left(\alpha^{T} x+\kappa\right), \\
\leq & \min _{\eta} \max _{x \in \mathcal{X}_{2}, z \in s\left(\mathcal{X}_{2}\right)} \eta^{T}(z-s(x))+\beta^{T}\left(w_{2} z+b_{2}\right)-\alpha^{T} x-\exp \left(\gamma^{T} x+\kappa\right), \\
\leq & \min _{\eta, \zeta} \max _{x \in \mathcal{X}_{2}, t} \mathbf{1}^{T} \max \left(\left(\eta+w_{2}^{T} \beta\right) \odot s\left(l_{2}\right),\left(\eta+w_{2}^{T} \beta\right) \odot s\left(u_{2}\right)\right)+\beta^{T} b_{2}-\alpha^{T} x \\
& -\eta^{T} s(x)-\exp (t)+\zeta\left(t-\gamma^{T} x-\kappa\right), \\
\leq & \min _{\eta, \zeta} \zeta(\log (\zeta)-1-\kappa)+\mathbf{1}^{T} \max \left(\left(\eta+w_{2}^{T} \beta\right) \odot s\left(l_{2}\right),\left(\eta+w_{2}^{T} \beta\right) \odot s\left(u_{2}\right)\right) \\
& \quad+\max _{l_{2} \leq x \leq u_{2}}-(\alpha+\zeta \gamma)^{T} x-\eta^{T} s(x), \\
\leq & \min _{\eta, \zeta} \zeta(\log (\zeta)-1-\kappa)+\mathbf{1}^{T} \max \left(\left(\eta+w_{2}^{T} \beta\right) \odot s\left(l_{2}\right),\left(\eta+w_{2}^{T} \beta\right) \odot s\left(u_{2}\right)\right) \\
& +\sum_{i} \max _{l_{2 i} \leq x_{i} \leq u_{2 i}}-\left(\alpha_{i}+\zeta \gamma_{i}\right) x_{i}-\eta_{i} s\left(x_{i}\right) .
\end{aligned}
$$

## C. 4 Inner Problem with Linear Multipliers

In its general form, the objective function of the inner maximization problem can be expressed as:

$$
\begin{equation*}
g_{k}\left(x_{k}, \lambda_{k}, \lambda_{k+1}\right)=\underset{y \sim \pi_{k}\left(x_{k}\right)}{\mathbb{E}}\left[\lambda_{k+1}(y)\right]-\lambda_{k}\left(x_{k}\right) \tag{11}
\end{equation*}
$$

We assume that the layer is in the form $y=W s(x)+b$, where $W$ and $b$ are random variables and $s$ is an element-wise activation function. Then we can rewrite $g_{k}\left(x_{k}, \lambda_{k}, \lambda_{k+1}\right)$ as:

$$
\begin{equation*}
g_{k}\left(x_{k}, \lambda_{k}, \lambda_{k+1}\right)=\underset{W, b}{\mathbb{E}}\left[\lambda_{k+1}\left(W \max \left\{x_{k}, 0\right\}+b\right)\right]-\lambda_{k}\left(x_{k}\right) \tag{12}
\end{equation*}
$$

We now use the assumption that $\lambda_{k+1}$ is linear: $\lambda_{k+1}: y \mapsto \theta_{k+1}^{\top} y$. Then the problem can equivalently be written as:

$$
\begin{align*}
g_{k}\left(x_{k}, \lambda_{k}, \lambda_{k+1}\right) & =\underset{W, b}{\mathbb{E}}\left[\theta_{k+1}^{\top}\left(W s\left(x_{k}\right)+b\right)\right]-\theta_{k}^{T} x_{k}, \\
& =\left(\mathbb{E}[W]^{\top} \theta_{k+1}\right)^{\top} s\left(x_{k}\right)+\theta_{k+1}^{\top} \mathbb{E}[b]-\theta_{k}^{T} x_{k}  \tag{13}\\
& =\theta_{k+1}^{T} \mathbb{E}[b]+\sum_{i}\left(\mathbb{E}[W]^{\top} \theta_{k+1}\right)_{i} s\left(x_{i}\right)-\left(\theta_{k}\right)_{i} x_{i} .
\end{align*}
$$

Maximizing the RHS subject to $l \leq x \leq u$, we obtain:

$$
\theta_{k+1}^{T} \mathbb{E}[b]+\sum_{i} \max _{z \in\left[l_{i}, u_{i}\right]}\left(\mathbb{E}[W]^{\top} \theta_{k+1}\right)_{i} s(z)-\left(\theta_{k}\right)_{i} z
$$

where the maximization over $z$ can be solved in closed form for most common activation functions $s$ as shown in Dvijotham et al. [2018b].
So we can simply apply the deterministic algorithm to compute the closed-form solution of this problem.

## D Relationship to Prior work

We establish connections between the functional Lagrangian framework and prior work on deterministic verification techniques based on Lagrangian relaxations and SDP relaxations.

## D. 1 Lagrangian Dual Approach

We assume that the network layers are deterministic layers of the form:

$$
\begin{align*}
\forall k, k \text { is odd } \pi_{k}(x) & =w_{k} x+b_{k} \\
\forall k, k \text { is even } \pi_{k}(x) & =s(x) \tag{14}
\end{align*}
$$

where $s$ is an element-wise activation function and that the specification can be written as:

$$
\begin{equation*}
\boldsymbol{\psi}\left(x_{K}\right)=c^{T} x_{K} \tag{15}
\end{equation*}
$$

Proposition 6 (Linear Multipliers). For a verification problem described by equations (14, 15), the functional Lagrangian framework with linear functional multipliers $\lambda_{k}(x)=\theta_{k}^{T} x$ is equivalent to the Lagrangian dual approach from Dvijotham et al. [2018b].

Proof. The final layer problem is

$$
\max _{x_{K}} c^{T} x_{K}-\theta_{K}^{T} x_{K}=\mathbf{1}^{T} \max \left(\left(c-\theta_{K}\right) \odot l_{K},\left(c-\theta_{K}\right) \odot u_{K}\right)
$$

For even layers with $k<K$, the optimization problem is

$$
\begin{aligned}
& \max _{x \in\left[l_{k}, u_{k}\right]} \theta_{k+1}^{T}\left(w_{k} x+b_{k}\right)-\theta_{k}^{T} x=\theta_{k+1}^{T} b_{k}+\left(w_{k}^{T} \theta_{k+1}-\theta_{k}\right)^{T} x \\
& =\mathbf{1}^{T} \max \left(\left(w_{k}^{T} \theta_{k+1}-\theta_{k}\right) \odot l_{k},\left(w_{k}^{T} \theta_{k+1}-\theta_{k}\right) \odot u_{k}\right)+\theta_{k+1}^{T} b_{k}
\end{aligned}
$$

For odd layers with $k<K$, the optimization problem is

$$
\max _{x \in\left[l_{k}, u_{k}\right]} \theta_{k+1}^{T} s(x)-\theta_{k}^{T} x=\sum_{i} \max _{z \in\left[l_{k i}, u_{k i}\right]}\left(\theta_{k+1}\right)_{i} s(z)-\left(\theta_{k}\right)_{i} z
$$

All these computations precisely match those from Dvijotham et al. [2018b], demonstrating the equivalence.

## D. 2 SDP-cert

We assume that the network layers are deterministic layers of the form:

$$
\begin{equation*}
\forall k \pi_{k}(x)=\operatorname{ReLU}\left(w_{k} x+b_{k}\right) \tag{16}
\end{equation*}
$$

where $s$ is an element-wise activation function and that the specification can be written as:

$$
\begin{equation*}
\boldsymbol{\psi}\left(x_{K}\right)=c^{T} x_{K} \tag{17}
\end{equation*}
$$

Proposition 7 (Quadratic Multipliers). For a verification problem described by equations (16, 17), the optimal value of the Functional Lagrangian Dual with

$$
\begin{aligned}
& \lambda_{k}(x)=q_{k}^{T} x+\frac{1}{2} x^{T} Q_{k} x \quad k=1, \ldots, K-1 \\
& \lambda_{K}(x)=q_{K}^{T} x
\end{aligned}
$$

and when an SDP relaxation is used to upper bound the inner maximization problems over $g_{k}$, is equal to the dual of the SDP relaxation from Raghunathan et al. [2018].

Proof. With quadratic multipliers of the form $\lambda_{k}(x)=q_{k}^{T} x+\frac{1}{2} x^{T} Q_{k} x \quad k=1, \ldots, K-2$ and $\lambda_{K}(x)=q_{K}^{T} x$, the inner maximization problems for the intermediate layers are of the form:

$$
\max _{\substack{x \in[l, u] \\ y=\operatorname{ReLU}(w x+b)}} \tilde{q}^{T} y+\frac{1}{2} y^{T} \tilde{Q} y-q^{T} x-\frac{1}{2} x^{T} Q x
$$

where $l=l_{k}, u=u_{k}, \tilde{q}=q_{k+1}, \tilde{Q}=Q_{k+1}, q=q_{k}, Q=Q_{k}, w=w_{k}, b=b_{k}$. Let $x \in \mathbb{R}^{n}, y \in$ $\mathbb{R}^{m}$ ( $n$ dimensional input, $m$ dimensional output of the layer). Further, let $\tilde{l}=l_{k+1}, \tilde{u}=u_{k+1}$.
We can relax the above optimization problem to the following Semidefinite Program (SDP) (following Raghunathan et al. [2018]):

$$
\begin{gather*}
\max _{P} \tilde{q}^{T} P[y]+\frac{1}{2} \operatorname{Trace}\left(\tilde{Q} P\left[y y^{T}\right]\right)-q^{T} P[x]-\frac{1}{2} \operatorname{Trace}\left(Q P\left[x x^{T}\right]\right)  \tag{18a}\\
\text { Subject to } P=\left(\begin{array}{ccc}
1 & (P[y])^{T} & (P[x])^{T} \\
P[y] & P\left[y y^{T}\right] & P\left[x y^{T}\right] \\
P[x] & \left(P\left[x y^{T}\right]\right)^{T} & P\left[x x^{T}\right]
\end{array}\right) \in \mathbb{S}^{n+m+1},  \tag{18b}\\
P \succeq 0,  \tag{18c}\\
\quad \operatorname{diag}\left(P\left[x x^{T}\right]-l(P[x])^{T}-P[x] u^{T}+l u^{T}\right) \leq 0,  \tag{18d}\\
\quad \operatorname{diag}\left(P\left[y y^{T}\right]-\tilde{l}(P[y])^{T}-P[y] \tilde{u}^{T}+\tilde{l}^{T} \tilde{u}^{T}\right) \leq 0,  \tag{18e}\\
P[y] \geq 0, P[w] \geq w P[x],  \tag{18f}\\
 \tag{18~g}\\
\operatorname{diag}\left(w P\left[x y^{T}\right]\right)+P[y] \odot b=\operatorname{diag}\left(P\left[y y^{T}\right]\right) .
\end{gather*}
$$

where the final constraint follows from the observation that $y \odot(y-w x-b)=0$.
The above optimization problem resembles the formulation of Raghunathan et al. [2018] except that it only involves two adjacent layers rather than all the layers at once. Let $\Delta_{k}$ denote the feasible set given the constraints in the above optimization problem. Then, the formulation of Raghunathan et al. [2018] can be written as:

$$
\begin{align*}
\max & c^{T} y_{K}  \tag{19a}\\
\text { subject to } & P_{k} \in \Delta_{k} \quad k=0, \ldots, K-1, l_{K} \leq y_{K} \leq u_{K}  \tag{19b}\\
& P_{k+1}\left[x x^{T}\right]=P_{k}\left[y y^{T}\right] \quad k=0, \ldots, K-2  \tag{19c}\\
& P_{k+1}[x]=P_{k}[y] \quad k=0, \ldots, K-2  \tag{19d}\\
& y_{K}=P_{K-1}[y] . \tag{19e}
\end{align*}
$$

Note that in Raghunathan et al. [2018], a single large $P$ matrix is used whose block-diagonal subblocks are $P_{k}$ and the constraint $P \succeq 0$ is enforced. Due to the matrix completion theorem for SDPs Grone et al. 1984, Vandenberghe and Andersen, 2015, it suffices to ensure postitive semidefiniteness of the sub-blocks rather than the full $P$ matrix.

Dualizing the last three sets of constraints above with Lagrangian multipliers $\Theta_{k} \in \mathbb{S}^{n_{k}+n_{k+1}+1}, \theta_{k} \in$ $\mathbb{R}^{n_{k}}$ and $\theta_{K} \in \mathbb{R}^{n_{K}}$, we obtain the following optimization problem:

$$
\begin{aligned}
\max & c^{T} y_{K}+\theta_{K}^{T}\left(-P_{K-1}[y]+y_{K}\right)+\sum_{k=0}^{K-2} \operatorname{Trace}\left(\Theta_{k}\left(P_{k+1}\left[x x^{T}\right]-P_{k}\left[y y^{T}\right]\right)\right) \\
& +\sum_{k=0}^{K-2} \theta_{k}^{T}\left(P_{k+1}[x]-P_{k}[y]\right)
\end{aligned}
$$

subject to $P_{k} \in \Delta_{k} \quad k=0, \ldots, K-1$,

$$
l_{K} \leq y_{K} \leq u_{K}
$$

The objective decomposes over $P_{k}$ and can be rewritten as:

$$
\begin{align*}
\max _{l_{K} \leq y_{K} \leq u_{K}} & \left(c+\theta_{K}\right)^{T} y_{K}+\sum_{k=0}^{K-1} \max _{P_{k} \in \Delta_{k}} \operatorname{Trace}\left(\Theta_{k-1} P_{k}\left[x x^{T}\right]\right)+\theta_{k-1}^{T} P_{k}[x]-\theta_{k}^{T} P_{k}[y]  \tag{20}\\
& -\operatorname{Trace}\left(\Theta_{k} P_{k}\left[y y^{T}\right]\right)
\end{align*}
$$

with the convention that $\Theta_{K-1}=0, \Theta_{-1}=0, \theta_{-1}=0$. If we set $Q_{k}=-\Theta_{k-1}, q_{k}=-\theta_{k-1}$ for $k=1, \ldots, K$, then the optimization over $P_{k}$ precisely matches the optimization in (18). Further, since $\lambda_{K}$ is linear, the final layer optimization simply reduces to:

$$
\max _{l_{K} \leq x_{K} \leq u_{K}} c^{T} x_{K}-q_{K}^{T} x_{K},
$$

which matches the first term in 20 .

Thus, the functional Lagrangian framework with quadratic multipliers $\lambda_{k}$ for $k=1, \ldots, K-2$ and a linear multiplier for $\lambda_{K}$ precisely matches the Lagrangian dual of $(19)$ and since $(19)$ is a convex optimization problem, strong duality guarantees that the optimal values must coincide.

## E Additional Experimental Details

## E. 1 Robust OOD Detection on Stochastic Neural Networks

Inner Optimization. All inner problems have a closed-form as shown in section C.4 except for the last one, which is handled as follows.

The last inner problem can be formulated as:

$$
\begin{equation*}
\max _{x_{K} \in \mathcal{X}_{K}} \mu^{\top} \operatorname{softmax}\left(x_{K}\right)+\nu^{\top} x_{K}, \tag{21}
\end{equation*}
$$

where $\mu$ is a one-hot encoded vector and $\nu$ is a real-valued vector.

- Projected Gradient Ascent (Training):
- Hyper-parameters: we use the Adam optimizer Kingma and Ba [2015], with a learningrate of 1.0 and a maximum of 1000 iterations.
- Stopping criterion: when all coordinates have either zero gradient, or are at a boundary with the gradient pointing outwards of the feasible set.
- In order to help the gradient method find the global maximum, we use a heuristic for initialization, which consists of using the following two starting points for the maximization (and then to take the best of the corresponding two solutions found):

1. Ignore affine part $(\nu=0)$, which gives a solution in closed form: set $x_{K}$ at its upper bound at the coordinate where $\mu$ is 1 , and at its lower bound elsewhere.
2. Ignore softmax part $(\mu=0)$, which also gives a solution in closed form: set $x_{K}$ at its upper bound at the coordinates where $\nu \geq 0$, and at its lower bound elsewhere.

- Evaluation: we use Algorithm 2 at evaluation time, which solves the maximization exactly.

Outer Optimization. We use the Adam optimizer, with a learning-rate that is initialized at 0.001 and divided by 10 every 250 steps. We run the optimization for a total of 1000 steps.

Gaussian-MLP. We use the ReLU MLP from Wicker et al. 2020] that consists of 2 hidden layers of 128 units each. The models are available at https://github.com/matthewwicker/ ProbabilisticSafetyforBNNs.

LeNet. We use the LeNet5 architecture with dropout applied to the last fully connected layer with a probability of 0.5 . To make the bound-propagation simpler, we do not use max-pooling layers and instead increase the stride of convolutions.

VGG-X. For VGG-X (where $\mathrm{X} \in\{2,4,8,16,32,64\}$ ), the architecture can be described as:

- Conv $3 \times 3$, X filters, stride 1
- ReLU
- Conv $3 \times 3$, X filters, stride 2
- ReLU
- Conv $3 \times 3$, 2 X filters, stride 2
- ReLU
- Conv $3 \times 3$, 2 X filters, stride 2
- ReLU
- Flatten
- Linear with 128 output neurons
- Dropout with rate 0.2
- Linear with 10 output neurons

Hardware The verification of each sample is run on a CPU with 1-2 cores (and on each instance, BP and FL are timed on the same exact hardware configuration).

## E. 2 Adversarial Robustness for Stochastic Neural Networks

Inner Optimization. We use a similar approach as in Appendix E.1. For the final inner problem (corresponding to the objective which is a linear function of the softmax and the layer inputs), we run projected gradient ascent during the optimization phase and then use Algorithm 2 to solve the maximization exactly. For projected gradient ascent, because of the non-convexity of the problem, we use the following heuristics to try and find the global maximum:

- Black-box attack (1st phase): we use the Square adversarial attack Andriushchenko et al. [2020], with 600 iterations, 300 random restarts and learning-rate of 0.1 .
- Fine-tuning (2nd phase): We then choose the best attack from the restarts, and employ projected gradient ascent, with a learning-rate of 0.1 and 100 iterations to fine-tune further.

Model Parameters. We use the 1 and 2 layer ReLU MLPs from [Wicker et al., 2020]. The models are available at https://github.com/matthewwicker/ProbabilisticSafetyforBNNs.

Outer Optimization. We use the Adam optimizer, with a learning-rate that is initialized at 0.001 and divided by 10 every 1000 steps. We run the optimization for a total of 3000 steps.

Hardware All experiments were run on a P100 GPU.

## E. 3 Distributionally Robust OOD Detection

Model. We train networks on MNIST using the code fromhttps://gitlab.com/Bitterwolf/ GOOD with the CEDA method, and with the default hyperparameters. We train a CNN with ReLU activations the following layers:

- Conv $4 \times 4$, 16 filters, stride 2 , padding 2 on both sides
- ReLU
- Conv $4 x 4$, 32 filters, stride 1 , padding 1 on both sides
- Relu
- Flatten
- Linear with 100 output neurons
- Relu
- Linear with 10 output neurons

Outer Optimization. For the outer loop of the verification procedure, we use Adam for 100k steps. The learning-rate is initially set to 0.0001 and then divided by 10 after 60 k and 80 k steps.

Hardware We run the experiments for this section on a CPU with 2-4 cores.

## F Additional Results with Interval Bound Propagation for Bilinear Operations

## F. 1 Robust OOD Detection for Stochastic Neural Networks

We repeat the experiments in Section 5.1 where we use IBP to handle bound-propagation through the layers where bilinear propagation is required (because of bounds coming from both the layer

Table 4: Robust OOD Detection: MNIST vs EMNIST (MLP and LeNet) and CIFAR-10 vs CIFAR100 (VGG-*). BP: Bound-Propagation (baseline), using IBP instead of Bunel et al. [2020] for bilinear operations; FL: Functional Lagrangian (ours). The reported times correspond to the median of the 500 samples.

| OOD Task | Model | \#neurons | \#params | $\epsilon$ | Time (s) |  | GAUC (\%) |  | AAUC (\%) |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  |  |  | BP | FL | BP | FL |  |
| (E)MNIST | MLP | 256 | 2k | 0.01 | 1.1 | +13.1 | 55.4 | 67.5 | 86.9 |
|  |  |  |  | 0.03 | 1.2 | +13.4 | 38.7 | 54.5 | 88.6 |
|  |  |  |  | 0.05 | 1.3 | +17.7 | 19.1 | 36.0 | 88.8 |
| (E)MNIST | LeNet | 0.3M | 0.1M | 0.01 | 50.1 | +13.1 | 0.0 | 28.4 | 71.6 |
|  |  |  |  | 0.03 | 54.7 | +13.7 | 0.0 | 11.7 | 57.6 |
|  |  |  |  | 0.05 | 79.4 | +24.8 | 0.0 | 2.3 | 44.0 |
| CIFAR | VGG-16 | 3.0M | 83k | 0.001 | 426.4 | +21.4 | 0.0 | 21.7 | 60.9 |
|  | VGG-32 | 5.9 M | 0.2M | 0.001 | 1035.2 | +21.3 | 0.0 | 23.8 | 64.7 |
|  | VGG-64 | 11.8M | 0.5M | 0.001 | 8549.7 | +42.1 | 0.0 | 28.6 | 67.4 |

Table 5: Adversarial Robustness for different BNN architectures trained on MNIST from Wicker et al. [2020]. BP: Bound-Propagation (baseline), using IBP instead of LBP for bilinear operations; FL: Functional Lagrangian (ours). The accuracy reported for FL and BP is the \% of samples we can certify as robust with probability 1 . For each model, we report results for the first 500 test-set samples.

| \#layers | $\epsilon$ | \#neurons | BP Acc. (\%) | FL Acc. (\%) | BP Time (s) | FL Time (s) | Adv Acc (\%) |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | 0.025 | 128 | 43.8 | $\mathbf{6 5 . 2}$ | 1.3 | +353.3 | 82.6 |
|  |  | 256 | 40.6 | $\mathbf{6 4 . 6}$ | 1.4 | +431.3 | 82.6 |
|  |  | 512 | 35.0 | $\mathbf{5 7 . 0}$ | 1.3 | +357.1 | 82.8 |
| 2 | 0.001 | 256 | 29.4 | $\mathbf{3 6 . 9}$ | 1.6 | +43.6 | 79.4 |
|  |  | 512 | 46.0 | $\mathbf{6 3 . 4}$ | 1.7 | +433.8 | 89.2 |
|  |  | 1024 | 18.4 | $\mathbf{1 9 . 6}$ | 1.6 | +440.9 | 74.8 |

inputs and the layer parameters due to the stochasticity of the model) instead of Bunel et al. [2020]. Bunel et al. [2020] usually results in significantly tighter bounds compared to IBP but we note that for MNIST-CNN and CIFAR-CNN, we expect IBP to perform competitively as the bilinear bound propagation is only applied for a single layer (dropout). The results are presented in Table 4 and we find that even while using IBP as the bound-propagation method, our framework provides significantly stronger guarantees.

## F. 2 Adversarial Robustness for Stochastic Neural Networks

For the verification tasks considered in Section 5.2, we use IBP instead of the tighter LBP as the bound-propagation method and report results in Table 2. We find that, similar to Section5.2, our framework is able to significantly improve on the guarantees the bound-propagation baseline is able to provide.

## Checklist

1. For all authors...
(a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] Our theoretical framework is explained in Section 3 and our empirical results are detailed in Section 5 .
(b) Did you describe the limitations of your work? [Yes] See Section 6 .
(c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section6
(d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
(a) Did you state the full set of assumptions of all theoretical results? [Yes] Formal statements are available in Section 3 and Appendix C.
(b) Did you include complete proofs of all theoretical results? [Yes] Formal proofs are available in Section $B$ and Appendix $C$.
3. If you ran experiments...
(a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] We have included the URL at which the code will be released.
(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] We provide experimental details in Appendix E .
(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] We follow standard practice in the neural network verification community, and our evaluation is performed on hundreds of samples to ensure that the results are statistically significant.
(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] We have included all runtimes and hardware platforms used.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
(a) If your work uses existing assets, did you cite the creators? [Yes]
(b) Did you mention the license of the assets? [N/A]
(c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]
