CS781 Quiz 1 (Autumn 2025)

Max marks: 20 Duration: 60 mins Roll No.

- *You are required to answer each question only in the space provided with each question.*
- *Only material written within the allotted answering space for each question will be graded.*
- *The spaces allotted for answering questions should give a rough indication of the relative lengths of correct answers to the questions.*
- *Please attach all your rough sheets.*
- *The exam is open book and notes. However, you are not allowed to search on the internet or consult others over the internet for your answers.*
- *Be brief, complete and stick to what has been asked.*
- *Unless asked for explicitly, you may cite results/proofs covered in class without reproducing them.*
- ***If you need to make any assumptions, state them clearly.***
- ***Do not copy solutions from others. Penalty for offenders: FR grade.***

1. *[10 marks]* A key step in the β-CROWN algorithm is the replacement of all split constraints, say $\boldsymbol{\mathcal{Z}}$, at a layer of ReLUs by a linear term $\boldsymbol{\beta^T S z}$, where $\boldsymbol{z}$ denotes an n-dimensional vector of pre-activation values, $\boldsymbol{S}$ is a $n \times n$ diagonal matrix, and $\boldsymbol{\beta}$ is a n-dimensional vector with all components ≥ 0.

   Recall that $S_{j,j} = -1$ if $z_j \geq 0$ is in $\boldsymbol{\mathcal{Z}}$, $S_{j,j} = +1$ if $z_j < 0$ is in $\boldsymbol{\mathcal{Z}}$, and $S_{j,j} = 0$ otherwise.

   A student claims that the β-CROWN algorithm would continue to work with a simplified diagonal matrix $\boldsymbol{R}$, defined as follows: $R_{j,j} = -1$ if $z_j \geq 0$ is in $\boldsymbol{\mathcal{Z}}$, and $R_{j,j} = 0$ otherwise.

   The soundness of the student's claim rests on whether the following inequality holds, where $\boldsymbol{A}$ is an $m \times n$ matrix, $\boldsymbol{b}$ is an m-dimenstional vector, $\mathcal{C}$ denotes an arbitrary region in the space of inputs $\boldsymbol{x}$, and $\boldsymbol{z} \models \boldsymbol{\mathcal{Z}}$ denotes that the components of $\boldsymbol{z}$ satisfy all constraints in $\boldsymbol{\mathcal{Z}}$:

$$\min_{\boldsymbol{x} \in \mathcal{C}, \boldsymbol{z} \models \boldsymbol{\mathcal{Z}}} \boldsymbol{Az} + \boldsymbol{b} \quad \geq \quad \max_{\boldsymbol{\beta} \geq \boldsymbol{0}} \min_{\boldsymbol{x} \in \mathcal{C}} \boldsymbol{Az} + \boldsymbol{\beta^T R z} + \boldsymbol{b}$$

   Either give a proof that the above inequality above holds, or provide a counter-example.

The inequality indeed holds.
Let $\vec{x}^*$ be the value of $\vec{x}$ in C that minimizes $A\vec{z} + b$, while ensuring $\vec{z} \models \mathcal{Z}$.
Let $\vec{z}^*$ be the corresponding value of $\vec{z}$.

Let $P = \{ i \mid z_i^* \geq 0 \}$, i.e., P is the set of indices of all those components of $\vec{z}^*$ that are non-negative.

Since $\vec{z}^* \models \mathcal{Z}$, it follows that for every j s.t. $z_j \geq 0$ is in $\mathcal{Z}$, $j \in P$.

By defn. of R, if $R_{jj} = -1$, then $z_j \geq 0$ is in $\mathcal{Z}$, and by above reasoning, $j \in P$

$$\therefore \quad \sum R_{jj} z_j^* \leq 0$$

$\therefore$ For every $\vec{\beta} \geq \vec{0}$, $\vec{\beta}^T . R \vec{z}^* \leq 0$

Hence, $\quad A\vec{z}^* + \vec{\beta}^T . R . \vec{z}^* + b \leq A\vec{z}^* + b.$

Since we know that when $\vec{x} = \vec{x}^*$, we get $\vec{z} = \vec{z}^*$, it follows that:

$$\min_{\vec{x} \in C} \left(A\vec{z} + \vec{\beta}^T R \vec{z} + b \right) \leq A\vec{z}^* + \vec{\beta}^T R \vec{z}^* + b \leq A\vec{z}^* + b = \min_{\substack{\vec{x} \in C \\ \vec{z} \models \mathcal{Z}}} (A\vec{z} + b)$$

Since the above holds for all $\vec{\beta}^T \geq \vec{0}$, we get

$$\max_{\vec{\beta} > \vec{0}} \; \min_{\vec{x} \in C} \left(A\vec{z} + \vec{\beta}^T R \vec{z} + b \right) \leq \min_{\substack{\vec{x} \in C \\ \vec{z} \in \mathcal{Z}}} \left(A\vec{z} + b \right)$$

2. *[10 marks]* You are given a decision tree as shown below. Find a set-minimal abductive explanation of the decision "NOK" corresponding to the feature values $a = 0, b = 1, c = 0, d = 1$. You must show all steps in your derivation of the minimal abductive explanation.

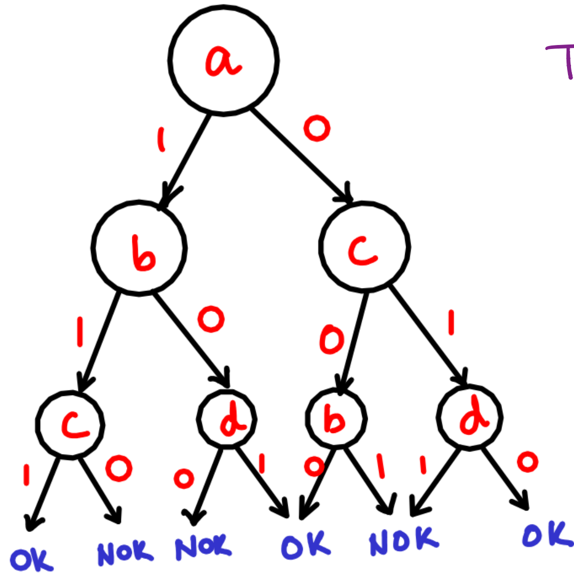

Figure 1: A decision tree

There can be multiple correct answers to this question. One of them is shown below.

Step 1: Is the empty set an AxP?

NO: $a = b = c = 0, d = 1$ yields OK.

Step 2: Difference features: $\{b\}$

Step 3: Minimal Hitting set of $\{b\} = \{b\}$.

Step 4: Is $\{b\}$ an AxP?

NO: $a = 0, b = c = d = 1$ yields OK.

Step 5: Difference features: $\{c\}$

Step 6: Minimal hitting set of $\{b\}, \{c\} = \{b, c\}$

Step 7: Is $\{b, c\}$ an AxP?

Yes it is: Regardless of value of a, d, if $b = 1$ and $c = 0$, the decision tree always yields NOK