A. Proof of Theorem 1

Given a point cloud T and its perturbed version T^*, we define the following two random variables:

$$\mathbf{W} = S_k(T),\ \mathbf{Z} = S_k(T^*), \tag{13}$$

where $\mathbf{W}$ and $\mathbf{Z}$ represent the random 3D point clouds with k points subsampled from T and T^* uniformly at random without replacement, respectively. We use Φ to denote the joint space of $\mathbf{W}$ and $\mathbf{Z}$, where each element is a 3D point cloud with k points subsampled from T or T^*. We denote by E the set of intersection points between T and T^*, i.e., $E = T \cap T^*$.

Before proving our theorem, we first describe a variant of the Neyman-Pearson Lemma [21] that will be used in our proof. The variant is from [9].

Lemma 1 (Neyman-Pearson Lemma)**.** *Suppose $\mathbf{W}$ and $\mathbf{Z}$ are two random variables in the space Φ with probability distributions σ_w and σ_z, respectively. Let $H : \Phi \to \{0, 1\}$ be a random or deterministic function. Then, we have the following:*

- *If $Q_1 = \{\varphi \in \Phi : \sigma_w(\varphi) > \zeta \cdot \sigma_z(\varphi)\}$ and $Q_2 = \{\varphi \in \Phi : \sigma_w(\varphi) = \zeta \cdot \sigma_z(\varphi)\}$ for some $\zeta > 0$. Let $Q = Q_1 \cup Q_3$, where $Q_3 \subseteq Q_2$. If we have $Pr(H(\mathbf{W}) = 1) \geq Pr(\mathbf{W} \in Q)$, then $Pr(H(\mathbf{Z}) = 1) \geq Pr(\mathbf{Z} \in Q)$.*

- *If $Q_1 = \{\varphi \in \Phi : \sigma_w(\varphi) < \zeta \cdot \sigma_z(\varphi)\}$ and $Q_2 = \{\varphi \in \Phi : \sigma_w(\varphi) = \zeta \cdot \sigma_z(\varphi)\}$ for some $\zeta > 0$. Let $Q = Q_1 \cup Q_3$, where $Q_3 \subseteq Q_2$. If we have $Pr(H(\mathbf{W}) = 1) \leq Pr(\mathbf{W} \in Q)$, then $Pr(H(\mathbf{Z}) = 1) \leq Pr(\mathbf{Z} \in Q)$.*

Proof. Please refer to [9]. □

Next, we formally prove our Theorem 1. Our proof is inspired by previous work [10, 9]. Roughly speaking, the idea is to derive the label probability lower and upper bounds via computing the probability of random variables in certain regions crafted by the variant of Neyman-Pearson Lemma. However, due to the difference in sampling methods, our space divisions are significantly different from previous work [9]. Recall that we denote $p_i = \Pr(f(\mathbf{W}) = i)$ and $p_i^* = \Pr(f(\mathbf{Z}) = i)$, where $i \in \{1, 2, \cdots, c\}$. We denote $y = \text{argmax}_{i=\{1,2,\ldots,c\}}\, p_i$. Our goal is to find the maximum r^* such that $y = \text{argmax}_{i=\{1,2,\cdots,c\}}\, p_i^*$, i.e., $p_y^* > p_e^* = max_{i \neq y} p_i^*$, for $\forall T^* \in \Gamma(T, r^*)$. Our key step is to derive a lower bound of p_y^* and an upper bound of $p_e^* = max_{i\neq y} p_i^*$ via Lemma 1. Given these probability bounds, we can find the maximum r^* such that the lower bound of p_y^* is larger than the upper bound of p_e^*.

Dividing the space Φ: We first divide the space Φ into three regions which are as follows:

$$\Delta_T = \{\varphi \in \Phi | \varphi \subseteq T, \varphi \nsubseteq E\}, \tag{14}$$
$$\Delta_{T^*} = \{\varphi \in \Phi | \varphi \subseteq T^*, \varphi \nsubseteq E\}, \tag{15}$$
$$\Delta_E = \{\varphi \in \Phi | \varphi \subseteq E\}, \tag{16}$$

where Δ_E consists of the subsampled point clouds that can be obtained by subsampling k points from E; and Δ_T (or Δ_{T^*}) consists of the subsampled point clouds that are subsampled from T (or T^*) but do not belong to Δ_E. Since $\mathbf{W}$ and $\mathbf{Z}$, respectively, represent the random 3D point clouds with k points subsampled from T and T^* uniformly at random without replacement, we have the following probability mass functions:

$$\Pr(\mathbf{W} = \varphi) = \begin{cases} \frac{1}{\binom{n}{k}}, & \text{if } \varphi \in \Delta_T \cup \Delta_E, \\ 0, & \text{otherwise}, \end{cases} \tag{17}$$

$$\Pr(\mathbf{Z} = \varphi) = \begin{cases} \frac{1}{\binom{t}{k}}, & \text{if } \varphi \in \Delta_{T^*} \cup \Delta_E, \\ 0, & \text{otherwise}, \end{cases} \tag{18}$$

where t is the number of points in T^* (i.e., $t = |T^*|$). We use s to denote the number of intersection points between T and T^*, i.e., $s = |E| = |T \cap T^*|$. Then, the size of Δ_E is $\binom{s}{k}$, i.e., $|\Delta_E| = \binom{s}{k}$. Given the size of Δ_E, we have the following probabilities:

$$\Pr(\mathbf{W} \in \Delta_E) = \frac{\binom{s}{k}}{\binom{n}{k}}, \tag{19}$$

$$\Pr(\mathbf{W} \in \Delta_T) = 1 - \frac{\binom{s}{k}}{\binom{n}{k}}, \tag{20}$$

$$\Pr(\mathbf{W} \in \Delta_{T^*}) = 0. \tag{21}$$

$$\Pr(\mathbf{Z} \in \Delta_E) = \frac{\binom{s}{k}}{\binom{t}{k}}, \tag{22}$$

$$\Pr(\mathbf{Z} \in \Delta_{T^*}) = 1 - \frac{\binom{s}{k}}{\binom{t}{k}}, \tag{23}$$

$$\Pr(\mathbf{Z} \in \Delta_T) = 0. \tag{24}$$

We have $\Pr(\mathbf{W} \in \Delta_E) = \frac{\binom{s}{k}}{\binom{n}{k}}$ because $\Pr(\mathbf{W} \in \Delta_E) = \frac{|\Delta_E|}{|\Delta_T \cup \Delta_E|} = \frac{\binom{s}{k}}{\binom{n}{k}}$. Since $\Pr(\mathbf{W} \in \Delta_T) + \Pr(\mathbf{W} \in \Delta_E) = 1$, we have $\Pr(\mathbf{W} \in \Delta_T) = 1 - \frac{\binom{s}{k}}{\binom{n}{k}}$. We have $\Pr(\mathbf{W} \in \Delta_{T^*}) = 0$ because $\mathbf{W}$ is subsampled from T, which does not contain any points from $T^* \setminus E$. Similarly, we can compute the probabilities of random variable $\mathbf{Z}$ in those regions.

Based on the fact that p_y and $p_i (i \neq y)$ should be integer multiples of $1/\binom{n}{k}$, we derive the following bounds:

$$\underline{p_y'} \triangleq \frac{\lceil \underline{p_y} \cdot \binom{n}{k} \rceil}{\binom{n}{k}} \leq \Pr(f(\mathbf{W}) = y), \tag{25}$$

$$\overline{p}_i' \triangleq \frac{\lfloor \overline{p}_i \cdot \binom{n}{k} \rfloor}{\binom{n}{k}} \geq \Pr(f(\mathbf{W}) = i), \forall i \neq y. \tag{26}$$

Deriving a lower bound of p_y^*: We define a binary function $H_y(\varphi) = \mathbb{I}(f(\varphi) = y)$, where $\varphi \in \Phi$ and $\mathbb{I}$ is an indicator function. Then, we have the following based on the definitions of the random variable $\mathbf{Z}$ and the function H_y:

$$p_y^* = \Pr(f(\mathbf{Z}) = y) = \Pr(H_y(\mathbf{Z}) = 1). \tag{27}$$

Our idea is to find a region such that we can apply Lemma 1 to derive a lower bound of $\Pr(H_y(\mathbf{Z}) = 1)$. We assume $\underline{p_y'} - \left(1 - \frac{\binom{s}{k}}{\binom{n}{k}}\right) \geq 0$. We can make this assumption because we only need to find a sufficient condition. Then, we can find a region $\Delta_y \subseteq \Delta_E$ satisfying the following:

$$\begin{aligned}
&\Pr(\mathbf{W} \in \Delta_y) && (28)\\
=&\underline{p_y'} - \Pr(\mathbf{W} \in \Delta_T) && (29)\\
=&\underline{p_y'} - \left(1 - \frac{\binom{s}{k}}{\binom{n}{k}}\right). && (30)
\end{aligned}$$

We can find the region Δ_y because $\underline{p_y'}$ is an integer multiple of $\frac{1}{\binom{n}{k}}$. Given the region Δ_y, we define the following region:

$$\mathcal{A} = \Delta_T \cup \Delta_y. \tag{31}$$

Then, based on Equation (25), we have:

$$\Pr(f(\mathbf{W}) = y) \geq \underline{p_y'} = \Pr(\mathbf{W} \in \mathcal{A}). \tag{32}$$

We can establish the following based on the definition of $\mathbf{W}$:

$$\Pr(H_y(\mathbf{W}) = 1) = \Pr(f(\mathbf{W}) = y) \geq \Pr(\mathbf{W} \in \mathcal{A}). \tag{33}$$

Furthermore, we have $\Pr(\mathbf{W} = \varphi) > \epsilon \cdot \Pr(\mathbf{Z} = \varphi)$ if and only if $\varphi \in \Delta_T$ and $\Pr(\mathbf{W} = \varphi) = \epsilon \cdot \Pr(\mathbf{Z} = \varphi)$ if $\varphi \in \Delta_y$, where $\epsilon = \frac{\binom{t}{k}}{\binom{n}{k}}$. Therefore, based on the definition of $\mathcal{A}$ in Equation (31) and the condition in Equation (33), we obtain the following by applying Lemma 1:

$$\Pr(H_y(\mathbf{Z}) = 1) \geq \Pr(\mathbf{Z} \in \mathcal{A}). \tag{34}$$

Since we have $p_y^* = \Pr(H_y(\mathbf{Z}) = 1)$, $\Pr(\mathbf{Z} \in \mathcal{A})$ is a lower bound of p_y^* and can be computed as follows:

$$\begin{aligned}
&\Pr(\mathbf{Z} \in \mathcal{A}) && (35)\\
=&\Pr(\mathbf{Z} \in \Delta_T) + \Pr(\mathbf{Z} \in \Delta_y) && (36)\\
=&\Pr(\mathbf{Z} \in \Delta_y) && (37)\\
=&\Pr(\mathbf{W} \in \Delta_y)/\epsilon && (38)\\
=&\frac{1}{\epsilon} \cdot \left(\underline{p_y'} - \left(1 - \frac{\binom{s}{k}}{\binom{n}{k}}\right)\right). && (39)
\end{aligned}$$

We have Equation (37) from (36) because $\Pr(\mathbf{Z} \in \Delta_T) = 0$, Equation (38) from (37) as $\Pr(\mathbf{W} = \varphi) = \epsilon \cdot \Pr(\mathbf{Z} = \varphi)$ for $\varphi \in \Delta_y$, and the last Equation from Equation (28) - (30).

Deriving an upper bound of $\max_{i \neq y} p_i^*$: We leverage the second part of Lemma 1 to derive an upper bound of $\max_{i \neq y} p_i^*$. We assume $\Pr(\mathbf{W} \in \Delta_E) > \overline{p}_i'$, $\forall i \in \{1, 2, \cdots, c\} \setminus \{y\}$. We can make the assumption because we aim to derive a sufficient condition. For $\forall i \in \{1, 2, \cdots, c\} \setminus \{y\}$, we can find a region $\Delta_i \subseteq \Delta_E$ such that we have the following:

$$\Pr(\mathbf{W} \in \Delta_i) = \overline{p}_i'. \tag{40}$$

We can find the region because $\overline{p}_i'$ is an integer multiple of $\frac{1}{\binom{n}{k}}$. Given region Δ_i, we define the following region:

$$\mathcal{B}_i = \Delta_i \cup \Delta_{T^*}. \tag{41}$$

For $\forall i \in \{1, 2, \cdots, c\} \setminus \{y\}$, we define a function $H_i(\varphi) = \mathbb{I}(f(\varphi) = i)$, where $\varphi \in \Phi$. Then, based on Equation (26) and the definition of random variable $\mathbf{W}$, we have:

$$\Pr(H_i(\mathbf{W}) = 1) = \Pr(f(\mathbf{W}) = i) \leq \overline{p}_i' = \Pr(\mathbf{W} \in \mathcal{B}_i). \tag{42}$$

We note that $\Pr(\mathbf{W} = \varphi) < \epsilon \cdot \Pr(\mathbf{Z} = \varphi)$ if and only if $\varphi \in \Delta_{T^*}$ and $\Pr(\mathbf{W} = \varphi) = \epsilon \cdot \Pr(\mathbf{Z} = \varphi)$ if $\varphi \in \Delta_i$, where $\epsilon = \frac{\binom{t}{k}}{\binom{n}{k}}$. Based on the definition of random variable $\mathbf{Z}$, Equation (42), and Lemma 1, we have the following:

$$\Pr(H_i(\mathbf{Z}) = 1) \leq \Pr(\mathbf{Z} \in \mathcal{B}_i). \tag{43}$$

Since we have $p_i^* = \Pr(f(\mathbf{Z}) = i) = \Pr(H_i(\mathbf{Z}) = 1)$, $\Pr(\mathbf{Z} \in \mathcal{B}_i)$ is an upper bound of p_i^* and can be computed as follows:

$$\begin{aligned}
&\Pr(\mathbf{Z} \in \mathcal{B}_i) && (44)\\
=&\Pr(\mathbf{Z} \in \Delta_i) + \Pr(\mathbf{Z} \in \Delta_{T^*}) && (45)\\
=&\Pr(\mathbf{Z} \in \Delta_i) + 1 - \frac{\binom{s}{k}}{\binom{t}{k}} && (46)\\
=&\Pr(\mathbf{W} \in \Delta_i)/\epsilon + 1 - \frac{\binom{s}{k}}{\binom{t}{k}} && (47)\\
=&\frac{1}{\epsilon} \cdot \overline{p}_i' + 1 - \frac{\binom{s}{k}}{\binom{t}{k}}. && (48)
\end{aligned}$$

By considering all possible i in the set $\{1, 2, \cdots, c\} \setminus \{y\}$, we have:

$$\begin{aligned}
&\max_{i \neq y} p_i^* && (49)\\
\leq&\max_{i \neq y} \Pr(\mathbf{Z} \in \mathcal{B}_i) && (50)\\
=&\frac{1}{\epsilon} \cdot \max_{i \neq y} \overline{p}_i' + 1 - \frac{\binom{s}{k}}{\binom{t}{k}} && (51)\\
\leq&\frac{1}{\epsilon} \cdot \overline{p}_e' + 1 - \frac{\binom{s}{k}}{\binom{t}{k}}, && (52)
\end{aligned}$$

where $\overline{p}'_e \geq \max_{i \neq y} \overline{p}'_i$.

Deriving the certified perturbation size: To reach our goal $\Pr(f(\mathbf{Z}) = y) > \max_{i \neq y} \Pr(f(\mathbf{Z}) = i)$, it is sufficient to have the following:

$$\frac{1}{\epsilon} \cdot \left(\underline{p}'_y - \left(1 - \frac{\binom{s}{k}}{\binom{n}{k}} \right) \right) > \frac{1}{\epsilon} \cdot \overline{p}'_e + 1 - \frac{\binom{s}{k}}{\binom{t}{k}} \tag{53}$$

$$\Longleftrightarrow \frac{\binom{t}{k}}{\binom{n}{k}} - 2 \cdot \frac{\binom{s}{k}}{\binom{n}{k}} + 1 - \underline{p}'_y + \overline{p}'_e < 0. \tag{54}$$

Since Equation (54) should be satisfied for all possible perturbed point cloud T^* (i.e., $n - r \leq t \leq n + r$), we have the following sufficient condition:

$$\max_{n-r \leq t \leq n+r} \frac{\binom{t}{k}}{\binom{n}{k}} - 2 \cdot \frac{\binom{s}{k}}{\binom{n}{k}} + 1 - \underline{p}'_y + \overline{p}'_e < 0. \tag{55}$$

When the above Equation (55) is satisfied, we have $\underline{p}'_y - \left(1 - \frac{\binom{s}{k}}{\binom{n}{k}}\right) \geq 0$ and $\Pr(\mathbf{W} \in \Delta_E) = \frac{\binom{s}{k}}{\binom{n}{k}} \geq \overline{p}'_i, \forall i \in \{1, 2, \cdots, c\} \setminus \{y\}$, which are the conditions that we rely on to construct the region Δ_y and $\Delta_i (i \neq y)$. The certified perturbation size r^* is the maximum r that satisfies the above sufficient condition. Note that $s = \max(n, t) - r$. Then, our certified perturbation size r^* can be derived by solving the following optimization problem:

$$r^* = \operatorname*{argmax}_r r$$
$$\text{s.t.} \max_{n-r \leq t \leq n+r} \frac{\binom{t}{k}}{\binom{n}{k}} - 2 \cdot \frac{\binom{\max(n,t)-r}{k}}{\binom{n}{k}} + 1 - \underline{p}'_y + \overline{p}'_e < 0. \tag{56}$$

B. Proof of Theorem 2

Similar to previous work [4, 10, 9], we show the tightness of our bounds via constructing a counterexample. In particular, when $r > r^*$, we will show that we can construct a point cloud T^* and a point cloud classifier f^* which satisfies the Equation (4) such that the label y is not predicted by our PointGuard or there exist ties. Since r^* is the maximum value that satisfies the Equation (56), there exists a point cloud T^* satisfying the following when $r > r^*$:

$$\frac{\binom{t}{k}}{\binom{n}{k}} - 2 \cdot \frac{\binom{\max(n,t)-r}{k}}{\binom{n}{k}} + 1 - \underline{p}'_y + \overline{p}'_e \geq 0 \tag{57}$$

$$\Longleftrightarrow \frac{\binom{t}{k}}{\binom{n}{k}} - \frac{\binom{\max(n,t)-r}{k}}{\binom{n}{k}} + \overline{p}'_e \geq \underline{p}'_y - \left(1 - \frac{\binom{\max(n,t)-r}{k}}{\binom{n}{k}}\right) \tag{58}$$

$$\Longleftrightarrow \frac{\overline{p}'_e}{\epsilon} + 1 - \frac{\binom{\max(n,t)-r}{k}}{\binom{t}{k}} \geq \frac{1}{\epsilon} \cdot \left(\underline{p}'_y - \left(1 - \frac{\binom{\max(n,t)-r}{k}}{\binom{n}{k}}\right)\right) \tag{59}$$

where t is the number of points in T^* and $\epsilon = \frac{\binom{t}{k}}{\binom{n}{k}}$. Let $\Delta_e \subseteq \Delta_E$ be the region that satisfies the following:

$$\Delta_e \cap \Delta_y = \emptyset \text{ and } \Pr(\mathbf{W} \in \Delta_e) = \overline{p}'_e. \tag{60}$$

Note that we can find the region Δ_e because we have $\underline{p}'_y + \overline{p}'_e \leq 1$. We let $\mathcal{B}_e = \Delta_{T^*} \cup \Delta_e$. Then, we can divide the the region $\Phi \setminus (\mathcal{A} \cap \mathcal{B}_e)$ into $c - 2$ regions and we use $\mathcal{B}_i$ to denote each region, where $i \in \{1, 2, \cdots, c\} \setminus \{y, e\}$. In particular, each region $\mathcal{B}_i$ satisfies $\Pr(T \in \mathcal{B}_i) \leq \overline{p}'_i$. We can find these regions since $\underline{p}'_y + \sum_{i \neq y} \overline{p}'_i \geq 1$. Then, we can construct the following point cloud classifier f^*:

$$f^*(\varphi) = \begin{cases} y, \text{ if } \varphi \in \mathcal{A}, \\ i, \text{ if } \varphi \in \mathcal{B}_i. \end{cases} \tag{61}$$

Note that the point cloud classifier f^* is well-defined in the space Φ. We have the following probabilities for our constructed point cloud classifier f^*:

$$\Pr(f^*(\mathbf{W}) = y) = \Pr(\mathbf{W} \in \mathcal{A}) = \underline{p}'_y, \tag{62}$$
$$\Pr(f^*(\mathbf{W}) = e) = \Pr(\mathbf{W} \in \mathcal{B}_e) = \overline{p}'_e, \tag{63}$$
$$\Pr(f^*(\mathbf{W}) = i) = \Pr(\mathbf{W} \in \mathcal{B}_i) \leq \overline{p}'_i, \tag{64}$$

where $i \in \{1, 2, \ldots, c\} \setminus \{y, e\}$. Note that the point cloud classifier f^* is consistent with Equation (4). Moreover, we have the following:

$$\Pr(f^*(\mathbf{Z}) = e) \tag{65}$$
$$= \Pr(\mathbf{Z} \in \mathcal{B}_e) \tag{66}$$
$$= \frac{\overline{p}'_e}{\epsilon} + 1 - \frac{\binom{\max(n,t)-r}{k}}{\binom{t}{k}} \tag{67}$$
$$\geq \frac{1}{\epsilon} \left(\underline{p}'_y - \left(1 - \frac{\binom{\max(n,t)-r}{k}}{\binom{n}{k}} \right) \right) \tag{68}$$
$$= \Pr(\mathbf{Z} \in \mathcal{A}) \tag{69}$$
$$= \Pr(f^*(\mathbf{Z}) = y), \tag{70}$$

where $\epsilon = \frac{\binom{t}{k}}{\binom{n}{k}}$. Note that we derive Equation (68) from (67) based on Equation (59). Therefore, for $\forall r > r^*$, there exist a point cloud classifier f^* which satisfies Equation (4) and a point cloud T^* such that $g(T^*) \neq y$ or there exist ties.

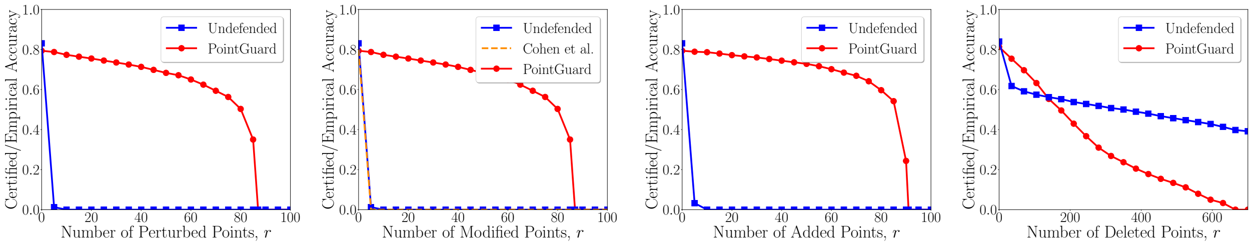

(a) Point perturbation attacks. (b) Point modification attacks. (c) Point addition attacks. (d) Point deletion attacks.

Figure 5: Comparing different methods under different attacks on ScanNet.

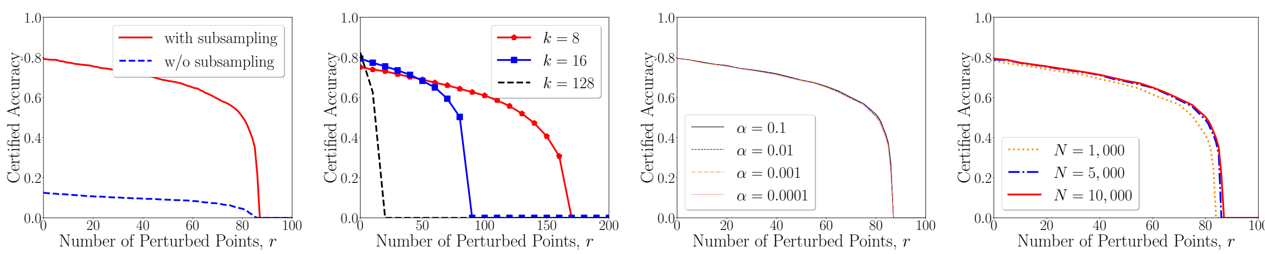

(a) Training with vs w/o subsampling (b) Impact of k (c) Impact of α (d) Impact of N

Figure 6: (a) Training the point cloud classifier with vs. without subsampling. (b), (c), and (d) show the impact of k, α, and N, respectively. The dataset is ScanNet.