STATISTICAL MACHINE LEARNING (WS2025/26)
HOMEWORK: GENERALIZATION BOUND

Assignment (2 points) *Assume a predictor $h_m \colon \mathcal{X} \to \mathcal{Y}$, where $\mathcal{Y} = \{0, 1, \dots, Y-1\}$, trained on i.i.d. sample $T_m = ((x_i, y_i) \in \mathcal{X} \times \mathcal{Y} \mid i = 1, \dots, m)$ by a learning algorithm $A \colon \cup_{m=1}^{\infty} (\mathcal{X} \times \mathcal{Y})^m \to \mathcal{H}$, where $\mathcal{H} = \{h_1, \dots, h_H\}$ is a finite hypothesis space. The trained predictor $h_m = A(T_m)$ is evaluated on the same training sample T_m.*

Given:

- *A sequence of class predictions $(h_m(x_1), \dots, h_m(x_m))$ on the training sample T_m,*
- *A sequence of true class labels $(y_1, \dots, y_m)$,*
- *A loss function $\ell \colon \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}_+$, which will be always the 0/1-loss,*
- *An error level $\delta \in (0, 1)$,*
- *The number of hypotheses H.*

Your task is to compute an upper bound R_{UB} such that

$$R(p, h_m) \leq R_{\text{UB}}$$

holds with probability at least $1 - \delta$. Here,

$$R(p, h) = \mathbb{E}_{(x,y) \sim p}[\, \ell(y, h(x)) \,]$$

denotes the true (expected) risk of predictor h.

You have to fill in code for the following Python function:

```
def generalization_bound(true_y, pred_y, loss, delta, H ):
    # Input:
    #  true_y is np.array(l,dtype=int) representing the true class labels;
    #   a label is an integer from 0 to Y-1
    #  pred_y is np.array(l,dtype=int) representing the predicted class labels;
    #   a label is an integer from 0 to Y-1
    #  Loss is np.array((Y,Y)) whose Loss[y,yy] represents  the loss
    #   incurred when the true label is y and prediction is yy
    #  delta is a scalar from (0,1) representing the probability of failure
    #  H  is a positive integer representing the number of hypothesis.
    #
    # Output:
    #  R_UB the upper bound on the expected risk

    return R_UB
```