


Goal: Perform classification from positive and unlabeled data

Image identification task

■ Identify unlabeled images that are in the same class as the labeled images

■ Dataset contains:

- Labeled images of **pandas** ($y=1$)
- Unlabeled images, either **panda** ($y=1$) or **not panda** ($y=-1$)



■ If the class prior $p(y = 1)$ is known, $p(y = 1|\mathbf{x})$ can be estimated [1]

How to estimate the class prior $p(y = 1)$ from the data?

[1] C. Elkan and K. Noto. Learning classifiers from only positive and unlabeled data. In *14th ACM SIGKDD*, pp. 213–220. ACM, 2008.