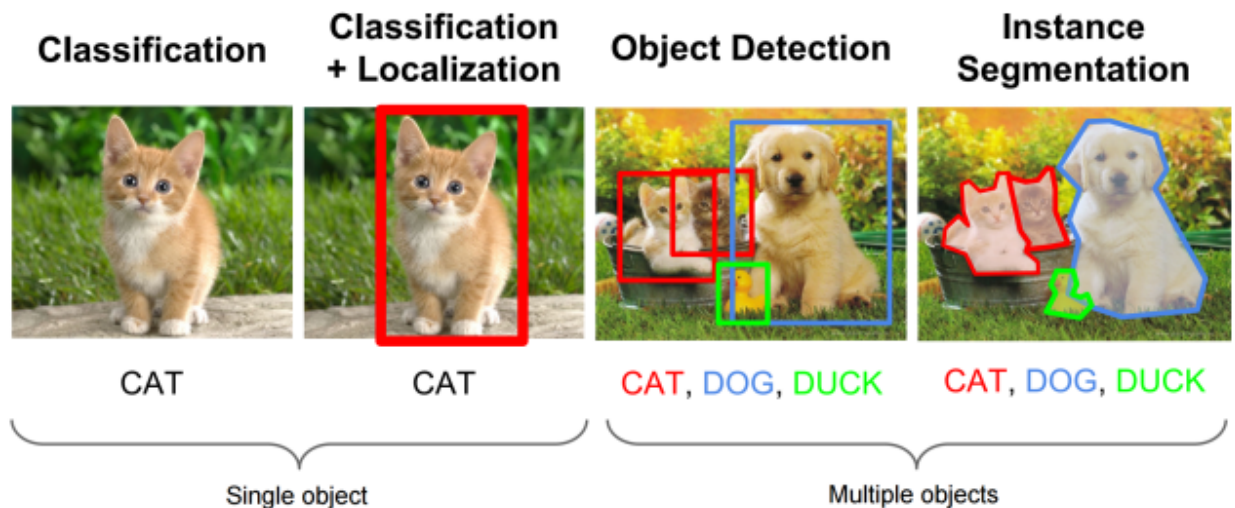


Basic object classification: So far in class we have been solving the basic classification, which is to determine the class the image belongs to. This class is specified by some object we are looking for in the image.

Going beyond this are several problems that arise in practice:

1. **Image localization:** Where is the object located? This is known as the object localization problem, also a task to be solved in the ImageNet contest. The location of the object is a bounding box around the object in the image. Thus we can solve the localization problem by adding four nodes in the last layer that give the upper left, upper right, lower left, and lower right coordinates of the box enclosing the object.
2. **Image segmentation:** What if we want something more specific than a bounding box? For this task we want the output to be of the same size as the input. In the output we put a pixel value of one if the pixel is contained in the object we are looking for and zero otherwise. We use tra




nsposed convolutions (also called deconvolutions) to increase the dimensionality of an image after successive convolutions.

Below is an example of classification (that we have been doing in class) and localization and segmentation.

Source: https://leonardoaraujosantos.gitbooks.io/artificial-intelligence/content/object_localization_and_detection.html

Examples of segmentation:

1. Identifying the tumor in brain MRIs



2. Identifying cars, roads, buildings, and people from public camera images



3. Identifying roads from satellite images

