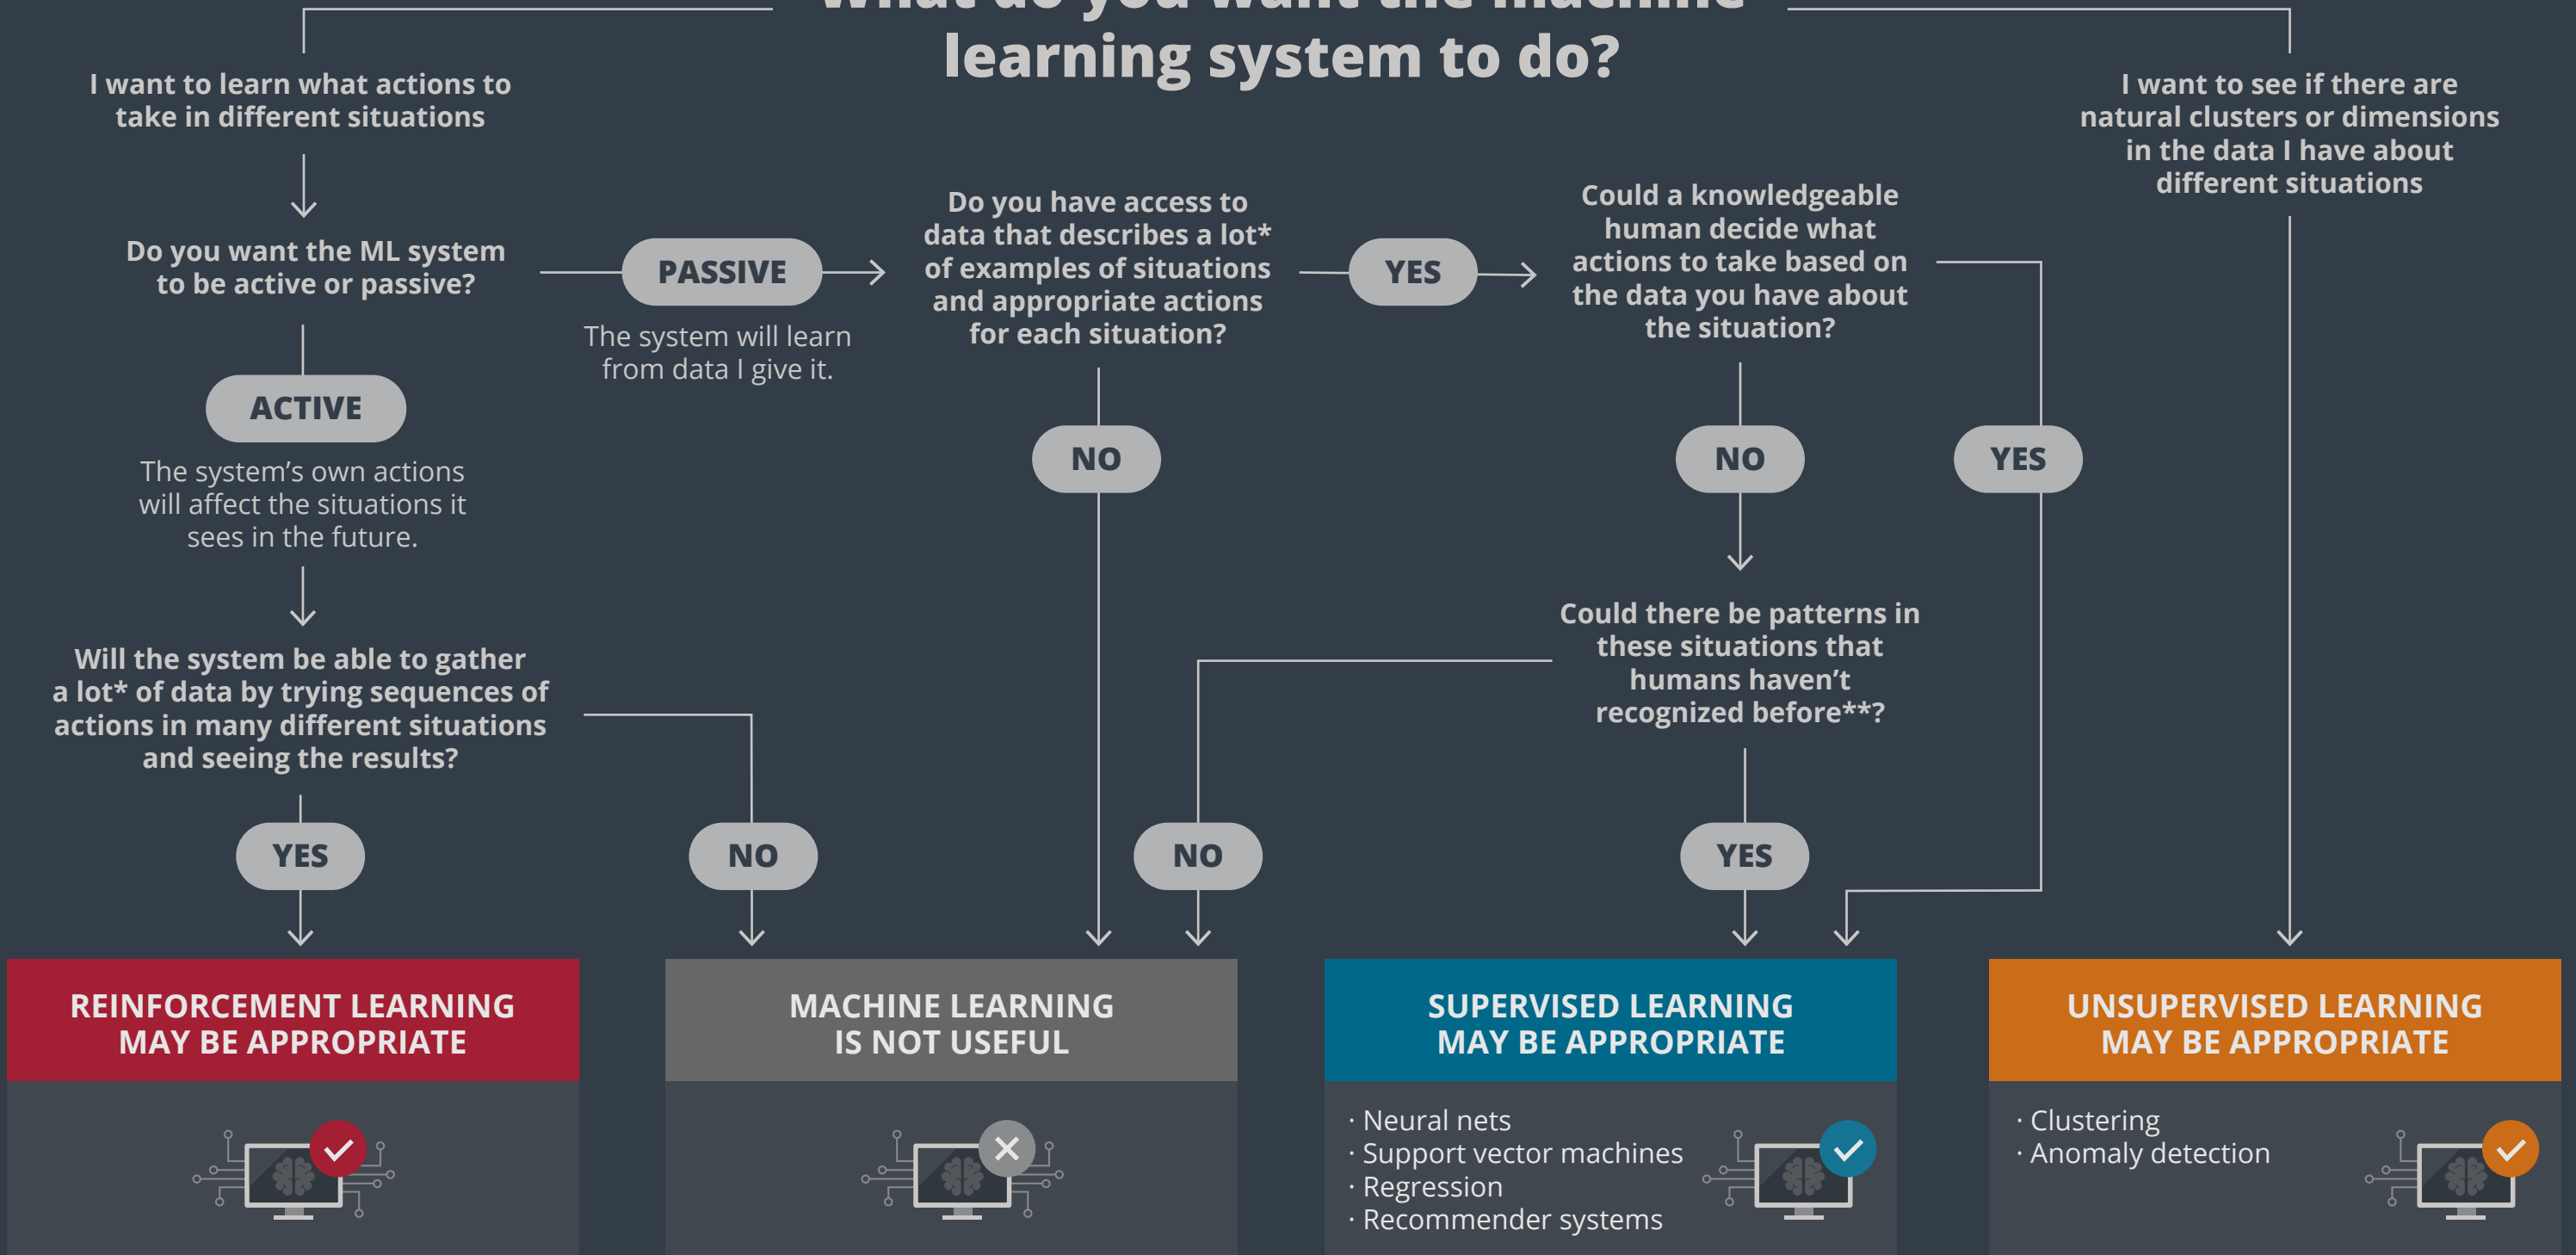


# What do you want the machine learning system to do?



**REINFORCEMENT LEARNING  
MAY BE APPROPRIATE**



**MACHINE LEARNING  
IS NOT USEFUL**



**SUPERVISED LEARNING  
MAY BE APPROPRIATE**

- Neural nets
- Support vector machines
- Regression
- Recommender systems



**UNSUPERVISED LEARNING  
MAY BE APPROPRIATE**

- Clustering
- Anomaly detection



\* It's very hard to say how much data will be needed until you start to analyze your problem in detail, but often you need hundreds, thousands, or millions of examples.

\*\* For example: Have humans not focused on this problem before? Was the data about this problem not available before? Can the patterns here, if they exist, only be detected by looking at more data than humans can really process?