

## Challenge Question

- When should we allow software to make an important decision?
  - When the humans have sufficient transparency to understand:
    - Performance thresholds
    - Intent of the system
    - Training/development pedigree
    - Traceability of the system's decision logic
  - When the humans have **sufficient** experience with the system in variety of contexts
  - When the system has a **strong** track record of reliable performance
  - When humans have observed the system make errors and those errors are **stable**
  - When the implications of human-machine **symbiosis** for decision making/interaction are known
  - When the system can **significantly & consistently** outperform humans and humans **accept** that behavior

# Building Trust

