Ricoh
Project 1: Reinforcement Learning Motion Policy for Pick-and-Place

Background & Context
At a commercial printing facility's cut line, a robot arm with a gripper currently executes pick-and-place motions using a manually programmed policy. This hand-crafted approach, built with low-code "free-drive" programming, works but may not be optimal or easily adaptable to new conditions. Reinforcement learning (RL) offers a way to automatically **learn** an improved motion policy through trial and error, potentially yielding smoother, faster, and more robust pick-and-place actions. Modern robot learning frameworks like NVIDIA **Isaac Lab** (built on Isaac Sim) make it easier to train such policies in simulation.

Goals and Approach
Goal: Train an RL-based motion policy for the robot arm to pick up paper stacks from a conveyor and place them accurately onto a pallet or designated area. By the end of summer, this policy should transfer from simulation to the real robot (running on an **NVIDIA Jetson AGX Orin** or an x86 PC with an RTX 4000 GPU) with minimal fine-tuning.

Approach:
- Create a simulated environment using Isaac Sim and USD models.
- Define reward functions, sensor inputs, and variations.
- Train policies using PPO and SAC algorithms.
- Evaluate generalization.
- Transfer to real-world robot and validate performance.

Tools & Resources
- **Isaac Sim**, **Isaac Lab**, and **Isaac Gym**
- RL libraries like PPO, SAC (Stable Baselines3, RLlib)
- Deployment on AGX Orin or x86 with RTX 4000
- Documentation and benchmarks: RLBench, Isaac Gym

Expected Outcomes
- A trained RL policy for pick-and-place
- Deployment-ready model with supporting code
- Documentation of training and transfer process

References
- Isaac Lab Documentation
- Isaac Gym Preview
- PPO (Proximal Policy Optimization)
- SAC (Soft Actor Critic)
- RLBench Benchmark
