Job Description

Summer Intern – Reinforcement Learning Environment Engineer
Alchedata
Location: Remote-friendly (or San Francisco Bay Area preferred)
Duration: 10–12 weeks (full-time, Summer 2026)

About Alchedata

Alchedata is building **Data Infra 2.0** — the agent-orchestrated data platform purpose-built for Physical AI. We integrate evaluation, post-training, and domain-specific RL environments into one continuous intelligence loop so the next generation of embodied AI and vision-language-action models can move from lab demos to real-world deployment at scale.

The Role

We are looking for a **highly motivated Summer Intern** to work directly on the design, implementation, and optimization of **Reinforcement Learning (RL) environments** for Physical AI and Vision-Language Models (VLMs / VLA). You will help create high-fidelity, multi-modal simulation environments that power our post-training and data flywheel. This is a hands-on role with real ownership — your environments will be used by our Nvex orchestration layer and customer models.

What You'll Do

- Design and implement **domain-specific RL environments** for Physical AI tasks (robot manipulation, locomotion, dexterous grasping, human-robot interaction, etc.).
- Build and extend multi-modal observation spaces (RGB + depth + tactile + proprioception + language instructions) using modern simulators (Isaac Lab, Isaac Gym, MuJoCo, SAPIEN, or custom).
- Develop reward functions, termination conditions, and curriculum learning pipelines that align with real-world robot behavior and VLM/VLA objectives.
- Integrate RL environments with our post-training stack (SFT → RLHF/PPO/DPO-style training loops) and evaluation harness.

- Run large-scale experiments, benchmark environment fidelity (sim-to-real gap), and iterate rapidly based on model performance feedback.
- Collaborate closely with ML engineers and the founding team on production-grade code and documentation.

What We're Looking For

Must-have - Currently pursuing Master or PhD degree in Computer Science, Robotics, AI/ML, or a related field (graduating 2026 or later). - Strong proficiency in **Python** and **PyTorch**. - Hands-on experience with at least one RL framework (Gymnasium, Stable-Baselines3, RLlib, or CleanRL). - Familiarity with at least one robotics simulator (Isaac Lab/Gym, MuJoCo, PyBullet, etc.). - Solid understanding of RL fundamentals (MDPs, policy gradients, value functions, PPO/SAC, etc.).

Nice-to-have (bonus points) - Experience with multi-modal or vision-language models (VLMs, VLA, diffusion policies, etc.). - Exposure to real robot hardware or sim-to-real transfer. - Knowledge of reward shaping, curriculum learning, or human preference data in RL. - Familiarity with modern RL environments for embodied AI (e.g., ManiSkill, RoboMimic, ORBIT, etc.).

What You'll Gain

- Ownership of production RL environments that ship to real customers.
- Exposure to the full stack: data pipeline → RL training → evaluation → agent orchestration.
- A fast-paced startup environment where your work directly impacts the next wave of Physical AI.
- Competitive summer stipend + potential for full-time offer or strong recommendation letters.

If you're excited about turning cutting-edge RL research into the infrastructure that powers the physical AI revolution, we want to hear from you.

To apply:
Please send your resume + GitHub link (or a short note about a relevant RL project you're proud of) to [intern@alchedata.ai] with subject line:
"Summer 2026 RL Environment Intern – [Your Name]"

We review applications on a rolling basis and aim to fill this role quickly. Early applicants have the best chance.

Alchedata is an equal opportunity employer. We celebrate diversity and are committed to creating an inclusive environment for all interns.

To apply, please send your resume to: hr@alchedata.ai