Workload	Container Orchestrator	Distributed compute engine	Training framework	Inference framework	Other comments
VeRL ByteDance	Documentation primarily for SLURM	**Ray** for scheduling and coordination of RL components	**PyTorch FSDP**, Megatron-LM (for very large models)	**vLLM**, SGLang	Single-controller programming paradigm. Easily switch algorithms (PPO, GRPO) with a few lines of code. Decouples data flow and compute allowing flexible placement on GPUs. Achieves 1.5x–20x throughput versus earlier frameworks. ✧ (GitHub)
SkyRL UC Berkeley Sky Lab	Experiments run on **Kubernetes**	**Ray** for scheduling and coordination of RL components	**PyTorch FSDP**	**vLLM**, SGLang	Single-controller programming paradigm. Extends VeRL with advanced agentic capabilities to perform tasks like SWE-Bench. Includes pre-built SkyRL-Agent models trained on SWE-Bench tasks. ✧ (GitHub, blog)
OpenRLHF OpenRLHF community	Scripts provided for SLURM	**Ray** for scheduling and coordination of RL components	**PyTorch,** DeepSpeed	**vLLM**	Easy-to-use and scalable RLHF framework. Designed to make RLHF training simple. Enables training for models up to 70B+ parameters. ✧ (GitHub, blog)
Open-Instruct AllenAI	Experiments run on **Kubernetes**-based Beaker platform	**Ray** for scheduling and coordination of RL components	**PyTorch,** DeepSpeed	**vLLM**	Post-training and instruction-tuning toolkit. Includes fine-tuning language models, DPO, preference tuning, and RL with verifiable rewards (RLVR). Includes scripts for measuring dataset contamination. Used to produce AllenAI's TÜLU-3 models. ✧ (GitHub)
NeMo-RL Nvidia	Documentation primarily for SLURM, but also **Kubernetes**	**Ray** for scheduling and coordination of RL components	**PyTorch FSDP**, Megatron-LM (for very large models)	**vLLM**	Single-controller programming paradigm. A scalable and efficient post-training library for models ranging from 1 to 1000s of GPUs and from tiny models to over 100B parameters. Emphasizes scalability. ✧ (GitHub)