

AI DevDay Agenda

TIME	MAINSTAGE	WORKSHOP ROOM	
8:30 – 9:30 AM	Registration & Breakfast		
9:30 – 9:35 AM	Opening Remarks		
9:35 – 9:55 AM	Luminary Fireside Chat with Ion Stoica, Professor		
9:55 – 10:15 AM	Building AI Agents That Work, Lin Qiao, Fireworks AI		
10:15 – 10:30 AM	AMD Technical Keynote		
10:30 – 10:50 AM	Break		
10:50 – 11:10 AM	Why is Everyone Doing Reinforcement Learning? <i>Daniel Han, CEO, Unsloth</i>	Beginner Workshop Developing Multi-Model Multi-Agent Systems	
11:10 – 11:30 AM	Easy, Fast, and Cheap LLM Serving for Everyone <i>Simon Mo, vLLM Project Co-Lead, vLLM</i> <i>Robert Shaw, Director of Engineering, Red Hat AI</i>		
11:30 – 11:50 AM	Bridging the Gap – Unlocking Access to the Scalable AI Infrastructure <i>Linda Yang, Staff. AI Solutions Management, Supermicro</i>		
11:50 – 1:00 PM	Lunch		
1:00 – 1:20 PM	The Easiest Way to Get Started Running Models Locally and In The Cloud <i>Michael Chiang, Co-founder, Ollama</i>	Beginner Workshop Teaching Models to Think: GRPO with Unsloth on a MI300X GPU <i>Partner: Unsloth</i>	
1:20 – 1:40 PM	gpt-oss. Bringing agentic reasoning models to AMD <i>Dominik Kundel, Developer Experience, OpenAI</i>		
1:40 – 2:00 PM	Blazing-Fast GenAI: Unleashing Performance with AMD Instinct GPUs <i>Erwan Mendard, SVP of Product Management, Crusoe</i>		
2:00 – 2:10 PM	Break		
2:10 – 2:30 PM	Frontier AI in your hand: on-device intelligence with Gemma and Gemini Nano <i>Tris Warkentin, Director of Product Mgmt., Google DeepMind</i>	Advanced Workshop Generating Video using Open-source Tools and Workflows (Wan2.2) <i>Partner: ComfyUI</i>	
2:30 – 2:50 PM	Open-Source Model Performance Optimization <i>Yineng Zhang, Core Maintainer, SGLang</i>		
2:50 – 3:10 PM	Utilizing Low-Precision Datatypes in PyTorch, TorchAO, and Triton <i>Driss Guessous, Meta</i> <i>Jeff Daily, Fellow, AMD</i>		
3:10 – 3:20 PM	Break		
3:20 – 3:40 PM	Introduction to Primus, <i>Zhenyu Gu, Sr. Director of SW Development, AMD</i>	Advanced Workshop How to Develop Optimized Kernels + From Models to Kernels: Powering AI with AMD <i>Partner: Hugging Face</i>	
3:40 – 4:00 PM	Agents to Generate Kernels for Faster AI <i>Sina Rafati, PhD – Senior Member of Technical Staff</i>		
4:15 – 5:00 PM	Closing Remarks, GPU Mode Contest Recognition Fireside Chat with Simran Arora, Academic Partner of Together AI		
5:00 – 6:00 PM	Reception & Raffle Winners		