

GPU Technology Conference 2010 Sessions on Machine Learning & Artificial Intelligence (subject to change)

IMPORTANT: Visit www.nvidia.com/gtc for the most up-to-date schedule and to enroll into sessions to ensure your spot in the most popular courses.

2132 - Accelerating Biologically Inspired Computer Vision Models

Join us for a discussion on applying commodity-server-based clusters and GPU-based clusters to simulating computer vision algorithms at a scale that approaches that of biological vision. We consider the limitations of each technology, survey approaches taken thus far, and suggest new hybrid models and programming frameworks to overcome current limitations and substantially improve performance.

Speaker: Tom Dean, Google Inc.

Topics: Computer Vision, Machine Learning & Artificial Intelligence

Time: Tuesday, September, 21st, 11:00 - 11:50

2204 - Bridging GPU Computing and Neuroscience to Build Large-Scale Face Recognition on Facebook.

Abstract coming soon.

Speaker: Nicolas Pinto, MIT

Topics: Computer Vision, High Performance Computing, Machine Learning & Artificial Intelligence, Neuroscience

Time: Wednesday, September, 22nd, 14:00 - 14:50

2114 - Cascaded HOG on GPU

We propose a real time HOG based object detector implemented on GPU. To accelerate the detection process, the proposed method uses two serially-cascaded HOG detectors. The first low dimensional HOG detector discards detection windows obviously not showing target objects. It reduces the computational cost of the second high dimensional HOG detector. This

method tested on 640x480 color image and the same size movie. The computation time decreases to 70ms per image. That is 4 times faster than a case of single detector. This method provides real time performance even on middle end GPUs such as GeForce GTS 250.

Speaker: Kento Tarui, AquaCast Corporation

Topics: Computer Vision, Machine Learning & Artificial Intelligence

Time: Thursday, September, 23rd, 16:00 - 16:50

1003 - Closing Keynote with Dr. Sebastian Thrun, Stanford

What really causes accidents and congestion on our roadways? How close are we to fully autonomous cars?

In his keynote address, Stanford Professor and Google Distinguished Engineer, Dr. Sebastian Thrun, will show how his two autonomous vehicles, Stanley (DARPA Grand Challenge winner), and Junior (2nd Place in the DARPA Urban Challenge) demonstrate how close yet how far away we are to fully autonomous cars. Using computer vision combined with lasers, radars, GPS sensors, gyros, accelerometers, and wheel velocity, the vehicle control systems are able to perceive and plan the routes to safely navigate Stanley and Junior through the courses. However, these closed courses are a far cry from everyday driving. Find out what the team will do next to get one step closer to the "holy grail" of computer vision, and a huge leap forward toward the concept of fully autonomous vehicles.

Sebastian Thrun is a professor of computer science and electrical engineering at Stanford, where he directs the Stanford AI Lab. He is also a distinguished engineer at Google. Thrun's team won the DARPA Grand Challenge, a US-Government sponsored autonomous robot race that took place in 2005. Thrun also pioneered the scientific field of probabilistic robotics, and he co-invented Google Street View. In recognition of his contributions, Thrun has been elected into the US National Academy of Engineering and the German Academy of Sciences. He is an elected fellow of the AAI, ECCAI, and WTN. Popular Science included Thrun in their "Brilliant Ten," Forbes Magazine in their "E-Gang" members, Scientific American in their list of 50 world technology and policy leaders, and Fortune selected him as one of the 50 smartest people in tech. Wired Magazine awarded Thrun's robot Stanley the top spot in the most influential robots of all times. The robot is now part of a permanent exhibition in the Smithsonian Museum of American History. Thrun has authored 11 books and over 300 scientific articles.

Topics: General Interest, Computer Vision, Machine Learning & Artificial Intelligence

Time: Thursday, September, 23rd, 17:00 - 18:30

2046 - Efficient Automatic Speech Recognition on the GPU

Gain insights into implementation techniques for the speech recognition inference process based on the state-of-art weighted finite state transducer methods. We will discuss the capabilities of the GPU for handling large, irregular graph-based models with millions of states and arcs. We will also present solutions for four challenges in the implementation of speech recognition on the GPU and talk about our more than an order of magnitude faster performance on one core of a CPU.

Speaker: Jake Chong, Parasians, LLC

Topics: Machine Learning & Artificial Intelligence, Algorithms & Numerical Techniques, Audio Processing

Time: Thursday, September, 23rd, 15:00 - 15:50

4008 - Emerging Companies: CEO on Stage featuring EDGE3, ICD, and Universal Robotics

See the hottest new technologies from startups that could transform computing.

In a lively and fast-paced exchange, the “Emerging Companies Summit - CEO on Stage” sessions will feature CEOs from three startups who will have 8 minutes to introduce their companies and 8 minutes to interact with a panel of industry analysts, investors and technology leaders.

This CEO on Stage session will feature EDGE3, ICD, and Universal Robotics - covering the fields of computer vision, mobile, and artificial intelligence.

Panelists will include Norman Winarsky (SRI), Savitha Srinivasan (IBM), and Rob Enderle (Enderle Group).

Speakers: Norman Winarsky, SRI, David Peters, Savitha Srinivasan, IBM, Rob Enderle, Enderle Group

Topics: General Interest, Computer Vision, Machine Learning & Artificial Intelligence, Mobile & Tablet & Phone

Time: Thursday, September, 23rd, 11:00 - 11:50

2030 - High-Throughput Cell Signaling Network Learning with GPUs

Explore how GPUs are being used to enable high-throughput cell signaling network discovery and data-intensive computational systems biology more generally. Systems biology is transitioning from a largely reductive discipline to one focused on building predictive models of large-scale biological systems. New instrumentation will provide the necessary raw data for

such an approach, the key challenge now is building the hardware and software tools to efficiently and interactively build these models. This session will describe how GPUs can and will play a key role in these efforts.

Speaker: Michael Linderman, Stanford University

Topics: Life Sciences, Algorithms & Numerical Techniques, Machine Learning & Artificial Intelligence

Time: Thursday, September, 23rd, 09:00 - 9:50

2207 - Playing Zero-Sum Games on the GPU

A Zero-Sum game is a match for which the gain of one results in loss of the other. Tic-Tac-Toe, Checkers and Chess are Zero-Sum board game examples. For realizing the best player move, the game is abstracted as a tree, often quite deep, consisting of all possible configurations. We present an efficient GPU implementation of the Mini-Max search algorithm, enhanced with Alpha-Beta pruning. We highlight challenges for deploying non-tail recursion of a highly irregular algorithm on GPUs, proposing a hybrid of compiler and user managed stack. We demonstrate superior performance for running many thousands of 3D Tic-Tac-Toe matches, simultaneously.

Speaker: Avi Bleiweiss, NVIDIA Corporation

Topic: Machine Learning & Artificial Intelligence

Time: Wednesday, September, 22nd, 11:00 - 11:50

2119 - Supercomputing for the Masses: Killer-Apps, Parallel Mappings, Scalability and Application Lifespan

Hear the latest on how supercomputing for the masses is changing the world. We will look at some of the one- to three-orders of magnitude faster killer apps and see how they do it. We will discuss specific mapping to GPGPU hardware and techniques for high performance and near-linear scalability both within and across multiple GPGPUs. We will also consider software investment and the decades long longevity of some successful massively parallel Investments in multithreaded software, scalability, balance metrics, lack of consensus on programming models, and lifecycle considerations.

Speaker: Robert Farber, PNNL

Topics: High Performance Computing, Algorithms & Numerical Techniques, Machine Learning & Artificial Intelligence, Physics Simulation

Time: Tuesday, September, 21st, 11:00 - 11:50

2140 - Superfast Nearest Neighbor Searches Using a Minimal kd-tree

Learn how to adapt a kd-tree spatial data structure for efficient nearest neighbor (NN) searches on a GPU. Although the kd-tree is not a natural fit for GPU implementation, it can still be effective with the right engineering decisions. By bounding the maximum height of the kd-tree, minimizing the memory footprint of data structures, and optimizing the GPU kernel code, multi-core GPU NN searches with tens of thousands to tens of millions of points run 10-40 times faster than the equivalent single-core CPU NN searches.

Speaker: Shawn Brown, UNC, Chapel Hill

Topics: Algorithms & Numerical Techniques, Databases & Data Mining, Machine Learning & Artificial Intelligence

Time: Wednesday, September, 22nd, 14:00 - 14:50

2091 - The GPU in the Reactive Control of Industrial Robots

Universal Robotics is using GPUs for real-time visual sensing in the reactive control of industrial robots. For a robot to work in a complex dynamic environment to achieve a more loosely specified goal, such as moving arbitrary boxes from a pallet to a conveyor, requires reactivity. Reactive control requires intensive, concurrent, low-latency computation for motion planning, exception handling, and sensing. We describe and demonstrate how GPU-based computation enables visual servoing and box moving. We also discuss the potential of the GPU to solve more difficult sensory problems such as multi-robot cooperation, multimodal sensor binding, attention, sensitization, and habituation.

Speaker: Dr. Alan Peters, Universal Robotics, Inc.

Topic: Machine Learning & Artificial Intelligence

Time: Thursday, September, 23rd, 16:00 - 16:50