

## GPU Technology Conference 2010 Sessions on Computer Vision (subject to change)

*IMPORTANT: Visit [www.nvidia.com/gtc](http://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.*

---

### **2236 - A Work-Efficient GPU Algorithm for Level Set Segmentation**

Explore a novel GPU level set segmentation algorithm that is both work-efficient and step-efficient. Our algorithm has  $O(\log n)$  step-complexity, in contrast to previous GPU algorithms which have  $O(n)$  step-complexity. We apply our algorithm to 3D medical images and we show that in typical clinical scenarios, our algorithm reduces the total number of processed level set field elements by 16x and is 14x faster than previous GPU algorithms with no reduction in segmentation accuracy.

Speaker: Mike Roberts, Hotchkiss Brain Institute, University of Calgary, Canada

Topics: Medical Imaging & Visualization, Algorithms & Numerical Techniques, Computer Vision, Computer Graphics

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

### **2298 - Accelerated Image Quality Assessment using Structural Similarity**

Explores the GPU porting and performance analysis of the image quality assessment algorithm based on structural similarity index (SSIM). This index is a powerful tool for image quality assessment and the algorithm is highly suitable for GPU architecture, offering a rapid image quality assessment in many image restoration applications.

Speaker: Mahesh Khadtare, Computational Research Laboratories, Pune, INDIA.

Topics: Computer Vision, Imaging

Time: Thursday, September, 23rd, 11:30 - 11:50

## **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

## **2209 - Accelerating Computer Vision on the Fermi**

### **Architecture**

GPUS have evolved from fixed function to general purpose, and continue to evolve with new features being added in every generation. This talk will discuss how to exploit the new features introduced by the Fermi architecture (such as concurrent kernel execution, writes to texture) to accelerate computer vision algorithms.

Speaker: James Fung, NVIDIA

Topics: Computer Vision, Tools & Libraries

Time: Thursday, September, 23rd, 14:00 - 14:50

## **2235 - Advanced Medical Volume Rendering and Segmentation on the GPU**

Learn how to speed up your interactive medical visualization pipeline by an order of magnitude and dramatically improve rendering quality at the same time. Leading researchers in medical imaging informatics describe recent advances in volume visualization and interactive segmentation. Emphasis is on the underlying parallel GPU algorithms and acceleration data structures.

Speakers: Mike Roberts, Hotchkiss Brain Institute, University of Calgary, Canada, Eric Penner, Hotchkiss Brain Institute, University of Calgary, Canada

Topics: Medical Imaging & Visualization, Algorithms & Numerical Techniques, Computer Vision,

Computer Graphics

Time: Tuesday, September, 21st, 15:00 - 15:50

## **2036 - Algorithms for Automated Segmentation of Medical Imaging Studies Utilizing CUDA**

Discover how GPU computing can help doctors make sense of modern imaging studies. This session is intended for a general audience as well as medical informatics specialists. The focus will be on algorithmic approaches to segmentation as it pertains to CTA (computed tomography angiography) studies. Topics covered will include specialized optimization algorithms and novel lumen tracking methodologies.

Speaker: Supratik Moulik, University of Pennsylvania

Topics: Medical Imaging & Visualization, Computer Vision

Time: Tuesday, September, 21st, 16:00 - 16:50

## **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 AAAI, 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

## **2093 - Computational Photography: Real-Time Plenoptic Rendering**

Get the latest information on GPU-based plenoptic rendering including a demonstration of refocusing, novel view generation, polarization, high dynamic range, and stereo 3D. Learn how

GPU hardware enables plenoptic rendering tasks with high-resolution imagery to be performed interactively, opening up entirely new possibilities for modern photography.

Speakers: Andrew Lumsdaine, Indiana University, Georgi Chunev, Indiana University, Todor Georgiev, Adobe Systems

Topics: Imaging, Computer Vision, Stereoscopic 3D

Time: Wednesday, September, 22nd, 16:00 - 16:50

## **2029 - Computer Vision Algorithms for Automating HD Post-Production**

Discover how post-production tasks can be accelerated by taking advantage of GPU-based algorithms. In this talk we present computer vision algorithms for corner detection, feature point tracking, image warping and image inpainting, and their efficient implementation on GPUs using CUDA. We also show how to use these algorithms to do real-time stabilization and temporal re-sampling (re-timing) of high definition video sequences, both common tasks in post-production. Benchmarking of the GPU implementations against optimized CPU algorithms demonstrates a speedup of approximately an order of magnitude.

Speaker: Hannes Fassold, JOANNEUM RESEARCH

Topics: Computer Vision, Video Processing

Time: Wednesday, September, 22nd, 15:00 - 15:50

## **2043 - Disparity Map Generation**

Explore the algorithms and implementation of disparity maps on the GPU. We will discuss how a disparity map facilitates stereoscopic content creation, applications and approaches tried, and final results of real time calculations on GPUs.

Speaker: Henry Gu, GIC

Topics: Stereoscopic 3D, Computer Vision, Imaging

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

## **2176 - Easy GPU Meta-programming: A Case Study in Biologically-Inspired Computer Vision**

Learn how to let the computer optimize your CUDA and OpenCL code for you with easy GPU Meta-programming and Scripting (e.g. PyCUDA). We will present a case study in which we consider the step-wise optimization of a 3D filter bank convolution, using a suite of open-source tools.

Speaker: Nicolas Pinto, MIT

Topics: Tools & Libraries, Computer Vision, High Performance Computing, Neuroscience

Time: Thursday, September, 23rd, 10:00 - 10:50

## **2021 - Efficient Volume Segmentation on the GPU**

Explore a new technique in the detection of common regions in a 2D/3D data array. Connected components along the axes are linked before actual label propagation starts. The algorithm is completely gather-based, which allows for several optimizations in the CUDA C implementation. It enables real-time frame rates for the analysis of typical 2D images and interactive frame rates for the analysis of typical volume data.

Speakers: Allan Rasmusson, University of Aarhus (NVIDIA intern), Gernot Ziegler, NVIDIA

Topics: Algorithms & Numerical Techniques, Computer Vision, Imaging, Medical Imaging & Visualization

Time: Wednesday, September, 22nd, 17:00 - 17:50

## **4003 - Emerging Companies Summit Panel: GPUs for Computer Vision**

This Emerging Companies Summit panel will feature analyst Jon Peddie, President at Jon Peddie Research, and will cover facial, gesture, human motion, and biometrics recognition; augmented reality; and robotic computing.

Speakers: Jon Peddie, Jon Peddie Research, Joe Stam, NVIDIA, Yoram Yaacovi, Microsoft Israel, R&D Center, Sam Cox, Milabra, Janko Mrcsic-Flogel, Mirriad, Tom Dean, Google Inc.

Topic: Computer Vision

Time: Wednesday, September, 22nd, 15:00 - 15:50

## **4005 - Emerging Companies: CEO on Stage featuring Empulse, Rocketick, and Softkinetic**

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 Jedox Business Intelligence, Rocketick, and Softkinetic - covering the fields of computer vision, data mining and high performance computing.

Panelists will include Drew Lanza (Partner, Morgenthaler), Charles Carmel (Cisco), and Nathan Brookwood (Principal Analyst, Insight64).

Speakers: Drew Lanza, Morgenthaler, Uri Tal, Rocketick, Charles Carmel, Cisco, Nathan Brookwood, Insight 64, Kristian Raue, Jedox AG, Michel Tombroff, Softkinetic

Topics: General Interest, Computer Vision, Databases & Data Mining, High Performance Computing

Time: Wednesday, September, 22nd, 17:00 - 17:50

## **4011 - Emerging Companies: CEO on Stage featuring Cinnafilm, Perceptive Pixel, and Total Immersion**

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 Cinnafilm, Perceptive Pixel, and Total Immersion - covering film, imaging, and computer vision.

Panelists will include Bill Tai (CRV), Paul Weiskopf (Adobe), and Tim Bjarin (Creative Strategies).

Speakers: Bill Tai, Charles River Ventures, Lance Maurer, Cinnafilm, Inc., Bruno Uzzan, Total Immersion, Paul Weiskopf, Adobe, Tim Bjarin, Creative Strategies

Topics: General Interest, Computer Vision, Film, Imaging

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

## **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

## **2123 - Enabling Augmented Reality with GPU Computing**

This talk will take a detailed look at Sportvision's “First and 10” system, perhaps the most widely experienced example of AR ever, with 106 million viewers during the 2010 Superbowl alone. We'll examine the current implementation and the GPU features that enable low latency, video-rate performance.

Speaker: Ryan Ismert, Sportvision, Inc.

Topic: Computer Vision

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

## **2173 - Enabling Large-Scale CCTV Face Recognition**

Learn how to use CUDA and GPGPU to perform large scale face search for both forensics as well as CCTV face recognition.

Speaker: Abbas Bigdeli, NICTA

Topics: Computer Vision, Video Processing

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

## **2215 - Extending OpenCV with GPU Acceleration**



OpenCV is a widely popular computer vision library, with millions of downloads and hundreds of thousands of users. Applications span many industries including robotics, industrial machine vision, automotive, film & broadcast, medical, and consumer applications. NVIDIA and the OpenCV development team are collaborating to provide CUDA implementations of the most demanding algorithms, thus enabling a new level of real-time capability and higher quality results.

This talk will introduce OpenCV, and summarize the new CUDA enabled capabilities, and provide an overview of future plans.

Speaker: Joe Stam, NVIDIA

Topics: Computer Vision, Imaging, Stereoscopic 3D, Video Processing

Time: Thursday, September, 23rd, 10:00 - 10:50

## **2087 - Fast High-Quality Panorama Stitching**

We present a panorama stitching application implemented with CUDA C on the GPU. The image processing pipeline consists of SIFT feature detection and matching and Graphcut image stitching to achieve high-quality results. We demonstrate live panorama creation with a Webcam.

Speaker: Timo Stich, NVIDIA

Topics: Video Processing, Algorithms & Numerical Techniques, Computer Vision, Imaging

Time: Thursday, September, 23rd, 14:00 - 14:50

## **2065 - Massively Accelerating Iterative Gauss-Newton Fitting**

To measure three-dimensional shape data of objects, we build up a measurement system that assigns three-dimensional coordinates to the position of projected measurement labels in a camera image. To achieve high measurement accuracy across high amounts of measurement points, we need a very quick routine to localize measurement labels with high precision. To speed up the computation, we evaluate the fits using the CUDA architecture. The final implementation speeds up the fitting of 104 two-dimensional Gauss functions by a factor of 90.

Speaker: Daniel Härter, University of Freiburg, IMTEK, Laboratory for Process Technology

Topics: Computer Vision, Stereoscopic 3D

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

## **2024 - NVIDIA Acceleration Engines with NVIDIA SceniX (Pre-Conference Tutorial)**

Learn how NVIDIA SceniX is a rapid start to building state of the art, realtime 3D applications, and how raytracing can be combined with raster graphics for new levels of interactive realism.

Speakers: Phillip Miller, NVIDIA, Holger Kunz, NVIDIA

Topics: Programming Languages & Techniques, Computer Vision, Ray Tracing

Time: Monday, September, 20th, 13:00 - 14:20

## **2293 - Scaling Up and Scaling Out GPUs with Supermicro's Twin™ Architecture (Sponsored by Supermicro)**

Find out how Supermicro scales up and scales out GPU performance by using Twin™ architecture. In this session, we outline Supermicro's Twin™ architecture advantages across 1U/2U GPU servers and the design of personal supercomputer, and how we are able to scale and optimize GPU technology for datacenter environment and for professional workstation.

Speaker: Don Clegg

Topics: High Performance Computing, Computer Vision

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

## **2242 - Swarming Bacteria and Diffusing Particles: High-Throughput Analysis of Microscopic 3D Motion**

Ever since the 1827 discovery of Brownian motion by observing pollen grains, quantifying motion under the microscope has led to breakthroughs in physics, biology and engineering. Here, I present methods we have developed using confocal microscopy to deduce 3D structure and dynamics from 2D image sequences. We analyze the motion of diffusing colloidal particles and swarms of bacteria free to swim in 3D, which we observe at the single-organism level. We rely heavily on GPU computing to process our large data sets, making extensive use of NPP, CuFFT and optical-flow CUDA algorithms originally developed for machine vision in automobiles.

Speaker: Peter Lu, Harvard University

Topics: Computer Vision, Imaging, Life Sciences

Time: Wednesday, September, 22nd, 17:00 - 17:50

## **2178 - Using GPUs to Track Changes in the Sun**

Learn how GPU computing is enabling astrophysicists to study our closest star. NASA's recently launched Solar Dynamics Observatory is continuously streaming full-disk images of the Sun at visible, UV and EUV wavelengths. This presentation will discuss ways that GPU computing is helping scientists cope with the analysis of the immense data volumes as well as in numerical modeling of the Sun.

Speaker: Mark Cheung, Lockheed Martin Solar & Astrophysics Laboratory

Topics: Astronomy & Astrophysics, Computer Vision, Computational Fluid Dynamics, Physics Simulation

Time: Wednesday, September, 22nd, 17:00 - 17:50