VMV '04 - FINAL PROGRAM

	Monday, November 15, 2004	Tuesday, November 16, 2004	Wednesday, November 17	Thursday, November 18
8:30-9:15		3D Modeling of Static and Dynamic Scenes *Avideh Zahkor, UC-Berkeley*	A Taxonomy of Visual Recognition *Pietro Perona, CalTech*	Computing maps and tracks in human visual cortex *Brian Wandell, Stanford University*
9:15-10:00		Visual 3D Modeling from Images *Marc Pollefeys, UNC-Chapel Hill*	Recognition and Synthesis of Human Actions from Video *Jitendra Malik, UC-Berkeley*	GPU-Based Visualization: New Approaches and Advanced Algorithms *Thomas Ertl, University of Stuttgart*
10:00-11:15 **POSTER SESSION**		Calibration, Registration, Tracking	Simulation and Rendering	Volume Data Processing and Scientific Visualization
11:15-12:00		High-performance imaging using dense arrays of cameras and projectors *Marc Levoy, Stanford University*	Pervasive Multi-Sensor Egomotion Estimation for Direct Interaction and Unstructured Robotics *Seth Teller, MIT*	Video-Based Rendering *Richard Szeliski, Microsoft Research*
12:00-13:30		**LUNCH**		
13:30-14:15		Digitizing the Parthenon: Estimating Surface Reflectance Properties of a Complex Scene under Captured Natural Illumination *Paul Debevec, USC*	Problems Revealed by Applications in Art History and Archaeology *Holly Rushmeier, Yale University*	
14:15-15:30 **POSTER SESSION**		Image and Video-based Modeling and Rendering	Geometry Processing	
15:30-16:15	**16:00 Stanford Campus Tour**	Enhancement of Shape Perception by Surface Reflectance Transformation *Tom Malzbender, Hewlett-Packard Laboratories*	3D Scanning: Improving Completeness, Processing Speed and Visualization *Roberto Scopigno, IST*	
16:15-17:00	**17:00-18:00 Registration**	Estimation of Indirect Physical Properties from Images *Eugene Fiume, University of Toronto*	Local and Global Analysis for Point Cloud Data *Leonidas Guibas, Stanford University*	
19:00			**Conference Dinner @ Cantor Arts Museum**	

3D Modeling of Static and Dynamic Scenes

Avideh Zakhor
Professor
Electrical Engineering and Computer Science
University of California- Berkeley
avz@eecs.berkeley.edu

In this talk, I will describe the efforts at U.C. Berkeley's video and image processing lab in the area of 3D model generation. The first half of the talk will focus on static 3D modeling and visualization of urban environments; our approach to this problem consists of generating airborne and ground based models separately and fusing them so as to arrive at a photo realistic 3D models suitable for virtual walk thrus, drive thrus and fly thrus. Our current system is fast, and entirely automated, with no human intervention. The second half of the talk will focus on a system for 3D modeling of dynamic scenes, and as such is work in progress. I will talk about our system architecture and algorithms for generating a time varying 3D depth map sequence.
This work is done jointly with Christian Frueh, who is currently a postdoc at U.C. Berkeley.

Visual 3D Modeling from Images

Marc Pollefeys
Professor
Department of Computer Science
University of North Carolina at Chapel Hill
marc@cs.unc.edu

In this talk we present our work on visual 3D modeling from images. This work has focused on developing automatic approaches that are as flexible as possible and require a minimum of information besides the images themselves. We model objects with a hand-held video camera and reconstruct dynamic events from a network of cameras. The geometric and photometric calibration, as well as the synchronization in the case of camera networks, is recovered from the video data itself.

High-performance imaging using dense arrays of cameras and projectors

Marc Levoy
Professor
Computer Science Department
Stanford University
levoy@cs.stanford.edu

The advent of inexpensive digital image sensors has led to widespread interest in building sensing systems that incorporate large numbers of cameras. To explore this domain, we have built an array of 128 custom CMOS-based video cameras. If our cameras are packed closely together, then the system functions as a single camera of unusually high performance along one or more imaging dimensions, such as spatial resolution, frame rate, dynamic range, signal-to-noise ratio, or spectral sensitivity. If the cameras are placed farther apart, then the system captures videos from multiple viewpoints at once, i.e. a video light field. If the cameras are placed at an intermediate spacing, then the system functions as a single camera with a large synthetic aperture, which allows us to see through partially occluding environments like foliage or crowds. Finally, if we augment the array of cameras with an array of video projectors, we can implement a discrete approximation of confocal microscopy, in which objects not lying on a selected plane become both blurry and dark, effectively disappearing. These techniques have potential application in scientific imaging, remote sensing, surveillance, and cinematic special effects.

Estimating Surface Reflectance Properties of a Complex Scene under Natural Illumination

Paul Debevec
Professor
Institute for Creative Technologies
University of Southern California
debevec@ict.usc.edu

I will present the process we used for capturing the geometry and spatially-varying surface reflectance of a complex scene observed under natural illumination conditions. The process uses a laser-scanned model of the scene's geometry, a set of digital images viewing the scene's surfaces under a variety of natural illumination conditions, and a set of corresponding measurements of the scene's incident illumination for each photograph. The process employs an iterative inverse global illumination technique to compute surface colors for the scene which, when rendered under the recorded illumination conditions, best reproduce the scene's appearance in the photographs. In our process we measure BRDFs of representative surfaces in the scene to better model the non-Lambertian surface reflectance. Our process uses a novel lighting measurement apparatus to record the full dynamic range of both sunlit and cloudy natural illumination conditions. We employ Monte-Carlo global illumination, multiresolution geometry, and a texture atlas system to perform inverse global illumination on the scene. The result is a lighting-independent model of the scene that can be re-illuminated under any form of lighting. We show that the technique can produce novel illumination renderings consistent with real photographs as well as the first reflectance property estimates that have been confirmed to be consistent with ground-truth measurements.

Enhancement of Shape Perception by Surface Reflectance Transformation

Tom Malzbender
Hewlett-Packard Laboratories
tom.malzbender@hp.com

The visual system interprets the shape of an object through a number of cues such as shading, occlusion, stereopsis and motion parallax. We present a simple imaging method that can provide improved perception of surface shape of objects by enhancing shading cues. This allows physical artifacts to be photographed with a conventional digital camera and then redisplayed with a new set of material properties more conducive to the interpretation of surface shape and detail. We demonstrate 3 separate approaches applied to artifacts from the fields of archeology, paleontology, forensics and fine arts. In addition, we demonstrate how our analytic representation of reflectance functions can be used for finding informative lighting directions and synthesizing 3D textures.

Estimation of Indirect Physical Properties from Images

Eugene Fiume
Professor
Computer Science Department
University of Toronto
elf@atlas.dgp.toronto.edu

Imagine that you have a video sequence of a sunlit meadow. Grasses, shrubs and trees are swaying in the wind; a stream gurgles and flows across the field. Now imagine that you wish to place a computer synthesized girl into that meadow. You want her to place a synthetic boat to bob in the stream and you want her hair and clothing to move consistently with the grass and trees. While the extraction of directly observable physical phenomena has received considerable attention in computer graphics and vision, research into the estimation of physical properties that are seen only through their effects is much less well developed. In this talk, we review the progress in this challenging area of research and examine some promising lines of attack.

A Taxonomy of Visual Recognition

Pietro Perona
Professor
California Institute of Technology
perona@caltech.edu

Visual recognition is the most interesting and difficult challenge in computer vision today. Machines that were able to recognize objects and object categories would change the way we live (mostly for the better). During the past few years much progress has been made in organizing our thoughts about recognition, although much remains to be done on the practical side. I will review our current understanding as well as give examples of the state of the art; I will also propose a list of open problems which I feel are particularly interesting.

Recognition and Synthesis of Human Actions from Video

Jitendra Malik
Professor
Computer Science Division
University of California-Berkeley
malik@cs.berkeley.edu

I will describe progress towards the goal of recognizing human actions from video data. We have developed two sets of techniques. In the case of medium resolution human figures, we rely on patterns of optical flow. When higher resolution data is available, we can extract stick figures corresponding to the locations of human joints in space. These representations support versions of "motion capture" for use in synthesis. More information is available at http://http.cs.berkeley.edu/projects/vision/human/index.html

Pervasive Multi-Sensor Egomotion Estimation for Direct Interaction and Unstructured Robotics

Seth Teller
Professor
Electrical Engineering and Computer Science Department
Massachusetts Institute for Technology
teller@lsc.mit.edu

For humans, knowledge of our own location is a basic kind of empowering information: as part of our mental model of the world, it enables us to navigate to desired places, to find resources, and to plan our movements more effectively. Until recently, people had to rely on experience and continuity to locate themselves, their assets and devices. In recent decades, however, position information from the Global Positioning System (GPS) infrastructure has wrought tremendous change in human and robotic activities outdoors, ranging from military operations (including autonomous aircraft), civilian navigation and surveying, to shipping and supply-chain management, resource exploration, and precision agriculture.

We envision an analogous indoor infrastructure, combining active radio and ultrasound beacons with passive receivers using machine vision, to provide fine-grained location and orientation ("pose" or "egomotion") information to human-held devices, autonomous robots, and ordinary objects. This infrastructure has the potential to substantially extend and improve human interfaces and robotic capabilities. For people, pose-awareness facilitates direct interaction with things in the world and their metadata. For robots, pose-awareness makes feasible tasks that are currently out of reach, such as complex household chores.

After motivating the infrastructure, we'll show some early deployments and proof-of-concept applications, and briefly discuss privacy concerns. We'll also show some early efforts toward making the infrastructure deploy itself autonomously.

Problems Revealed by Applications in Art History and Archaeology

Holly Rushmeier
Professor
Computer Science Department
Yale University
holly@acm.org

Art history and archaeology would seem to be ideal applications for showcasing the power and utility of computer graphics. The results of such projects are appealing because they are either digital versions of masterpieces, or they reveal something about our own history. In my experience however, these applications have primarily illustrated how computer graphics techniques fall short. I will discuss problems encountered in obtaining data, processing data into 3D objects, editing large objects, using objects in applications, and documenting and validating the accuracy of digital versions of physical object. I will describe ideas for addressing some of these problems by adopting an "image based everything" approach.

3D Scanning: Improving Completeness, Processing Speed and Visualization

Roberto Scopigno
Professor
Visual Computing Laboratory
ISTI-CNR, Pisa
roberto.scopigno@isti.cnr.it

The construction of detailed and accurate 3D models is made easier by the increasing diffusion of 3d scanning devices. These allow to build digital models of real 3D objects in a cost- and time-effective manner. The talk will present the capabilities of this technology focusing mainly on some issues which are preventing a wider use of this technology, such as for example the considerable user intervention required, the usually incomplete coverage of the object surface and the complexity of the models produced. Another emerging issue is how to support the visual presentation of the models (local or remote) with guaranteed interactive rendering rates. Some examples of the results of current projects, mainly in the Cultural Heritage field, will be shown.

Local and Global Analysis for Point Cloud Data

Leonidas Guibas
Professor
Computer Science Department
Stanford University
guibas@cs.stanford.edu

Digital shapes are becoming ubiquitous and require new tools for their analysis. While audio, images, or video, consist of regularly sampled signals, scanned shapes typically start their life as nothing more than an unorganized collection of points irregularly sampled from the surface of an object -- so called point cloud data. We investigate techniques for local feature detection, segmentation, and more global shape analysis, classification and comparison of such data sets. The irregular sampling creates new challenges and leads to methods with a distinctly more combinatorial and topological character that in traditional signal processing.

Computing Maps and Tracks in Human Visual Cortex

Brian Wandell
Professor
Psychology Department and Neuroscience Program
Stanford University
wandell@stanford.edu

Visual cortex has been an excellent model system for developing a quantitative understanding of brain function. We understand a great deal about the physical signals that initiate vision, and this knowledge has led to a relatively advanced understanding of the organization of major structures in visual cortex, such as visual field maps. This talk will explain several measurements and computational methods that are used to understand human brain development and plasticity.

First, we have developed functional magnetic resonance imaging (fMRI) methods for measuring and quantifying the properties of maps in individual human and macaque brains. To understand the development and plasticity of these maps, we have made measurements in several cases of abnormal development as well as in controlled experiments using macaque.

Second, we are combining fMRI with diffusion tensor imaging (DTI), a method that can be used to study the white matter fibers, to understand visual development. Specifically, as children develop and learn to read certain visual recognition skills become highly automatized and the brain develops specialized visual circuitry to support skilled reading. We are measuring how certain parts of these circuits develop, and how the signals from these circuits are transmitted to other cortical systems.

Video-Based Rendering

Richard Szeliski
Professor
Interactive Visual Media Group
Microsoft Research
szeliski@microsoft.com

Image-based rendering has been one of the hottest areas in computer graphics in recent years. Instead of using CAD and painting tools to construct graphics models by hand, IBR uses real-world imagery to rapidly create extremely photorealistic shape and appearance models. However, IBR results to date have mostly been restricted to static objects and scenes.

Video-based rendering brings the same kind of realism to computer animation, using video instead of still images as the source material. Examples of VBR include facial animation from sample video, repetitive video textures that can be used to animate still scenes and photos, 3D environment walkthroughs built from panoramic video, and 3D video constructed from multiple synchronized cameras. In this talk, I survey a number of such systems develped by our group and by others, and suggest how this kind of approach has the potential to fundamentally transform the production (and consumption) of interactive visual media.

GPU-based Visualization: New Approaches and Advanced Algorithms

Thomas Ertl
Professor
Visualization and Interactive Systems Institute
University of Stuttgart
thomas.ertl@vis.uni-stuttgart.de

Although the recent progress made by modern graphics cards with respect to performance and functionality was targeted to another market segment, scientific visualization was obviously able to benefit from this development. By mapping more and more algorithms of the visualization pipeline to programmable GPUs, interactive visualization of large datasets now becomes possible for many application scenarios. After first results for straightforward algorithms on simply structured data, researchers now investigate new approaches for mapping more advanced algorithms to the GPU. This talk presents some of the recent results of the VIS group of the University of Stuttgart in the area of volume visualization of unstructured meshes, visualization of functionally encoded multifield data, and feature extraction and texture advection for vector field visualization.

Poster Session 1: Calibration, Registration, Tracking

- *Torsten Rohlfing, Joachim Denzler, Daniel Russakoff, Christoph Graessl, Calvin Maurer:*
  Markerless Real-Time Target Region Tracking: Application to Frameless Serotactic Radiosurgery
- *Sanjit Jhala, Suresh Lodha:*
  Stereo and Lidar-Based Pose Estimation with Uncertainty for 3D Reconstruction
- *Jochen Schmidt, Florian Vogt, Heinrich Niemann:*
  Vector Quantization Based Data Selection for Hand-Eye Calibration
- *Stefan Gehrig, Jens Klappstein, Uwe Franke:*
  Active Stereo for Intersection Assistance
- *Narcis Pares, Anna Carreras, Miquel Soler:*
  Non-invasive attitude detection for full-body interaction in MEDIATE, a multisensory interactive environment for children with autism
- *Roland Brockers, Marcus Hund, Barbel Mertsching:*
  A Fast Cost Relaxation Stereo Algorithm with Occlusion Detection for Mobile Robot Applications
- *Yannick Caulier, Klaus Spinnler:*
  Calibration of 1D cameras - Determination of 3D reconstruction accuracy
- *Nizam Anuar, Igor Guskov:*
  Extracting Animated Meshes with Adaptive Motion Estimation
- *Geremy Heitz, Torsten Rohlfing, Calvin Maurer:*
  Automatic Generation of Shape Models Using Nonrigid Registration with a Single Segmented Template Mesh
- *Christoph Graessl, Timo Zinsser, Heinrich Niemann:*
  A Probabilistic Model-based Template Matching Approach for Robust Object Tracking in Real- Time
- *Hui Chen:*
  Gradient-based Approach for Fine Registration of Panorama Images

Poster Session 2: Image and Video-based Modeling and Rendering

- *Ingo Bauermann, Eckehard Steinbach:*
  Low-complexity Image-based 3D Gaming
- *Benjamin Deutsch, Ingo Scholz, Christoph Graessl, Heinrich Niemann:*
  Extending Light Fields using Object Tracking Techniques
- *Volker Scholz, Marcus Magnor:*
  Cloth Motion from Optical Flow
- *Emric Epstein, Martin Granger-Piche, Pierre Poulin:*
  Exploiting Mirrors in Interactive Reconstruction with Structured Light
- *Harlyn Baker, Donald Tanguay:*
  Graphics-Accelerated Panoramic Mosaicking from a Video Camera Array
- *Richard Roussel, Andre Gagalowicz:*
  Realistic Face Reconstruction From Uncalibrated Images
- *Thomas Schiwietz, Ruediger Westermann:*
  GPU-PIV
- *Martin Granger-Piche, Emric Epstein, Pierre Poulin:*
  Interactive Hierarchical Space Carving with Projector-based Calibrations
- *Jan Meseth, Reinhard Klein:*
  Memory Efficient Billboard Clouds for BTF Textured Objects
- *Peter Eisert, Yong Guo, Anke Riechers, Juergen Rurainsky:*
  High-Resolution Interactive Panoramas with MPEG-4

Poster Session 3: Simulation and Rendering

- *A. Zinke, G. Sobottka, A. Weber:*
  Photo-Realistic Rendering of Blond Hair
- *Nils Thuerey, Ulrich Ruede:*
  Free Surface Lattice-Boltzmann fluid simulations with and without level sets
- *Veronica Sundstedt, Alan Chalmers, Kirsten Cater, Kurt Debattista:*
  Top-Down Visual Attention for Efficient Rendering of Task Related Scenes
- *Jens Krokowski, Harald Raecke, Christian Sohler, Matthias Westermann:*
  Reducing State Changes with a Pipeline Buffer
- *S. Kimmerle, M. Wacker, C. Holzer:*
  Multilayered Wrinkle Textures from Strain
- *Rosario De Chiara, Ugo Erra, Vittorio Scarano, Maurizio Tatafiore:*
  Massive Simulation using GPU of a distributed behavioral model of a flock with obstacle avoidance
- *Nicolas Fritz, Philipp Lucas, Philipp Slusallek:*
  CGiS, a new Language for Data-parallel GPU Programming
- *Kirill Dmitriev, Hans-Peter Seidel:*
  Progressive path tracing with lightweight local error estimation
- *Manfred Ernst, Christian Vogelgsang, Guenther Greiner:*
  Stack Implementation on Programmable Graphics Hardware
- *Martin Sunkel, Jan Kautz, Hans-Peter Seidel:*
  Rendering and Simulation of Liquid Foams
- *Marco Winter, Marc Stamminger:*
  Depth-Buffer based Navigation

Poster Session 4: Geometry Processing

- *Stefan Roettger, Stefan Guthe, Andreas Schieber, Thomas Ertl:*
  Convexification of Unstructured Grids
- *Stefan Gumhold:*
  Hierarchical Shape-Adaptive Quantization for Geometry Compression
- *Thomas Lewiner, Luiz Velho, Helio Lopes and Vinicius Mello:*
  Simplicial Isosurface Compression
- *S. Kimmerle, M. Nesme, F. Faure:*
  Hierarchy Accelerated Stochastic Collision Detection
- *Richard Keiser, Matthias Mueller, Bruno Heidelberger, Matthias Teschner, Markus Gross:*
  Contact Handling for Deformable Point-Based Objects
- *Michael Guthe, Akos Balazs, Reinhard Klein:*
  Real-time out-of-core trimmed NURBS rendering and editing
- *Mirko Sattler, Ralf Sarlette, Gabriel Zachmann, Reinhard Klein:*
  Hardware-accelerated ambient occlusion computation
- *Bruno Heidelberger, Matthias Teschner, Richard Keiser, Matthias Mueller, Markus Gross:*
  Consistent penetration depth estimation for deformable collision response
- *Rhaleb Zayer, Christian Roessel, Hans-Peter Seidel:*
  Efficient Iterative Solvers for Angle Based Flattening

Poster Session 5: Volume Data Processing and Scientific Visualization

- *Balazs Csebfalvi, Jozsef Koloszar, Zsolt Tarjan:*
  Vector Quantization for Feature-Preserving Volume Filtering
- *Dorit Merhof, Peter Hastreiter, Grzegorz Soza, Marc Stamminger, Christopher Nimsky:*
  Non-linear Integration of DTI-based Fiber Tracts into Standard 3D MR Data
- *Soeren Grimm, Stefan Bruckner, Armin Kanitsar, Eduard Groeller:*
  Flexible Direct Multi-Volume Rendering in Dynamic Scenes
- *Thomas Klein, Thomas Ertl:*
  Illustrating Magnetic Field Lines using a Discrete Particle Model
- *Thomas Jansen, Bartosz von Rymon-Lipinski, Nils Hanssen, Erwin Keeve:*
  Fourier Volume Rendering on the GPU Using a Split-Stream-FFT
- *Itay Cohen, Dan Gordon:*
  The Voxel-Sweep: A Boundary-based Algorithm for Object Segmentation and Connected- Components Detection
- *Xiaohong Bao, Renato Pajarola, Michael Shafae:*
  SMART: An Efficient Technique for Massive Terrain Visualization from Out-of-core
- *Holger Theisel, Tino Weinkauf, Hans-Christian Hege, Hans-Peter Seidel:*
  Grid-independent Detection of Closed Stream Lines in 2D Vector Fields
- *Gerd Marmitt, Andreas Kleer, Ingo Wald, Heiko Friedrich, Philipp Slusallek:*
  Fast and Accurate Ray-Voxel Intersection Techniques for Iso-Surface Ray Tracing
- *Michael Meissner, Klaus Engel:*
  Pre-Integrated Non-Photorealistic Volume Rendering

Workshop Venues

Tour of Stanford Campus:

Mon, 4 pm in front of *David Packard Building* (lasts about 1 hour)

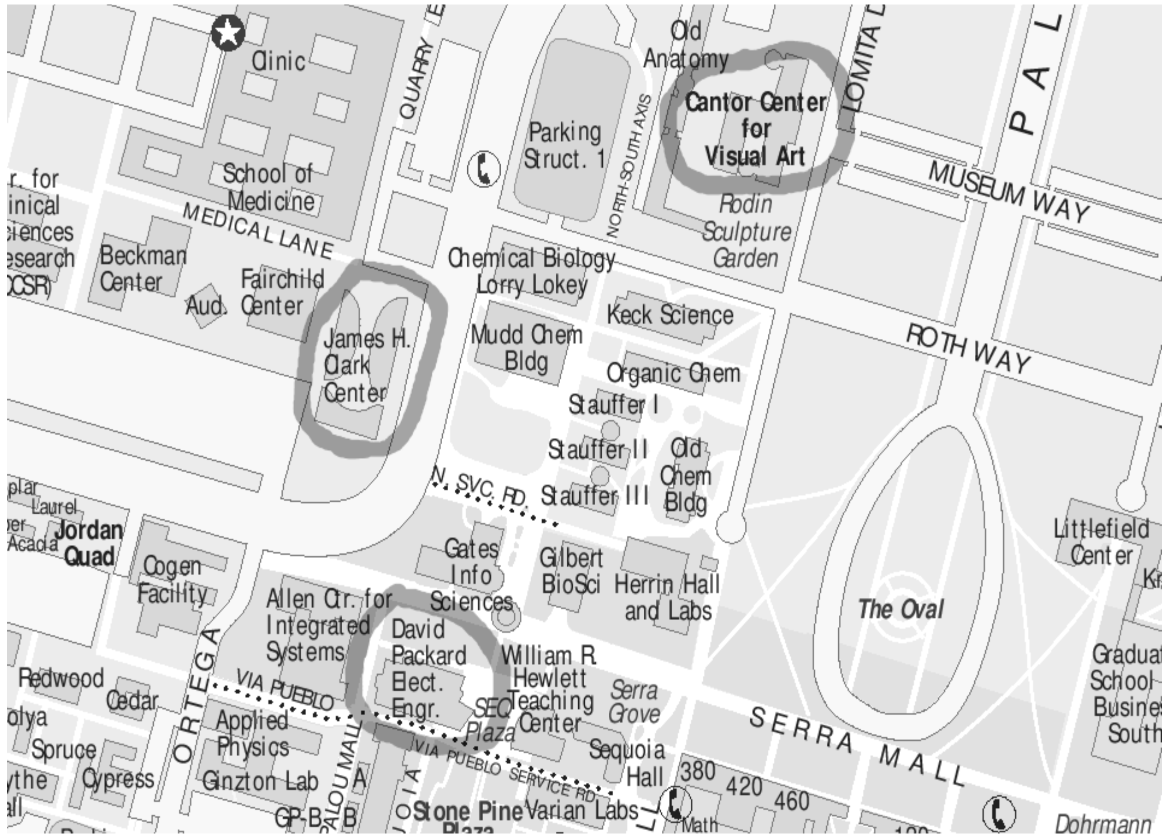

Registration:

Monday, 5pm-6pm in *David Packard Building*, Tue-Thu, at Workshop Site in *Clark Center*

Welcome Reception:

Monday, 5pm-6pm in *David Packard Building*

Conference Dinner:

Wednesday, 7pm in the *Cantor Center for Visual Art*

Workshop Sessions

Tue-Thu, *Clark Center* Auditorium/Linx Cafe

Getting to Stanford

DRIVING TO CAMPUS:

From Highway 101 North & South:

Take the Embarcadero Road exit west toward Stanford. At El Camino Real, Embarcadero turns into Galvez Street as it enters the university. Stay in the left lane and continue toward the center of campus. There is metered parking in a lot on the right side Galvez Street, just past Memorial Way. A few hundred feet beyond this lot, Galvez ends at Serra Street, where there is more metered parking. The visitor information center is in Memorial Hall, which is across from Hoover Tower on Serra Street.

From Highway 280 North & South:

Exit Sand Hill Road east toward Stanford. Continue east, turning right at the traffic light on Santa Cruz Avenue. Make an immediate left onto Junipero Serra Boulevard. Turn left at the second stoplight, Campus Drive East. Continue around Campus Drive East and turn left when you reach Serra Street, at the gas station. Follow Serra Street until it ends at Galvez Street. Turn right onto Galvez and look for the first parking lot on the left which is a public pay parking lot that is adjacent to the side of Memorial Hall. The Visitor Center is in the front of Memorial Hall. Please note that parking is monitored Monday - Friday, 8 a.m. - 4 p.m.

From El Camino Real:

Exit El Camino Real at University Avenue. Turn toward the hills (away from the center of Palo Alto). As you enter Stanford, University Avenue becomes Palm Drive. Go through one traffic light, and turn left onto Campus Drive at the first stop sign. Turn right at Galvez Street, the next stop sign. There is metered parking in the lot between Memorial Hall and the Economics building. The Visitor Center is in the front of Memorial Hall. Please note that parking is monitored Monday - Friday, 8 a.m. - 4 p.m.

PUBLIC TRANSPORTATION

By train (Caltrain):

Caltrain is a commuter rail service that runs between Gilroy and San Francisco. There are two stops close to Stanford: one on California Avenue and another at the end of Palm Drive in downtown Palo Alto. http://www.caltrain.org

The Stanford shuttle, called Marguerite, meets most trains at both the Palo Alto and California Ave. stations from 6 a.m. to 7:45 p.m Monday - Friday.

From points south of Stanford (Santa Clara Valley Transportation Authority):

The Valley Transportation Authority (VTA) runs buses to Palo Alto and most points south of Stanford. It also operates the light rail line in Santa Clara County. http://www.vta.org

From points north of Stanford (SamTrans):

SamTrans runs buses to most points north of Stanford. Five of their lines serve Stanford well. http://www.samtrans.org

From the East Bay (Dumbarton Express):

The Dumbarton Express, used by many East Bay commuters, runs weekdays from the Union City BART station across the Dumbarton Bridge to the Palo Alto Caltrain station, where the Stanford shuttle, called Marguerite, meets most trains. http://www.vta.org

---

FROM LOCAL AIRPORTS

Stanford is roughly equidistant from San Jose and San Francisco airports. Sam Trans' Route KX might be the cheapest way to get to and from San Francisco International Airport. That bus stops in front of the Stanford Shopping Center on El Camino Real. The Stanford shuttle, called Marguerite, connects there. Caltrain and BART (Bay Area Rapid Transit) connect at the Millbrae station and for an additional fee, there are BART trains that run directly to the airport. More information about public transportation from the Oakland, San Jose and San Francisco airports is available through the Transit Info website http://www.transitinfo.org.

A number of private companies also provide shuttle transportation from the San Francisco and San Jose airports, and taxis serve the Stanford campus from both locations.

For additional questions or trip planning assistance, please call the Bay Area Transit Information Hotline by dialing 511 from within the San Francisco Bay Area or (510) 817-1717 from outside the San Francisco Bay Area.

Stanford Parking & Transportation Services (P&TS) is also available if you need assistance planning your trip. They can talk you through the best route - from your front door to campus. Contact the P&TS office at 723-9362, email at transportation@stanford.edu. http://transportation.stanford.edu