Computational Neuroscience 2002 Symposium

Visual Processing of Natural Images

April 5 - 6, 2002

University of Minnesota
Minneapolis, Minnesota USA

Table of Contents

This information is available in alternative formats upon request by individuals with disabilities. Please send email to clinton@compneuro.umn.edu or call (612) 625-8424.

Program Committee

Daniel Kersten, Chair
Department of Psychology, University of Minnesota

Tom Carlson
Graduate Program in Psychology, University of Minnesota

Sing-Hang Cheung
Graduate Program in Psychology, University of Minnesota

Sheng He
Department of Psychology, University of Minnesota

Gordon Legge
Department of Psychology, University of Minnesota

Thomas Naselaris
NSF-IGERT Fellow, Graduate Program in Neuroscience
University of Minnesota

Cheryl Olman
NSF-IGERT Fellow, Graduate Program in Neuroscience
University of Minnesota

Paul Schrater
Department of Psychology, University of Minnesota

Sponsors

National Science Foundation (Integrative Graduate Education and Research Training)

University of Minnesota Center for Cognitive Sciences

University of Minnesota Graduate School

University of Minnesota Institute of Technology

University of Minnesota Medical School

University of Minnesota Supercomputing Institute for Digital Simulation and Advanced Computation

Third Annual Computational Neuroscience Symposium
Visual Processing of Natural Images

Program Schedule

All activities will be held at the Basic Sciences and Biomedical Engineering Building, 312 Church St. SE, Minneapolis, Minnesota 55455

Friday, April 5, 2002

7:30 - 8:00 a.m.	Set up posters in BS&BE 2-101 Atrium
8:00 - 8:45 a.m.	Continental breakfast and registration in BS&BE 2-101 Atrium
8:45 - 9:00 a.m.	Welcome Daniel Kersten, Chair

Morning session: Moderator--Paul Schrater, Department of Psychology

9:00 a.m.	**Bruno Olshausen**, University of California at Davis, "Sparse coding as a principle of image representation in visual cortex"
9:35 a.m.	**David Field**, Cornell University, Ithaca, "Natural scenes and the depth of statistical knowledge"
10:10 a.m.	Discussion
10:30 a.m.	**Pamela Reinagel**, Harvard Medical School, Boston, "Coding of natural visual stimuli in the thalamus: What exactly is optimized? "
11:05 a.m.	**Sheng He**, University of Minnesota, Minneapolis, "What adaptation can tell us about processing of natural images?"
11:35 a.m.	Discussion
11:45 a.m.	Walk to McNamara Center
12:00-1:40	**Lunch in Ski U Mah Room, McNamara Center**

Afternoon session: Moderator--Patty Costello, Graduate Program in Neuroscience

2:00 p.m.	**Aapo Hyvarinen**, Helsinki University of Technology, Finland, “Beyond independence and sparseness in models of natural image statistics”
2:35 p.m.	**Don MacLeod**, University of California at San Diego, “Color discrimination, color constancy, and natural scene statistics”
3:10 - 3:40 p.m.	Discussion and Break
3:40	**Jack Gallant**, University of California, Berkeley, “The role of area V4 in natural vision”
4:15 p.m.	**Geoff Boynton**, The Salk Institute for Biologiical Studies, La Jolla, “Using fMRI to compare cortical magnification factors in human V1 to visual acuity”
4:50 p.m.	Discussion
5:00 - 6:00 p.m.	Posters
6:00 p.m.	Self-organized dinners

Saturday, April 6, 2002

8:00 - 8:45 a.m.	Continental breakfast in BS&BE 2-101 Atrium

Morning session: Moderator--Tom Carlson, Graduate Program in Psychology

9:00 a.m.	**Eero Simoncelli**, New York University, New York, “Neural gain control: ecological justification and characterization”
9:35 a.m.	**Wilson Geisler**, University of Texas, Austin, Natural scene statistics and the evolution of perceptual systems

10:10 - 10:40 a.m.	Discussion and break
10:40 a.m.	**Dario Ringach**, University of California, Los Angeles, “Testing theories of natural image representation in primary visual cortex”
11:15 a.m.	**Dae-Shik Kim**, University of Minnesota, Minneapolis, “What does neuroimaging tell us about visual coding strategies?”
11:50 - Noon	Discussion
12:00 p.m.	**Lunch boxes in BS&BE 2-101 Atrium** Seating available in atrium and BS&BE floors 4, 5, 6 and 7

Afternoon session: Moderator--Cheryl Olman, Graduate Program in Neuroscience

1:00 p.m.	**William Freeman**, Massachusetts Institute of Technology, Cambridge, “Learning to estimate missing high-resolution details”
1:35 p.m.	**Gregor Rainer**, Max Planck Institute for Biological Cybernetics, Tubingen, Germany, “Non-monotonic noise tuning of the BOLD signal to natural images in monkey primary visual cortex”
2:10 - 2:40 p.m.	Discussion and break
2:40 p.m.	**Ron Dror**, Massachusetts Institute of Technology, Cambridge, “Illumination statistics and surface reflectance recognition”
3:15 p.m.	**Daniel Kersten**, University of Minnesota, Minneapolis, “From image statistics to visual inference”
3:50 - 4:30 p.m.	Discussion
5:15 - 6:00 p.m	Weisman Museum permanent art collections open to symposium participants
6:00 p.m.	Dinner in the Weisman Museum Dolly Fiterman Riverview Gallery

Addresses

Basic Sciences and Biomedical Engineering Building
312 Church Street SE
Minneapolis, Minnesota 55455

McNamara Alumni Center
University of Minnesota Gateway
200 Oak Street SE
Minneapolis, MN 55440

Radisson Hotel Metrodome
615 Washington Avenue SE
Minneapolis, Minnesota 55414

Frederick R. Weisman Art Museum
333 East River Road
Minneapolis, Minnesota 55455

Speaker Abstracts

Geoffrey Boynton
The Salk Institute for Biological Studies
La Jolla, California

Using fMRI to compare cortical magnification factors in human V1 to visual acuity

We studied the relationship between visual acuity and areal cortical magnification factor in human primary visual cortex (V1) by comparing Vernier acuity thresholds with retinotopic maps measured with fMRI. Vernier acuity thresholds were measured at eccentricities of 3, 6, 9 and 12 degrees in ten subjects using a staircase procedure and a 2-AFC paradigm. As expected, Vernier acuity thresholds increase with eccentricity in a roughly linear fashion. Area V1 was localized in the same observers by projecting fMRI responses to standard retinotopic mapping stimuli (expanding rings and rotating wedges) onto a computationally flattened representation of the each subject's occipital cortical surface. Next, the eccentricity dimension of these retinotopic maps was carefully measured using flickering checkerboards restricted to annuli of 1.5, 3, 6, 9 or 12 degrees. These annuli alternated with uniform gray fields every 20 seconds. Similarly, the polar angle dimension was measured with flickering wedges presented along the vertical and horizontal meridians in alternation. We quantified the topology of activity maps produced by these stimuli within each subject using a complex logarithmic transform. This provided an estimate of the area of cortex within V1 that represents a given patch of visual space. For each subject's cortical hemisphere, the area of V1 that represents the Vernier acuity stimulus was compared to Vernier acuity thresholds in that subject's contralateral visual hemifield. Our results show that the spatial offset of the Vernier acuity stimulus at threshold is associated with roughly the same distance in primary visual cortex, regardless of stimulus eccentricity.

Ron Dror
Massachusetts Institute of Technology
Cambridge, Massachusetts

Illumination statistics and surface reflectance recognition

How can one distinguish between metal, plastic, and paper from a photograph? If one knew the amount of light incident on the surface from all directions, one could invert the computer graphics rendering process to determine reflectance properties such as shininess and gloss. If one does not know the illumination, however, the problem is underconstrained, because different combinations of illumination and reflectance could produce the same image. For example, a chrome sphere reflects the world around it, so if the illumination were just right, it could appear to be a ping-pong ball. Yet, in the real world, humans effortlessly recognize surfaces of different reflectance.

We have found that the spatial structure of real-world illumination possesses a great deal of statistical regularity, akin to that described in the literature on natural image statistics. We argue that this regularity facilitates recognition of surface reflectance properties by both humans and machines. We have conducted psychophysical experiments suggesting that the human visual system makes use of real-world priors on illumination in recognizing surface reflectance properties. We have also designed a computer vision system to classify reflectance from a single monochrome image under unknown illumination. The system succeeds because of the predictable relationships between surface reflectance and certain statistics of the observed image.

David Field
Cornell University
Ithaca, New York

Natural scenes and the depth of statistical knowledge

Statistical dependencies in natural scenes occur at all levels: from pixel-pixel relationships to object-object relationships. All levels of these dependencies can be shown to important to recognition. Recent work in statistical learning has shown that human observers and infants can learn a wide range of different statistical information . However, the skills in picking up statistical regularities creates potential problems in identifying illusory vs. meaningful structure. In this talk, evidence of these illusions will be presented along with questions regarding the statistics used in identifying information in scenes.

William Freeman
Massachusetts Institute of Technology
Cambridge, Massachusetts

Learning to estimate missing high-resolution details

We address the super-resolution problem: how to estimate missing high spatial frequency components of a static image. We seek to estimate missing high-resolution details from a given low-resolution image. We use an example-based approach, building a database of corrsponding high- and low-resolution image information. The database from a generic set of training images applies well to estimate missing details from images of different subjects. Given a new low-resolution image to enhance, we select from the training data a set of candidate high-frequency patches for each low-resolution input patch. We define a Markov random field for which the most probable state is the desired estimated high resolution image. We find an approximation to the most probable state using Bayesian belief propagation, or other approximations. The algorithm maintains sharp edges, and makes visually plausible guesses in regions of texture.

Jack Gallant
University of California
Berkeley, California

The role of area V4 in natural vision

Natural vision is governed by two critical factors: natural scenes, which have characteristic statistical structure that is exploited by the visual system to facilitate perception; and natural eye movements, which constrain how the scenes are sampled and processed. My laboratory has simultaneously investigated how these factors influence processing in area V4, an important intermediate stage of visual processing. The task requires search for a natural image patch hidden in a large array of similar patches. Eye movements were permitted during search. Arrays were arranged so that whenever any patch was fixated, the receptive field of a recorded V4 neuron would be centered over a different patch. The overall activity of many V4 cells was modulated by the search target; in some cells modulation was additive, in others it was multiplicative. Fourier-domain tuning profiles were also modulated by the search target, revealing top-down influences on spatial frequency and orientation tuning. Finally, the activity of many neurons predicted the direction of subsequent eye movements: the larger the response of the cell, the more likely it was that the next saccade would draw the fovea toward the receptive field. These observations suggest that area V4 acts as a parallel bank of spatially localized multidimensional filters. Their sensitivity and tuning is modulated by top-down influences to optimize performance. During natural vision, filters respond according to their match with the scene and their output, in combination with other areas, is used to drive subsequent eye movements.

Wilson Geisler
University of Texas
Austin, Texas

Natural scene statistics and the evolution of perceptual systems

Recent advances in measuring natural scene statistics raise the possibility of developing more complete and testable theories of evolution by combining precise statistical descriptions of the environment with precise statistical descriptions of genetics. We propose a formal framework for analyzing how the statistics of natural stimuli and the process of natural selection interact to determine the design of perceptual systems. The framework consists of two parts. One is a Bayesian ideal observer with a utility function appropriate for natural selection. The other is a Bayesian formulation of natural selection. In the Bayesian formulation of natural selection, each allele vector (polymorphism) in each species under consideration is represented by a fundamental equation, which describes how the number of organisms carrying that allele vector at time t+1 is related to (1) the number of organisms carrying that allele vector at time t, (2) the prior probability of a state of the environment at time t, (3) the likelihood of a stimulus given the state of the environment, (4) the likelihood of a response given the stimulus, and (5) the birth and death rates given the response and the state of the environment. The process of natural selection is represented by iterating these fundamental equations in parallel over time, while updating the allele vectors using appropriate probability distributions for mutation and sexual recombination. We show that simulations of Bayesian natural selection can yield new insights, for example, into the co-evolution of camouflage and color vision. We believe that Bayesian natural selection offers an appropriate theoretical framework for investigating perceptual systems, as well as many other biological systems.

Sheng He
University of Minnesota
Minneapolis, Minnesota

What adaptation can tell us about processing of natural images?

Adaptation is a powerful mechanism of the sensory system to deal with the unchanging aspects of the changing world. Vision scientists have used this as a tool to study where and how visual information is represented. In this talk I will discuss two adaptation experiments, one psychophysical and one fMRI, both used somewhat unnatural stimuli. In the first experiment (He and MacLeod, nature, 2001), with the help of a laser interferometer, we showed that visual cortical neurons can adapt to gratings of ultra-high spatial frequencies, so high that they normally can not pass the eye's optics. Surprisingly the adaptation is orientation selective, implying that orientation information of unresolvable spatial pattern is registered at the cortex, at least at the input stage. In the second experiment, we studied cross adaptation between luminance and disparity defined patterns with fMRI. A luminance grating is followed by either another luminance grating or a disparity defined grating, in the same or orthogonal orientation. Results show that the different orientation combination generally increased the measured BOLD signal in multiple visual areas in both the luminance-luminance and luminance-disparity pairs. This pattern of results can be seen in V1 and both the dorsal and ventral extrastriate areas. Implications of these experiments on natural image processing will be discussed.

Aapo Hyvarinen
Helsinki University of Technology
Helsinki, Finland

Beyond independence and sparseness in models of natural image statistics

A fundamental model of natural image statistics is independent component analysis which is essentially similar to linear sparse coding. In this talk, I will discuss some extensions and alternatives to these models. First, modeling the dependencies of the simple cell outputs leads to models of complex cells and topography. Second, I consider natural video data, and replace the sparseness criterion by temporal correlations. This leads to similar properties than sparseness: advantages of the two approaches are considered.

Dae-Shik Kim
University of Minnesota
Minneapolis, Minnesota

What does neuroimaging tell us about visual coding strategies?

Cognitive neuroimaging based on functional magnetic resonance imaging (fMRI) provides important information about networks of co-activated areas, thus yielding information about the "where" of the brain's information processing. What is the significance, however, of such localization data for the understanding of the spatial coding strategies used by the brain? We performed three studies to link fMRI in animal and human visual cortices in greater detail with existing body of visual neuroscience data. In the first, "Head-From-Motion" experiment, we could map the detailed, intra-areal spatial structure of human ventral visual areas with respect to different perceived object categories induced by almost identical physical stimuli. In the second experiment in animals, we were able to elucidate the spatial extent to which the observed neuroimaging data might correlate with the underlying spiking and subthreshold activities. In the third series of experiment, we obtained diffusion tensor (DTI) based axonal connectivity pattern in vivo in combination with functional neuroimaging data. The results of our studies suggest that MR imaging at ultra-high fields may provide important insights into the visual coding strategies in vivo.

Daniel Kersten
University of Minnesota
Minneapolis, Minnesota

From image statistics to visual inference

A major challenge for vision is the development of theories that make quantitative predictions of visual performance from natural image input. Such theories require the development of generative models of image formation tailored to the functional tasks of vision. Generative models explicitly represent the factors or causes important to estimate for a specific task. These models also make explicit the information in the image that may confound those estimates. Generative models can either be developed close to the level of the image in terms of sufficient statistcs that support inference (image-based models), or at the level of the objects, materials, and illumination that generate the image (scene-based models). I will describe recent progress in understanding human curve perception that makes use of image-based generative models. I will also show how Bayesian decision theory can be applied to scene-based models to understand how the visual system resolves ambiguity in the perception of surface material color.

Don MacLeod
University of California
San Diego, California

Color discrimination, color constancy, and natural scene statistics

Does the visual system allocate discriminative ability to different regions in colour space in a way that optimizes discrimination among natural colours? If so, discrimination should satisfy "a cube root rule": in an optimized system differential sensitivity will be greatest for the most commonly encountered conditions, dropping to half its maximum under conditions of relative frequency 1/8. Quantitatively, this principle is only very roughly consistent with psychophysical data, but it does account for some of the salient findings, such as the relative sensitivity for different directions in colour space; the fit between theory and psychophysical observation is improved by considering the stimulus to be the local contrast between test field and background, rather than absolute luminance and chromatic values of individual pixels. Comparison with physiological data shows less satisfactory agreement: M cells appear to be too nonlinear, and P cells too linear, for optimal metric representations of luminance and colour respectively. The good colour discrimination of some strongly anomalous trichromats may result from an optimization, during development, of postreceptoral nonlinearity to match the limited range of inputs delivered by the anomalous photoreceptors. For natural colours under natural illuminants, the cone excitations for all surfaces in an image are scaled by approximately the same factor with a change of illumination. This allows the effect of varying illumination to be simply corrected by reciprocal adjustments of sensitivity in the different cone types. The resulting representation is illumination-invariant, but also fails to preserve information about the overall chromatic cast of a scene. Experimentally, colouration of the image is perceptually attributed in part to the illuminant and in part to the viewed surfaces, resulting in "underconstancy". When the statistical variation among natural illuminants and scenes is considered, underconstancy can be viewed not as a failure of constancy, but as a best guess about illuminant colour appropriately based on knowledge of relevant environmental statistics. Natural images do generate small deviations from the scaling principle. These can provide useful cues to the illuminant: statistics (other than the mean) of the distribution of an image's elements in cone excitation space can in principle resolve the ambiguity inherent in the mean alone. Experiment suggests that vision does exploit these cues.and gives them statistically justifiable weight.

Bruno Olshausen
University of California
Davis, California

Sparse coding as a principle of image representation in visual cortex

I shall discuss the principle of sparse coding as it has been applied toward modeling the response properties of V1 neurons, and especially its relation to alternative coding objectives such as independent components analysis (ICA) and temporal coherence. I shall also present evidence from functional magnetic resonance imaging showing that activity in V1 is reduced as a result of perceptual grouping. These results together with the modeling studies suggest that cortical neurons attempt to provide a succinct description of the structures in natural scenes using a small number of active units.

Gregor Rainer
Max Planck Institute for Biological Cybernetics
Tubingen, Germany

Non-monotonic noise tuning of the BOLD signal to natural images in monkey primary visual cortex

The perceptual ability of humans and monkeys to identify natural images declines monotonically as noise is added to the Fourier phase spectrum of the images. Unlike behavioural performance, blood-oxygenation-level-dependent (BOLD) signal levels in the primary visual cortex (V1) of the anesthetized monkey were non-monotonic. The BOLD response had a V-shaped tuning function with local maxima for non-degraded natural images and pure noise patterns (random phase with natural amplitude spectrum), and a minimum for intermediate images with Fourier phase interpolated between natural and random phase. I will describe a simple model, which suggests that the V-shaped tuning function may be a result of a trade-off between the number of responsive V1 neurons and the strength of the response. In addition, BOLD activity in extrastriate cortex will be compared with single unit data recorded in awake monkeys using the same stimuli. Together, these studies aim to shed light on the representation of natural images in visual cortical areas and to advance our understanding of the relation between single unit activity and BOLD signal levels for complex natural inputs.

Pamela Reinagel
Harvard Medical School
Cambridge, Massachusetts

Coding of natural visual stimuli in the thalamus: What exactly is optimized?

Given the task of representing natural stimuli, some coding strategies can be shown to be "better" than others by various different criteria. If efficiency of information transfer has been an important constraint in the evolution of sensory neurons, we should be able to see evidence of this in the structure of neural codes. In the mammalian visual system, there is an information transmission channel between the retina and the visual cortex, passing through a relay in the thalamus, the Lateral Geniculate Nucleus (LGN). At the level of this relay, the neural code is optimized for natural scenes in at least one respect: natural scenes are represented with less redundancy than matched artificial stimuli. Other things being equal, less redundancy is better because redundancy limits the total amount of information that can be carried over a channel. However it does not follow that the LGN actually encodes more information about natural scenes, and I will present and discuss experiments showing that the opposite is true.

Dario Ringach
University of California
Los Angeles, California

Testing theories of natural image representation in primary visual cortex

The idea that primary visual cortex (V1) is designed to efficiently represent natural images has generated specific predictions about the shape of simple-cell receptive fields. Here, I compare some of the theoretical predictions to corresponding measurements of receptive field profiles in macaque primary visual cortex. I will discuss several features of the data that are not explained by the theories and the possible reasons behind these discrepancies. If time permits, I will share some thoughts on how the Hubel and Wiesel hierarchical model of simple and complex cells has influenced our search for a "theory of visual cortex", and how, perhaps, we might have been lead astray.

Eero Simoncelli
New York University
New York, New York

Neural gain control: ecological justification and characterization

One of the most widely noted nonlinear properties of neurons is gain control: the modulation of response sensitivity based on environmental conditions. I'll describe some of our recent work linking gain control behaviors in visual and auditory neurons to the statistics of natural images through Barlow's Efficient Coding Hypothesis. Time permitting I'll also describe techniques for experimental characterization of gain control behaviors in neurons, and the implications of these ideas for perception.

Posters

Poster Session Information

All posters will be presented in the atrium of the Basic Sciences and Biomedical Engineering (BS&BE) Building. The BS&BE is located at 312 Church Street SE on the East Bank Campus of the University of Minnesota.

The poster session will be held on Friday, April 5 from 5:00 to 6:00 p.m. Presenters may set up their posters Friday morning during the registration period or morning break.

We will provide stick pins for you to mount the posters to the board. Please address any questions you might have to the Poster Session Coordinator, Cheryl Olman.

Poster titles and presenters

1

Lateral connections from natural image statistics

Hirokazu Tanaka and Ning Qian
Center for Neurobiology & Behavior
Columbia University

2

Predicting human perception based on a single fMRI time acquisition

Thomas A. Carlson, Paul Schrater, Sheng He
Department of Psychology
University of Minnesota

3

Learning to recognize novel camouflaged objects

Mark Brady, Stacy Ziegenhagen, Dan Kersten
Department of Psychology
University of Minnesota

4

Shape perception reduces activity in human primary visual cortex

Scott Murray and Bruno Olshausen
Center for Neuroscience
University of California at Davis

5

Functional imaging of visual activity during viewing of natural images

Cheryl Olman
Neuroscience
University of Minnesota

6

Perception of the direction of target motion

Leigh A. Mrotek, Martha Flanders & John F. Soechting
Department of Neuroscience
University of Minnesota

7

Coherency: Assessing the quality of a neuronal ensemble representation

Jadin Jackson
Neuroscience
University of Minnesota

8

Retinal contrast coding and natural images: intracellular recordings from cones, bipolar and amacrine cells

Dwight A. Burkhardt, Patrick K. Fahey,
Michael Sikora
Departments of Psychology and Neuroscience
University of Minnesota

Maps

Basic Sciences & Biomedical Engineering

Radisson
615 Washington Avenue S.E.
Minneapolis
Minnesota 55414

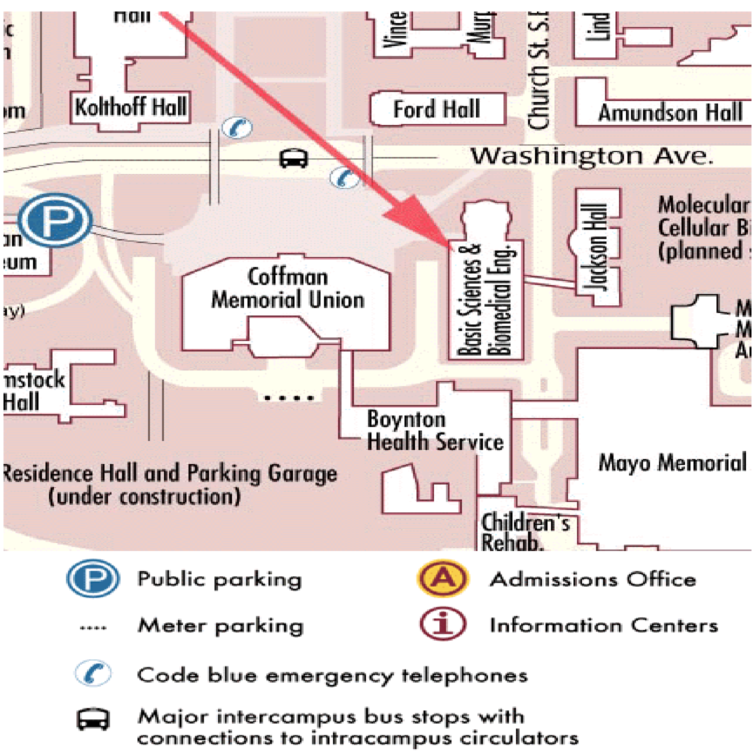

Suggested restaurants within walking distance of the campus

Annie's Parlour
315 14th Ave. SE
Minneapolis, MN
(612) 379-0744
Burgers, fries, malts

Bona Vietnamese Restaurant
802 Washington Ave. SE
Minneapolis, MN
(612) 331-5011

Caspian Bistro
2418 University Avenue SE
(612) 623-1113
Middle Eastern

Loring Pasta Bar
327 14th Ave SE
Minneapolis, MN
(612) 378-4849
American

Shuang Cheng
1320 4th St. SE
Minneapolis, MN
(612) 378-0208
Chinese

Parking

Parking near the Basic Sciences and Biomedical Engineering Building

Basic Sciences and Biomedical Engineering Building
312 Church Street SE
Minneapolis, MN 55455

The building is located on the corner of Washington Avenue and Church Street just east of Coffman Memorial Union.

Parking is available in the Washington Avenue Parking Ramp located next to the Radisson Hotel Metrodome or in the Frederick R. Weisman Garage, 333 East River Road. The entrance to the Washington Avenue Parking Ramp is on Union Street SE.

Parking at the Radisson Hotel Metrodome

Radisson Hotel Metrodome
615 Washington Avenue SE
Minneapolis, MN 55414
(612) 379-8888

Parking is available behind the Radisson Hotel Metrodome. To reach the parking lot drivers enter Walnut Street SE from Washington Avenue.