Journal of Undergraduate Research

University of Rochester

Volume Two
Issue Two
Spring 2004

The *Journal of Undergraduate Research* (*jur*) is dedicated to providing the student body with intellectual perspectives from various academic disciplines. *jur* serves as a forum for the presentation of original research thereby encouraging the pursuit of significant scholarly endeavors.

Acknowledgements

Advisor, Director of Undergraduate Research
Thomas Krugh, Ph.D.

Dean of The College
William Scott Green, Ph.D.

Professorial Review Board
Asish R. Basu, Ph.D.
Andrew J. Berger, Ph.D.
Dana H. Ballard, Ph.D.
Jonathan Geen, Ph.D.
Kenneth Gross, Ph.D.
Daniel G. Harrison, Ph.D.
Michael K. Tanenhaus, Ph.D.

Contributors

Students' Association; Art & Art History;
Biology; Brain and Cognitive Science;
Chemistry; Earth and Environmental Sciences;
English; Economics; History; Music;
Political Science; Psychology; Religion

From the Editor

Often we move in and out of college in such a rush that it seems only a blur. Year after year we emerge from our studies in May only to ask where all the time has gone. What ultimately guides our scholarly endeavors? Dr. Rudolf Virchow, in an address to graduating medical students posits, "Medicine is a social science… if medicine is to fulfill her greater task, then she must enter the political and social life." In truth, this is the case of all fields.

Seeing our respective fields in their societal contexts compels us to be more academically and socially proactive. Research is the foremost socially-conscious activity because it creates knowledge, the primary component of social advancement. We undergraduates must not belittle our accomplishments, for our research resembles graduate level work more and more every year. In the university environment, research opportunities surround us, and we must be earnest in taking advantage of them. Irrespective of one's field, research is the key to fulfilling education's "greater task."

jur began two years ago with a simple, yet ambitious, goal: to encourage scholarly endeavors by undergraduates; today we are proud to say we are not only encouraging but also facilitating undergraduate research. In this issue you will find an expanded "Letters to the Editor" section as well as the new "Undergraduate Research Programs" section. In future issues, we will publish more letters regarding the social, ethical and scientific concerns of research; in addition, each issue will showcase three additional research programs providing the vital details and comments from recent participants of each. You will find an extensive list of research opportunities on the jur website.

This issue of jur extends into the realms of computer science, music, earth and environmental science, optics and American Sign Language. I am also proud to announce that, due to the quality of Rochester's research, and the tireless efforts of our staff, the journal is now a national publication with copies sent to over 125 research institutions.

We thank the Rochester community for its continued support and invite submissions of original research, reviews and letters. For additional information please visit our website, http://jur.rochester.edu, or email jur@mail.rochester.edu with any questions or comments.

Sincerely,

Deepak Sobti, '04
Editor–in–Chief

University of Rochester · Volume 2 · Issue 2 · Spring 2004 jur

Letters to the Editor

Undergraduate Research at the University of Rochester

This issue of *jur* includes seven articles describing research contributions of Rochester undergraduates. Recognizing that these articles present only the "tip of the iceberg" in terms of undergraduate research, the editor invited me, in my role of Director of Undergraduate Research in the College, to comment on Undergraduate Research and Creativity on campus. Rochester faculty have a responsibility to nurture and support creativity in all our activities. Although readily evident in our contributions to our professional areas of expertise, this responsibility is an integral component of interactions with students in all disciplines.

The Rochester Curriculum has been characterized (see the Undergraduate Studies Bulletin) in terms of development of Curiosity, Competence, and Community. It seems natural to use these same terms to describe an undergraduate research project. Curiosity stimulates creative work that requires a level of competence to move forward, and is best done by working within a community. Okay, but "How do I begin?"

The first stage in developing a project is to make a connection with a faculty member with expertise in your area of interest. Such connections often originate from faculty-student interactions within courses or as part of the requirements of a major, and on occasion, even from first-year Quest courses. A faculty adviser is an essential component of an undergraduate research project. You will find a faculty adviser associated with each project described in this issue of *jur*. Once the adviser connection is established, subsequent stages of a project benefit from critical review from the adviser. Thus, the first step, after deciding you wish to engage in undergraduate research, is to refine your area of interest and talk with a faculty member with expertise in this area. Departmental web sites and search engines can be valuable resources in this process.

Once completed, the work should be shared with others, either in the form of a thesis, and/or in a public forum such as our annual meeting on Undergraduate Research & Creativity (URC) held on campus each April. Completing a senior thesis, a senior project, or joining the Senior Scholars program represents a capstone to the Rochester Curriculum. For some, the opportunity to participate in research may come earlier, by way of an Independent Studies course, or a summer research opportunity. Regardless of when you find the time to venture beyond traditional coursework, I encourage you to do so.

jur editors have decided to include two new features in this issue to help readers appreciate breadth of opportunities, and thus stimulate the next generation of student participants. On page 49 we list the members of the class of 2003 who received the honor of "Distinction in Research" as part of their degree. Although an impressive list, you should realize that not all departments choose to bestow the honor "Distinction in Research" so this list is only a fraction of 2003 graduates who justifiably earned distinction in research or creativity. For example, my department (chemistry) awards levels of distinction based on overall performance and chooses not to use the "Distinction in Research" option, and thus chemistry graduates do not make the list on page 49. We recognize this limitation, but include the list to highlight students who did receive the honor.

The second new feature in this issue begins on pages 50-51 where the editors describe three undergraduate research programs: The Camelot Project, The de Kiewiet Summer Research Program, and the McNair Program. A statement from a participant in each program adds first-hand insights into the value of participation. Once again, this list is not meant to be inclusive. Future issues of *jur* will highlight other programs other programs.

—Thomas R. Krugh
Director of Undergraduate Research in the College,
and Professor of Chemistry

Interdisciplinarity at the University of Rochester

We will graduate with an interdepartmental PhD in physics and history. This is not the sum of two doctorates, nor is it a PhD in the history of science. What we have here is a unique dialogue bridging two seemingly disparate and incommensurable departments at the University of Rochester. It is a potentially fruitful enterprise that should add new insights to our current understanding of the development of physics in the 20th century. This is our story of interdisciplinarity at the University of Rochester.

What does the term "interdisciplinarity" mean? To define this concept it is best to start with its root. Discipline is defined by the Merriam-Webster Dictionary as "a field of study" or "a rule or system of rules governing conduct or activity". This agrees fundamentally with the working definition proposed here: a discipline, in academia, is the study of a well-defined subject matter using consistent methods and concepts. From this definition it is easy to extrapolate the definition of interdisciplinarity. This is the deliberate weaving together of diverse or disparate sets of methods and concepts in order to study a well defined subject. The specialization of disciplines began in earnest in the centuries leading up to the scientific revolution of the 16th and 17th centuries and was an underlying framework of all academia by the early 19th century. As our understanding of the universe, including our humanity, evolved, our descriptions and conceptualizations became ever more complex forcing individual disciplines to splinter and become highly specialized fields and subfields. In the 20th century we have witnessed, in parallel, an explosion of specialization and complexity of our knowledge. This complexity is now forcing us to reconsider the 'normal' evolution of this specialization. Is there only room for an evolution along ever finer delineations between fields? Or can we begin to study the complexities of our universe by bridging disciplines and moving towards interdisciplinary studies?

I think the lively debate between the postmodernists and the 'hard' scientists such as Sokal and Bricmont are an indication that this interdisciplinary evolution is in progress and it will not be smooth sailing. In 1959 CP Snow, an English author and physicist, foreshadowed this debate by introducing the concept of the *two cultures* (hard sciences and humanities) on a collision course. It is clear that in this evolution, especially in the 20th century, we have already witnessed many instances of both the positive results of a well done interdisciplinary study (Erwin Schrödinger's *What Is Life?*) and the negative repercussions of a forced, poorly executed, interdisciplinary study (studies by the French psychiatrist Jacques Lacan).

How interdisciplinary are our proposed studies? If we take a purely interdisciplinary study as one where a novel framework and novel methodologies have been established by the marriage of two or more disciplines then I believe our studies will be closer to a multidisciplinary study where instead of creating a new framework, multiple frameworks are used to study the same subject. It seems feasible to create a successful pure interdisciplinary study with two similar fields such as physics and chemistry or economics and political science; however, it seems quite daunting to try and do this with two disparate fields such as history and physics. Ours will be studies that fall somewhere on the continuum between purely interdisciplinary and multidisciplinary.

— Jose Perillan & Drew R. Abrams
University of Rochester Graduate Students

Learning the Language

Interdisciplinary research at the Center for Future Health

jur interviews Philippe Fauchet, Ph.D.

Professor Philippe Fauchet holds appointments as Chair and Distinguished Professor in the Department of Electrical and Computer Engineering, as a Professor of Optics and Biomedical Engineering, and as a Senior Scientist at the Laboratory for Laser Energetics. He is the director of the Center For Future Health, a joint venture between several departments in the College of Arts and Sciences, the School of Engineering and Applied Science, and the School of Medicine and Dentistry. He spoke with *jur* about cross-disciplinary research at the Center for Future Health.

jur: What is the Center for Future Health?

Fauchet: Dr. Alice Pentland, Chair of Dermatology, and I officially created the Center in 1998. It really spans both the School of Medicine and Dentistry and the College. We have participating faculty members from Electrical and Computer Engineering, Computer Science, Chemistry, and other departments as well – we even had somebody from religion at one time – and several of the departments in the School of Medicine and Dentistry; Dermatology, and so forth. We looked at the landscape and found that there are many places in universities and industry where engineers and medical doctors work alongside to develop, for example, a better CAT scan machine, a better MRI machine, a better this or a better that. When you go to the hospital all those gadgets really help you. But we decided that is not what we want because it has been done, so we asked ourselves, 'How could we use technology in a different way?' The analogy that I like is the following: suppose you buy a new car, and that's your body. You buy a new car and it comes with all sorts of advice: every three thousand miles you change the oil, after ten thousand miles you have a major checkup and all that kind of stuff. But suppose that instead of following that advice, all those preventative measures, you just drove your car on the highway at whatever speed you wanted, ignoring the change oil sign, the maintenance, all that stuff, what will happen? Your car is going to have a major breakdown, and then you have to haul it to the garage to fix it. It will certainly cost you a lot of money, and then you start again. Well, that's sort of what we do with our bodies. We just don't take many preventive measures. We don't monitor subtle changes that may happen at a young age, when it doesn't really matter. At my age, it starts to matter ,and when you are sixty, seventy, or eighty years old it really matters a lot. Basically, we only intervene when something is obvious. That's sort of the equivalent of waiting until the car is broken on the highway.

When you bring yourself in for a repair, they use the CAT scan machine, and all those kind of things. I'm not saying those machines should be replaced, I'm saying they should be complemented by something else. That something else is using technology to keep an eye on your health status every day.

Of course, you can't do that consciously. You have to develop a system that, for example, measures simple things like heart rate, and more complicated things like the way you walk. If you do that over a long period of time, we can know what is normal for you, not just for the average population.

What this is really moving towards is medicine based not on the average population, but one that uses you as the reference. You're wearing the sensors, and you can monitor and see subtle changes over periods of time. Instead of waiting until a disease like cancer is so fully developed that you can feel all your symptoms before going to the physician, at which point the physician says "we have to treat you right away," the idea is to develop systems that can alert you to the fact that something is happening, even if you don't realize it. We thought such systems are exactly what technology ought to do, and because it's not that way today, we assembled this group and got a big gift from a major foundation on the West Coast to get these ideas started.

A panoramic view of the Smart House.

jur: What are some of the challenges in working with such a broad array of disciplines?

Fauchet: Well, the first challenge was learning the language. It's tough enough for a mechanical engineer to understand what a chemist has to say, but we had to put the two of them together and then try to talk to a medical doctor. What we have done is have several groups of people start meeting. We have these meetings weekly, with the ground rule of 'no jargon.' If anyone uses jargon, immediately someone yells "What is that?" and whoever said it has to be able to explain it to the other scientists. And, by the way, if you can explain something to somebody from a totally different discipline, then you can explain it to the public, which is a big advantage. Plus, with the explanations comes an appreciation of the other discipline. So I think that's a major accomplishment.

The other challenge is that on the global and national level, we were the first to start a project like this. And now there are at least a dozen groups nationwide and even internationally. I'm not saying that they have copied everything, but they are in the same domain. There are groups at MIT, the University of Virginia, Florida, Georgia Tech, someplace in Oregon, a couple of groups in Asia, and at least one in England. Every time we look around, there are more and more groups that are working on the same thing. It has begun to be accepted as an approach for the future by professional organizations such as the American Association of Housing and Services for the Aged. Those people want the same technology to help them because they don't have the manpower. For that reason, they have endorsed the approach that we began, and now other people know it is the direction to take.

Another one of the big challenges is that to do research, we have to be funded. In academia, you get funding from sources like the National Science Foundation (NSF), the Department of Defense, the Department of Energy, or the National Institutes of Health (NIH). Well, how are we going to find the "in between"? The NIH says what we do is science, so it should be funded by the NSF, and the NSF says the application is medicine and health, so should be funded by the NIH. As a result, all the groups have suffered the past few years from the fact that each agency says, "It's not uniquely in our domain, so you pay for it."

jur: How feasible do you think it is to get this technology into the majority of homes?

Fauchet: Some of it is feasible, not today, but probably within, say, five years. I think the benefit would be for people at high risks. With someone who is at high risk for something, it makes a lot of sense to invest money in monitoring whatever condition they have. There would be a progression of that kind of technology. We have had focus groups indicate willingness to use the gait detection system involving cameras in homes with people who are prone to falling. Of course, people very often don't like to have cameras in their homes. But if the choice is between that and leaving home for good after a fall, they are usually open to it. That's not that difficult, but I think again, the older population is the world that will accept that first. Also, some of those things need clinical trials to demonstrate that there are no dangers when the technology fails.

jur: Once this system is employed, what kind of impact do you see it having on the medical community?

Fauchet: You might think the medical community would say, "Well, this is going to take patients away from me," right? But in fact, the vast majority of the primary care, general practitioners, love this. They hate their job today; they see patients for five minutes and it becomes an assembly line kind of thing. And then they spend more and more time with insurance paper work. They would love to have this technology. Also, they very often will deal with patients who come and the conversation goes, "This is painful." "Well, how painful?" "Well, I don't know, but it's painful." What kind of information is that giving you? Suppose you have a couple of little machines that measure pain; when you come in, you can say, "It's painful, and this is what

Top Left: Cameras looking at a women through a Smart Mirror. **Above**: The Smart Bandage
Top Right: Schematic of a Smart House.

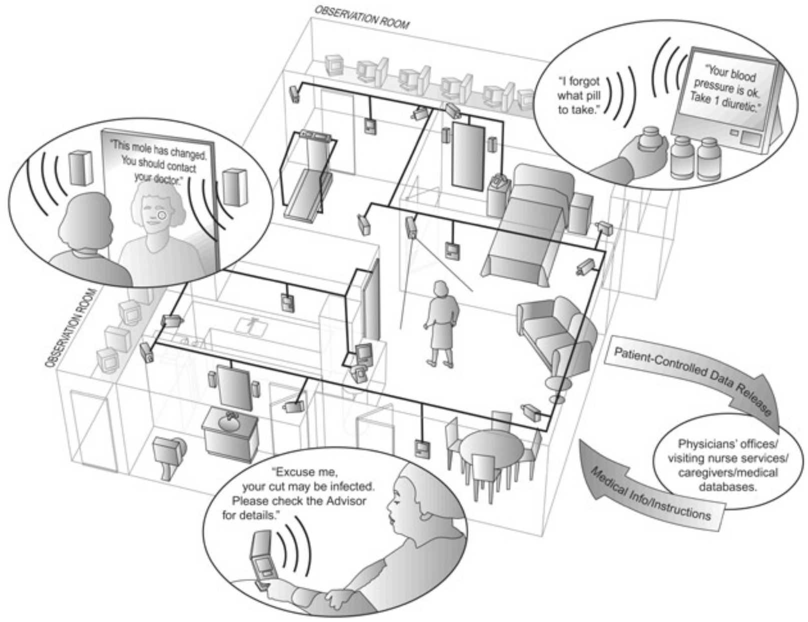

my chip says." Put that information into the computer and the doctor can understand the problem right away, then apply the medical training. Right now, your health is being measured every time you go see a physician, which means you can go a full year or more without being checked. That's pretty stupid. Your health can change drastically during a year. But also, not only do you have regular updates, but fine-grained information is available, and the finer the grain of the information, the better informed your decisions can be.

jur: So it helps to individualize medicine?

Fauchet: That's it! That's the big key. We don't want medicine to be the same for everybody. It's going to be your own thing. What I envision is this, sort of; Radio Shack, or a store like Radio Shack, and you have a shelf with a whole array of devices for your house. And as your house changes and you have some new condition that you need to monitor, you can go buy the appropriate thing. You put it in your house, or a company puts it in your house, and it is integrated with whatever else is already there, just adding more value. That's my vision of it. It may be a service industry like cable TV, you may never own the devices, they may come and install it as you need and you pay a monthly fee.

jur: Can you tell us about one product and how different departments are contributing to it?

Fauchet: I'll give you two examples. One is the gait detection system, where there are several cameras inserted in the ceiling; that's a collaboration between Professor Tekalp who is in Electrical and Computer Engineering and Dr. Mike Berg who is a professor in Neurology in the School of Medicine and Dentistry. They have this joint project because Dr. Berg is in charge of Parkinson's and neurological problems and Prof. Tekalp is an engineer. They work together and supervise the students. Another example is where I collaborated with Prof. Miller, a chemist who is now in dermatology. I'm the device guy and he's the one who makes the smart molecules, puts them in a device and we get a biosensor. Well, five years ago, I didn't know what DNA looked like. But, now I have given talks at the American Chemical Society. It's multi-disciplinary; you learn about a different language and different discipline.

jur: Are there any undergraduates working on it? How can undergraduates start working on it?

Fauchet: Yes, we have undergrads in the summer and a couple during the academic year. We have a couple from MCC also, so the students are not just from U of R. We also have had seniors in Biomedical engineering who have actually started doing work on specific parts right in the center. The way to find out about getting involved is to contact me or Cecilia Horwitz, who is the associate director of the Center for Future Health. There are always possibilities for independent projects, senior theses, or summer employment.

Mabel

Extending Human Interaction and Robot Rescue Designs

Thomas Kollar, Jonathan Schmid, Eric Meisner, Micha Elsner, Diana Calarese, Chikita Purav, Jenine Turner, Dasun Peramunage, Gautam Altekar, and Victoria Sweetser

Advised by Chris Brown, Ph.D.

Department of Computer Science

Artificial Intelligence (AI) has at its core the creation of intelligent systems. The holy grail of this entire subject is to create a general form of intelligence that can reason, learn and interact intelligently in its environment much as a human does. However, the solution to this problem has turned out to be more difficult than anyone could have imagined. In this project, therefore, we have tried to move toward a better understanding of what it takes to create a truly intelligent system.

Mabel (the Mobile Table) is a robotic system developed primarily by undergraduates at the University of Rochester that can perform waypoint navigation, speech generation, speech recognition, natural language understanding, face finding, face following, nametag reading, and robot localization. Mabel uses these components of intelligence to interact with its environment in two distinct ways. It can act as a robot host at a conference by giving people information about the conference schedule. Mabel can also act as a search and rescue robot by entering a mock disaster scene, locating mock human victims, mapping the victims' location, and returning safely out of the scene. Mabel was the winner of the robot host event and tied for third place in the robot search and rescue event at the 2003 International Joint Conference on Artificial Intelligence (IJCAI) in Acapulco, Mexico.

The overall design philosophy for Mabel was to integrate human interaction with robotic control. We achieved this goal by using an interactive speech system, a pan/tilt/zoom camera that actively follows faces, and another pan/tilt/zoom camera that reads patrons' nametags. In this paper, we present the changes that have occurred since the 2002 American Association for Artificial Intelligence (AAAI) conference and the methods that we used to integrate the hardware and software into a usable system.[1]

Robot Host

A general overview of the architecture for the information-serving robot can be seen in Figure 1. One can see that there are multiple layers in this system: the sensor level, the interpretation level, and the decision level. The sensors include sonar, cameras, keyboard, the dead reckoning position and a microphone. The drivers for using these sensors, as well as a library for writing behaviors, were provided with the ActivMedia robot.[2]

We produced a map of the competition areas in advance using a tape measure. On this map we drew a graph of places that the robot could go from any given place. This map was not only used to visually direct people from one place to another at the conference center, but it was also used for localization purposes.

At the next level, the interpretation level, the programs interpret the sensors to provide high-level assessments about the environment and intentions of the person with whom the robot is interacting. For example, if the vision component detects a face in the scene, it reports the person's location and name, based on his nametag. The intentions of the patron can also be determined from dialog. For example, the robot can terminate a conversation if a patron says "goodbye."

At the top level is the control aspect (which is basically a finite state machine). If the robot detects someone in front of it then it starts dialog with him. Our system can detect a person's presence if he speaks, types, or if his face is being tracked by the camera. If a person is not detected by the robot then it will wander around to various points on our map of the competition area searching for people, all the while the robot is keeping track of its location (using Monte Carlo Localization, which is discussed later). Should it find a person, then it will continue driving toward him until it successfully reaches him, or, if unsuccessful, it will go back to wandering to various points on the map.

When a person interacts with the robot, many things can happen. Immediately, the robot starts trying to find the person's nametag. Once it has attained that information it uses the information when speaking or displaying other information on the screen. At the same time, the robot continues to track the patron's face with the higher camera. Anytime the patron speaks to it, Mabel parses the words using Sphinx speech recognition then uses a Bayes net to decide what the person's intentions are when talking to it.

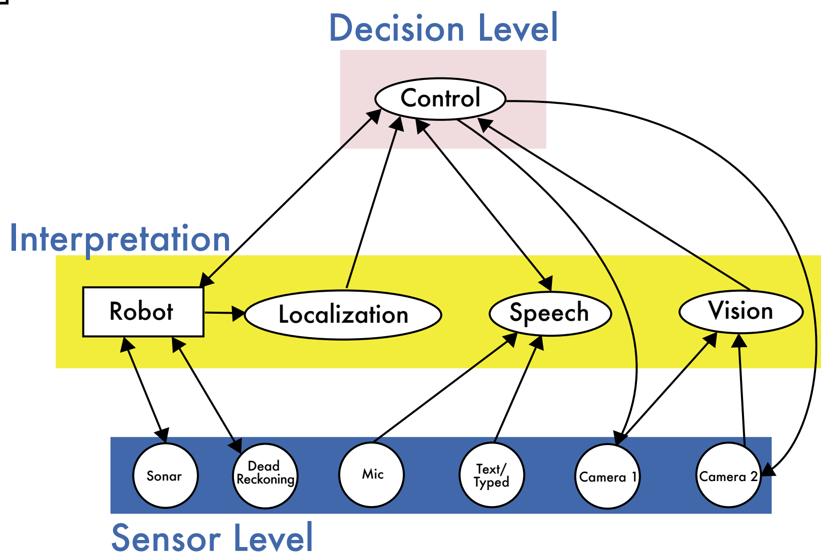

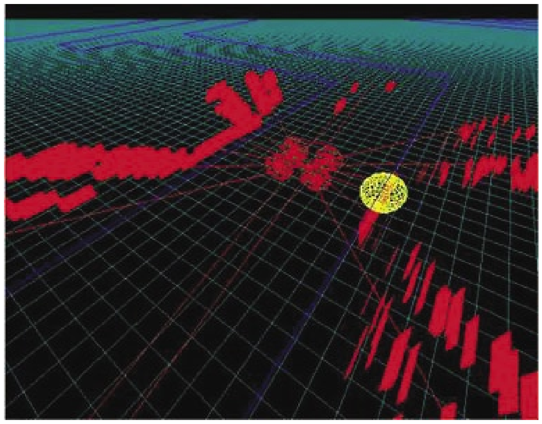

Above, Figure 1: The Mabel system structure consists of three layers: sensory input and output, interpretation of sensor data, and decision-making about what to do given the analysis of input data. **Above Right, Figure 2**: Real-time 3D mapping aids the teleoperator in robot search and rescue. The red lines are sonar readings and the yellow sphere is a mapped victim. **Right, Figure 3**: Overview of the robot search and rescue system.

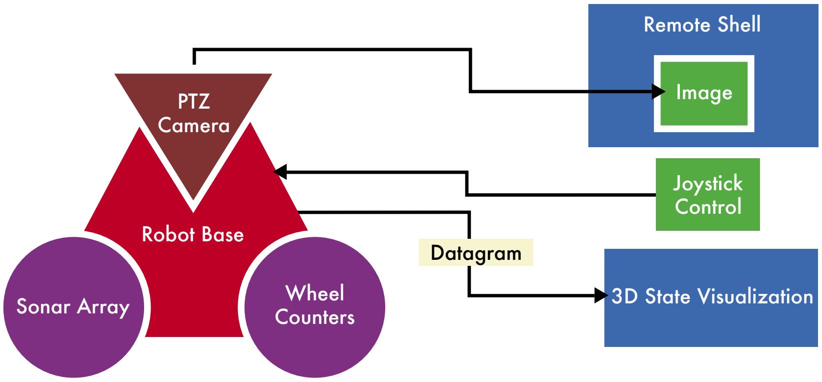

Given this information, the robot can perform a query on a database of events and produce a response about an event, speaker, time that an event occurred, time of day, or date. Moreover, if any events matched the query, then they would appear on the screen.

Robot Search and Rescue

Using the robot search and rescue system, one can teleoperate the ActivMedia Pioneer 2AT robot to locate and map victims in a mock disaster scene and provide a map of victim locations to mock rescuers. The robot and its sensors are the only things that the teleoperator can use to navigate the scene. Points are awarded in the competition for successfully locating victims and creating useful maps of the disaster area. Points are deducted for making contact with victims and for interacting with the environment in a dangerous manner, i.e. causing secondary collapses of structures.

Our system consists of three core components: the control, mapping, and joystick applications. An overview of this system can be seen in Figure 3. The control application, which interfaces directly to the robot's actuators and sensors, is responsible for obstacle avoidance, linear and angular velocity control and communication of sensor information back to the user workstation. This program runs completely autonomously on the robot and there is no need for the intervention of the teleoperator.

The mapping application is one of the teleoperator's tools for retrieving information from the disaster scene. It provides the user with a live interactive 3D rendering of the robot environment by mapping the raw sonar readings and the internal robot position $<x, y, \theta>$ to a 3D world. One can see an example of the mapping interface, as rendered by OpenGL, in Figure 2. The teleoperator is also able to load maps of the environment or save maps for printing.

Monte Carlo Localization

Localization is the problem of determining the spatial position of a robot from its sensor readings (vision, sonar, and dead reckoning). It is implemented here using a technique called Monte Carlo Localization (MCL).[3] As in other works, we are only concerned with the readings from the robot's wheel encoders and range devices.

There are various forms of localization that vary in the ambition of the problems that they attack. Local techniques aim to correct drift from the inaccuracy of dead reckoning (the most well-known of these are Kalman Filters). Global techniques aim to locate the robot even when catastrophic things happen. There are two problems usually solved by the global techniques: the wake-up robot problem and the kidnapped robot problem.[3,4] The wake-up robot problem occurs when the robot is given no initial information about its position and is expected to localize itself in some environment. The kidnapped robot problem occurs when the robot is carried to another location during its operation and is expected to recover from this strange occurrence where the dead reckoning does not match up with its range data. Thus, global techniques must be more powerful than local ones.

Localization is often considered to be

one of the most fundamental problems of mobile robotics, since without it a robot will have an inaccurate internal representation of its current position.[3,4] This can lead to some disastrous consequences, especially if there are stairs, ledges, or other obstacles that cannot be readily seen by the range devices and which could have been avoided if the perceived and actual positions coincided.

MCL is a relatively new method that provides a solution to the global localization problem. We implemented the MCL algorithm with distance filtering as per other studies.[3] The major benefits of using MCL over the other approaches are that it can globally localize the robot, it greatly reduces the memory requirements in comparison to Markov Localization, it is more accurate than Markov Localization, and it is implemented easily.[4]

For MCL, we divide the new sensor data into two groups: a new odometry reading and a new range sensor reading. $S = \{s_i\}$ such that $i = 1, \ldots N$ is a set of N weighted and random samples over a space. For example, in our implementation of Monte Carlo Localization this distribution would initially be uniformly distributed. Each sample is a two-tuple with a pose and a probability of that pose. Thus, a sample is: $<< x, y, \theta >, p >$. Moreover, we assume $\Sigma p_i = 1$.

Each time the robot obtains a new odometry reading a, MCL generates N new random samples that approximate the robot's new position. Let s_i be a given sample from S and l' denote its position, so that "each sample is generated by *randomly* drawing a sample from the previously computed sample set [S], with likelihood determined by their p-values [probability]."[4] Thus, the value for the new sample's l can be generated by sampling according to $P(l|l', a)$.

In our implementation this probability ($P(l|l', a)$), often called the motion model, is computed by taking the ideal movement of any sample, and sampling a new position with a probability coming from a Gaussian distribution in the *rho* and *theta* components of this movement. The standard deviation and the mean of this distribution were computed from experimentation.

Moreover, for a new sensor reading s and a normalization constant α that enforces $\Sigma p_i = 1$, we re-weight the sample set S. We let $< l, p >$ be a sample and we recalculate p such that $p \leftarrow \alpha P(s|l)$. This can be done in $O(N)$ time.[4]

The probability $P(s|l)$ is called the sensor model. In our implementation it is merely a Gaussian distribution over the ideal sensor reading of the robot given that sample's location. The $P(s_i|l)$ is measured by ray-tracing to get the ideal reading of the *ith* sensor and then sampling from the Gaussian using the difference of the ideal and measured readings. Moreover, we integrate these probabilities for each sonar by multiplying them together. This integrated probability gives us $P(s|l)$. The standard deviation and the mean of this Gaussian distribution were derived from experimentation.

There are some cases where we run into problems with Monte Carlo Localization. First, since MCL uses finite sample sets, it can (and sometimes does) happen that none of the samples are generated close to the robot position.[4] This causes the robot to lose its location, never to recover. There are some techniques to prevent this from happening, but many of these methods are not necessarily mathematically sound.[3] The solutions usually involve introducing artificially high amounts of noise into the sensor measurements or by generating samples near the most recent sensor reading. We use a simple technique of generating random samples around a sample s_i using a progressively larger Gaussian distribution, should the hypothesized position move into an unmodeled space. We also take some uniformly distributed random samples on each round of MCL.

There is also the problem of the dynamic environment. Regarding our assumption that the environment is static, "clearly this conditional independence can be violated in the presence of people (which often block more than one sensor beam). In such cases it might be advisable to sub-sample the sensor readings and use a reduced set for localization."[3] Thus, we sub-sampled the sonar readings using a distance filter to get the readings that correspond to the world and not the dynamic obstacles in the world. Thereby the sensor readings that come from unmodeled obstacles were ruled out.[4]

Finally, there is the problem of the map that is used as a model of the world. In our system, we used a simple mapping system that only allows lines. Moreover, we had no way of automatically generating these maps from real data. In other words, we used a tape measure to make the maps by hand. Now when one is using a simulator there are no problems and everything works as expected, since your model of the world matches exactly with the world used by the simulator. However, when one is using MCL with real world data, then the data will often not match the modeled one due to sometimes quite large errors in generating a map by hand. This is a problem that we did not have time to solve, and thus our algorithm would have many problems when working in the real world. In the simulator, however, the convergence and tracking of MCL worked very well.

Vision

Mabel's vision system is focused on finding people in the environment, and reading their nametags. Both techniques make use of a general filter cascade architecture that searched for a pattern throughout the image using a constant size detection window. The filter cascade detects a single rectangular bounding box around the largest target pattern in the source image, which can then be tracked over multiple frames with an Alpha-Beta filter to smooth over any noisy or incorrect detections.

A filter cascade is a search technique that focuses its searching effort on small areas of the image that show initial promise for matching the target pattern. These small areas of the image are called sub-windows and the initial promise of a sub-window is determined by the success of the first classifier. If a sub-window of the image passes the initial classifier, then additional classifiers are applied to it.[5] We apply a cascade of these filter-classifiers to each sub-window, testing for the presence of several simple rectangular features. While most filter cascade architectures to date have utilized only features within grayscale images, we apply color models to create additional feature channels as shown in the two central images of Figure 4.

While searching through a feature channel image, the sub-window often detects a spatial pattern in several adjacent locations. Also, the target pattern might appear in more than one location in the image. Jones and Viola perform a disjoint-sets algorithm to combine locally adjacent detections and select the largest target region. Our filter cascade algorithm instead fills in each detected rectangle with 1s in an initially empty binary image. This binary image shares the same dimensions as the signal images and the starting color image. A contour finding algorithm is applied to the binary image. The largest contour from this detection image becomes the filter cascade target. Adjacent detections from the source image overlap when filling the binary image, and thus form only a single contour. See the bottom of Figure 4. Spurious detections from other similar objects tend to create smaller contours and are thus often ignored.

The person-finding algorithm used both a skin colored channel and an intensity channel (see the right of Figure 4) for locating faces in the detection window. To generate the binary skin channel (where 1s represent skin pixels), we test for the presence of each pixel from the image in a binary Hue/Saturation model (see the left of Figure 4).

We first generate this model from images of human skin tones captured in the environment by using a series of pictures that were taken of various individuals. The skin regions of each individual were then isolated. To make the model robust to all individuals, a sample of different skin pigmentations was carefully selected. From the training set of skin images, the value of the color in each pixel was calculated using the HSV scale and plotted on a Hue/Saturation graph. We save this as a bitmap so that we can fill in missing skin tone regions using a standard image-editing program. This improves the robustness of the skin detection.

The first level of the filter cascade for faces drags a sub-window throughout the binary skin-tone map. Sub-windows are eliminated as possible faces if the sum of their pixels is not over 25% of the sub-window, a result that would suggest that there are insufficient skin pixels on the object to deem it a face. The second and third levels both operate on a grayscale intensity image. In the second level of the filter cascade, we look for the eyes–a trait that distinguishes faces from other objects in most cases. We look for the eyes by (1) summing the number of skin-tone pixels within a rectangle covering the forehead, (2) summing the number of skin-tone pixels within a rectangle covering the eyes, and (3) subtracting the result of the second step from the result of the first step. If the result of the subtraction is a large positive number (thereby suggesting the rectangle covering the forehead and the eyes describe two dissimilar entities), then we gain confidence that the sub-window spans a face. If the sum tends to be near zero, then we lose confidence that the sub-window spans a face and we terminate the cascade. The result from step 3 must constitute at least 8% of the sub-window to allow the cascade to continue. Intuitively, this requirement captures the idea that the forehead should consist of many skin tone pixels and the eye region should consist of no skin-tone pixels (thereby producing a high number in the subtraction of step 3). In the final level in the cascade, we compare a rectangle covering the area of the lips to the chin below in a similar manner as above. In this case, the chin region is subtracted from the darker lips region above it. We have empirically found that this technique works well.

The nametag reading process employs two different zoom levels using a Canon PTZ camera. The central control system activates the tag reader during the person approach phase. The tag is found at this outer zoom level using the filter cascade, and the camera is centered on its position. When the alpha-beta filtered tag location is centered in the image, the camera zooms in. When the zoom is completed, the image resolution is increased from the usual 180 x120 to a full 760 x 480. Ten frames of the nametag are stored at high resolution and read by the Abbyy Fine Reader engine. The most frequent occurrence of the first two string tokens is assumed to be the person's first and last names.

The nametag filter cascade consists of four levels, each level paying close attention for the presence of a particular feature. The cascade's first level eliminates all filter sub-windows that lack a high percentage of white pixels. In most instances, the first level eliminates close to half the sub-windows in the image, thereby narrowing down our options significantly.

The second level checks for the presence of a colored bar at the top of the nametag (see the bottom right of Figure 5). We distinguish a bar from other objects through the use of a hue/saturation color model for the bar, which was created by sampling several images containing the bars. Once the second level is complete, the majority of the sub-windows are centered vertically on the tag, but remain uncentered horizontally.

The third and fourth levels attempt to horizontally center sub-windows on the tag. The third level begins the process by (1) summing the pixels in the sub-window containing the tag text, (2) summing the pixels in the sub-window to the left of the tag text, and (3) subtracting the result of step 2 from the result of step 1. If the result of the subtraction in step 3 is a high number, then we gain confidence that the sub-window is adequately centered from the left of the tag. If the result of the subtraction is a low number, then we lose confidence that the sub-window is adequately centered from the left. In the fourth level of the operation, we perform the same series of steps as done in the third level. This time, however, we consider the sub-window to the right. If the cascade passes the fourth level, then we can be reasonably confident that the sub-window is centered on the tag.

In order to ensure quality text reading, we implemented rotation invariance on the high resolution frames of the nametag. This is done by calculating the angle formed between the blue rectangular bar and the bottom of the image. The image is rotated to cancel the calculated angle.

Speech

Interaction is initiated whenever the robot recognizes that someone is in front of it, when someone speaks to the robot loudly enough to trigger speech

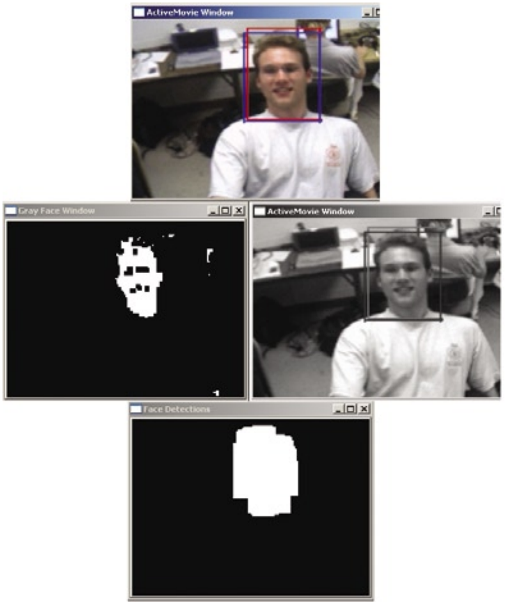

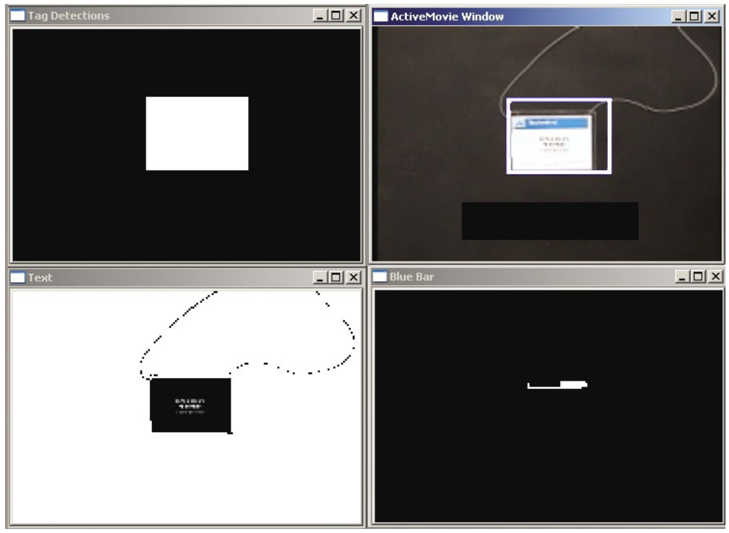

Left, Figure 4: The process of finding a face. **Above, Figure 5**: The process of finding a nametag.

recognition, or when someone types or clicks in the GUI. The communication system never initiates utterances of its own; it responds to user's input by directly answering their questions, then waits for more input. Interaction is terminated when the person stops using input functions and walks away.

The robot handles typed input by filtering out keywords relevant to the conference, then using a Bayesian tagger to find the type of information the user is asking for.[1] The graphical interface transforms clicks in the window to tags and filtered output identical to the result of Bayesian filtering. The system then processes both typed text and clicks identically.[6]

After tagging, the robot immediately handles some conversational pleasantries with randomly chosen rote responses. This allows it to deal with 'Hello', 'Thank you' and other statements that are not requests for information. Real questions are turned into SQL queries that retrieve information from a database, which was hand-constructed from the conference schedule.

The robot uses graphical display, text display and text-to-speech as output modes. Graphical output is an unranked table containing the information the user requested. The user can click on any item in the table for more information. Text display uses a natural language generation algorithm. This algorithm first constructs a core sentence of the form:

< speaker >< verb >< event >

It guesses the verb by examining the event, and if there is no speaker, it uses an expletive construction such as 'there was' to mimic the same basic form, then it recursively adds prepositional phrases until it has used up the remaining information. If there is too much information to fit on the screen, it omits some and adds a phrase containing 'more' or 'others' to the end of the sentence. Text to speech speaks the same output as is displayed graphically using Microsoft Speech SDK.

Because of the technical problems inherent when using speech recognition in noisy settings with unfamiliar speakers, speech recognition was handled separately. It uses CMU Sphinx.[1]

Conclusion

This is the second of the Undergraduate Robot Research Team's papers, which now have appeared two years in a row in the AAAI robotics workshop proceedings. Throughout the years the personnel has been diverse, having an unusually high proportion of women for computer science fields. Another unique aspect of this project has been to have a student-run class created with the help of Professor Chris Brown's continuous and overwhelming support. The result has been an incredible experience for undergraduates: an opportunity to work in a team environment, to work on unique research, interact with professors, and to gain personal recognition while an undergraduate.

Acknowledgments

Thanks to Abbyy Finereader for providing their optical character recognition at a greatly discounted price. This research was funded by NSF Research Infrastructure and associated Research Experiences for Undergraduates grant "Spatial Intelligence," number EIA-0080124.

During the fall of 2002 and the spring and summer of 2003, there were about six people working on the project full time. Thomas Kollar focused on localization, navigation, speech recognition, TTS, robot search and rescue, and global integration. Eric Meisner and Chikita Purav also worked on the MCL algorithm. Jon Schmid, Dasun Peramunage, and Gautam Altekar focused on the vision components of the robot. Micha Elsner, Diana Calarese, and Jenine Turner worked on the natural language understanding and GUI. Eric Meisner worked on mapping and robot search and rescue designs.

Painting of Siva decorating the inner precinct gallery of the Brhadīśvara Temple.

Sonic Architecture and the Creation of Space

Approaches to Dīkshitar's Brhadīśvaram Kriti and the Brhadīśvara Temple in Tañjāvūr

Alfred Vitale, 2004

Advised by D.R. Brooks, Ph.D. (Religion) and Robert Morris, Ph.D. (Eastman Composition Dept.)

Departments of Religion and Music

Introduction: A Subject, An Object, A Relationship

Musical forms, like architectural structures, can describe our experience of spatial relationships. Our study will demonstrate this through a particular musical construction of the Carnatic (South Indian) tradition that is further enhanced by the composer's religious perspective. The piece is a composition ("kriti") by Muthusvāmy Dīkshitar (1775-1835) entitled "Brhadīśvaram." The title refers to a particular temple, the Rajarajesvara Temple in Tañjāvūr, Tamilnadu, also known as the "Brhadīśvara Temple." Our purpose here is to consider Dīkshitar's profoundly Tantric conception of space and relationships in this kriti and see how the musical structure both mirrors and embodies the creation of temple spaces delineated by the relationships of the worshippers to the Divine. It is clearly reasonable to perform such examinations through the windows of musicology, religion, or anthropology, but we should also open up the possibilities for other disciplinary engagements that can perform comparative or supportive functions. Among these paradigms there is a commonality: each seeks to reconcile, address, or otherwise elucidate on the relationship between the subject and the object; the perceiver and the perceived. Since this commonality aims at the heart of the matter, we will treat these paradigms as possible windows through which to examine this composition. It is the hope of this paper to make a reasonable case for such a comparative approach to understanding the relationships and spaces within "Brhadīśvaram".

Let us begin our work with one of these models in order to create a framework from which to place some useful ideas.

Creation of Space: The Metaphor of Architecture

Architecture is the creation of space. The physical materials we think of when we think of architecture are only carefully formed metaphorical devices. For how else can we define space? When we look at a landscape, we see something versus something else; a duality between what we consider tangible and that which is everything else. The "everything else" is that which we notice but cannot describe. Architectural forms are constructed frames around these ineffable spaces. By making such structures, we make the indescribable approachable.

But there is yet another important feature of architecture that is of utmost importance in our examination of Dīkshitar's composition. This feature, the relationship of perceiver to the perceived, is in fact the very motivation of cognizable architecture and defines its structural relationships. As the philosopher Martin Heidegger wrote, "The bridge gathers the earth as landscape around the stream…It does not just connect banks that are already there. The banks emerge as banks only as the bridge crosses the stream."[1]

In Heidegger's view, it is the disconnected elements of the natural world that is re-cognized (they "emerge") as elements of the architecture when the architecture takes its particular form.

If we consider architecture not as the building of structures but as the creation of cognizable space, we are also given a useful perspective in which to examine Dīkshitar's remarkable skills as both a performer and as a Tantrika. For, as we will see, the delineation of space is a particularization of experience that is crucial to the Tantric perspective.

The Hindu Temple: Monument of manifestation…

The gargantuan task of ascertaining the meaning of the numerous aesthetic elements in and around the Hindu temple would not be productive here. Yet it should suffice to say that as with most aspects of Hinduism, there is very little that can be done to extricate the religious from the mundane. It can be safely asserted that the Hindu temple holds a prominent place in the social and religious milieu of India, and that it is a regular feature of existence for Hindus.

We can also assert that there would most likely be a deeper significance of Hindu temples to the devotees whose lives were intertwined with its existence. Depending on the devotee's particular understanding, the Hindu temple fulfilled many roles on many levels. Stella Kramrisch, in her book The Hindu Temple, touches on this idea: "The temple is the concrete shape of the Essence; as such it is the residence and vesture of God. The masonry is the sheath (kosa) and body. The temple is the monument of manifestation."[2]

This is an elegant passage that does not miss the point that there is more to the Hindu temple than its visible architecture; but the visible architecture is that which presents the means to appreciate its depth. Let us look at an example of this approach in an article on a particular sculptured element of Hindu temples. M.A. Dhaky writes that "…the bhutas (elementals) or, more precisely, bhutanayakas (captains of the elementals)… represent more than the five elements; they personify the elemental fragments of creation, infinite in number."[3] He then elaborates further on their meanings:

The bhutas (or bhutanayakas) and ganas (or gananayakas) on Saiva temples transfer the perception of the subjective-objective reality of Creator-Creation or Causation-Causality to a tangible plane, through the symbolism of their concrete imagery.[3a]

If there is this much complexity and depth to the consideration of a temple's simpler elemental features, it is not difficult to see why philosophical and metaphysical elaboration might have be rightly relegated to the domain of the Tantric or Yogi. This is not to say that such understandings were out of reach for everyone else, however. What we will now examine is how a devotee's conception of the divine is informed by their relationship to the temple.

The Temple and Me: Approach and Consideration of Relationships

Looking far ahead to the edge of the village; the rising sun causes a boldness of colors to burst from the natural landscape, and a devotee sees the familiar shape of the temple as it punctures the sky. The devotee apprehends the horizon and the shape made where the earth meets the sky. The temple is as natural and evident as the trees and hills, and for the time they are all one singular expression of Brahman. The devotee walks with a quick but focused pace, heading up the road to the temple complex. It looms larger as he nears; his footsteps mingle with the other footsteps that dance around his own. And soon he is standing in front of the entrance, as close as he can get before going inside. He notices the tarnished statues of Śiva Nataraja and TripuraSundari and many others, his eyes roam across the contours of the temple and rise up with its many levels. The shapes and colors and textures are creating an overwhelming sensation; they demand his action. He must enter the temple or else he must leave. As he crosses the threshold, the world of the vibrant sun and of the temple shape have collapsed into a new set of experiences. The devotee is where he is supposed to be and doing what he is supposed to do. As he performs his puja, he is aware that he is truly inside and that it is here where he can experience the promises of the temple. His puja finished, he re-emerges from the temple and enters the world again, startled by the action around him. He regains his composure and travels forth into the outer; into the world of actions and people until at last he watches the sunset over the temple and how natural it all looks.

The above was included to illustrate something important about the temple: it gains different consideration and different meanings as the devotee approaches and as the devotee exits. The philosopher Emmanuel Levinas said in an interview, "Human experience is not some self-transparent substance or pure cogito; it is always intending or tending towards something in the world which preoccupies it."[4]

Humans, then, are perpetually motivated to re-establish (re-cognize) their relationships with the world around them (and as we will see, within them as well). The temple establishes a useful model to examine just what these relationships are and how they may be established and maintained.

In the fictional devotee story above, there appear to be four distinct phases that occur in his perception. Actually, there are only three. The fourth experience, when he re-emerges, is a continuation of the third phase and merely starts the process all over again. So how are these phases considered? Let us examine them now in more detail.

Phase One: Unity

Those who saw the temple from a distance consensually perceived it as a Single Entity, the Purusa or Universal Self who is reality (sat), Consciousness (cit), and Infinite Bliss (ananda); who is also beyond sky or space…[3b] The temple in the distance is unobtainable; an impression. But it is a natural feature of the landscape, which is all one organic whole. The distant temple is a space where the seemingly infinite potential meanings outweigh any particular assumptions. It remains ineffable and unreachable.

Phase Two: Plurality

Those who viewed the temple at close quarters saw in its organization and its stratified divisions, its details, voids, and masses, the embodiment of Prakriti or Nature—Cosmos, Creation, Manifest, or Empirical Reality—with its interminable, though coherent, amalgam of tangible and intangible, seen and unseen, sensed and unsensed verities.[3c]

It is in this phase that we come to understand bhakti. The distance between devotee and temple is no longer one of unity; it has become a duality. Here the confrontation with the temple emphasizes the issue of separation, because there is no synthesis of the inner with the outer. There is always the distance between the subject and object, no matter how much contact there may be. In a sense, one can only get close to the temple; one cannot internalize the temple and thus a crisis occurs. The crisis forces one to negotiate a relationship with the Divine.

Consequently, it is only when one term of the dialectic is upset that the part usually played by both terms becomes visible, revealing simultaneously that direct power which the world holds over our body and the reciprocal power which the body has in anchoring itself in a world, in demanding 'certain preferential planes'.[5]

The temple's entrance, as a threshold, is also seen as a boundary: "A boundary may also be understood as a threshold, that is, an embodiment of a difference."[6] And as a boundary, it implies selectivity and further forces a negotiation: "Perhaps the time/place of the templum is the time/place of a threshold that cannot be crossed or erased: something like an invisible sieve…a filter that allows the eye to see."[7]

Yet it is precisely the power of a boundary that creates the motivational tension necessary to experience something. Heidegger states: "A boundary is not that at which something stops but, as the Greeks recognized, the boundary is that, from which something begins its presencing."[1a]

Phase Three: Transcendence and Re-Cognition

There is an implicit relation between revelation and concealment upon crossing the threshold, which becomes solidified as an explicit relation when the temple is re-considered by the devotee from inside. Here, the duality between inner and outer temple is resolved. However, once inside the temple, an even deeper negotiation of relationship between temple and devotee occurs. What is at stake is the experience itself, whereas in Phase Two what was at stake was the nominal form of the relationship.

The inner metaphors create particular experiences of space, and the devotee recognizes that the temple, with its great immensity of space, becomes metaphor in order to reveal its deepest secrets. It also becomes fully tangible in order to resolve the problem of subject and object. As Dhaky writes, "They seem inspired by a vision, directly perceived by the senses, of the prasada, or temple as a concrete symbol of both the dualism and monism of all that exists and of the ultimate which subsists…"[3d]

It is vital that the devotee realize the ineffable in the tangible form of the temple, but it is even more important that the devotee internalize this revelation.

…the unity of experience can no longer be considered to lie 'out there' or 'in here' but must, rather, originate in that dynamic relationship between body-subject and world through which 'objects' and 'subjects' come into being for us.[5a]

The experience must be personalized and particular, giving the devotee the full awareness of his existence both in the temple and outside the temple:

In itself, it is present in the inconspicuousness of things at hand being taken care of by a circumspection absorbed in them for that circumspection. Space is initially discovered in this spatiality with being-in-the-world. On the basis of the spatiality thus discovered, space itself becomes accessible to cognition.[1b]

So what we find is that the temple is not just the external object, it is also the very relationship between the subject and the object, as well as the considerations of this relationship and the roles imparted:

Instead of a mechanistic deterministic relationship of causality, we have an organic relation of motivation between subject and world, such that the body possesses the world in a certain way while gearing itself to that world….[5b]

The Kriti and Me: Approach and Consideration of Relationships

Turning to the structure of the Kriti form, of which the Carnatic composer and musician Muthusvāmy Dīkshitar was a master, we are confronted with perspectives and considerations that are analogous to those shown in the previous section on temples.

The kriti, in general, has three parts (angas): Pallavi, Anupallavi, and Charanam. A brief comparison may suffice to show how the relationships between devotee and the temple as expressed earlier are found within the kriti structure.

The Pallavi ("sprout; bud") is a simple statement usually introducing the contemplation of the supreme divine presence in the kriti. In comparison with perception of the temple form, the pallavi begins as the outward temple seen from a distance and is oriented towards a single, broad experience of the Divine in a superlative form and within a broad space in which all is contemplated as one organic whole.

The Anupallavi illustrates the relationship with the Divine; it tells us what is at stake and provides the definition of a relationship between the worshipper and the worshipped. This is the temple experience up close, when devotees are recognized in the relationship, and we find an emphasis on the plurality of temple details. There is a new delineation of space as being the connection between the devotee and the Divine. The devotional contract is established.

Finally, in the Charanam, the action is detailed and the instructions are unfolded; it becomes the place of real secrets and real revelation. Returning to the perception of the temple, the Charanam is the point where the temple is entered and worship is performed. It is here that the method of experiencing the divine is shown to be personal and particular. The space is now delineated from the inside out; from the inward experience to the outward.

Now that we have a simple understanding of this form, we can examine the "Brhadīśvaram" kriti in more detail in order to plunge deeper into its meanings.

Dīkshitar's Re-Creation of Space: The Metaphor of the Brhadīśvaram Kriti

The Rajarajesvara Śaivite Temple in Tañjāvūr, also known as the "Brhadīśvara Temple," is considered one of the greatest built during the Middle Chola period around the 11th century. The temple is of a massive, but unusual, design which seems to emphasize the contemplation of spaces. There are three perimeters within the temple; the first is the outer temple entrance to the east, next is the middle perimeter and finally to the great lingam itself in the center. It evokes an experience unlike any other temple.

Dīkshitar certainly visited the temple, as is clear by the five kritis which are named for this temple. Of these four, one, Brahdisvaram, provides a rather elegant glimpse into Dīkshitar's Tantric play. And though his compositions often detail the theological and ritual aspects of a temple, the physical descriptions of its presiding deities, or the particular architectural form of the temple itself, an elaboration of those aspects are not the subject of this paper.

In order to create an appropriate analysis, we will need to examine a number of important elements that are relevant to the larger understanding of this kriti and its play of creation. First, we will briefly discuss the concept of "sound" in Indian theology.

Sound and Music: Nada Brahman

Architecture is not about building structures, it's about creating space. The "space" is always there, of course, and always has been there. By "creating" it, we are actually "re-creating" it. In the process of delineating it physically, we are materially manifesting everything

Left: Granite Tower of the Brhadīśvara temple. **Above Left**: Perspective View of the portals of the Brhadīśvara Temple. **Above Right:** Monumental gateway of the Brhadīśvara Temple.

that space is not. Yet, in doing so, we can impart the experience of what space may be. In the same way, music can be seen as creating a cognition of the Nada-Brahman (The Divine or Causal Sound, which is the source of all). But Nada-Brahman has always been there. By "creating" music, Dīkshitar is "re-creating" it. He is describing everything that Nada-Brahmin is not by delineating it through a limited, impermanent form. Yet in doing so, his music can impart the experience of Nada-Brahman. According to Lewis Rowell:

Sound, according to the testimony of Sarngadeva, is manifested, not caused, and the sounds that are uttered represent only one category among its many manifestations. By sound the gods exist, with sound they are worshipped, and the entire world process of continuous creation and dissolution subsists by means of sound.[8]

The musician, who is not an object in this cognitive field but is undifferentiated from this field (hence the anonymity aspect of Carnatic music), does not create music, music creates itself through him; through his pulsation. Music is subsumed within the causal vibration of the universe.

Again we refer to Rowell: "Sound…is one, universal, eternal, causal (but not caused), permeating both personal and transpersonal consciousness, and manifested along the human pathway from inner to outer space."[8a]

Dīkshitar's idealization of spatiality in the kriti is evident in the relationships he presents and the approach to the divine in each anga of the kriti:

Pallavi

In the pallavi, Dīkshitar presents space that is untenable, infinite, and ineffable. It is, like the distant temple, undifferentiated. He implores, "O Mind", though he seems to be implying the Mind that is at One with the great lord. For he immediately invokes the superlative: "great lord" (Brhadīśvaram), worshipped by deities (i.e., out of the realm of human experience). The text of the pallavi reads:

Oh Mind! Sing the praises of Lord BRHADĪŚVARA, who is worshipped by BRAHMA, INDRA and others.

The pallavi begins at the beginning of the kriti, but what else? The temple is called "Brhadīśvara" as is the great linga in the center of this vast temple. Dīkshitar approaches the temple in two ways here: it is both the outer temple heading inward and the inner linga, the bhija, heading outward. The reference to Brahma and Indra (and others) is also dual. First, they are among the images in the central portion of the temple which houses the linga, the garbhagrha, and may indicate an outward heading perspective from the point of view of the linga center. But they may also refer to the inner reflection of the great lord from the outer world of bhakti and worship; Brahma, Indra and others reflect the outer world where the gods and men are separated until united inside. The spaces and relationships are pulsing inwardly and outwardly.

A SriVidya tantrika (as we will learn later) would also find mantric (the recitation of sacred words/sounds) significance with these names, though that is best left to initiates in their particular cult. One other possible significance comes from the use of Indra in the passage and may be making a reference to the so-called talu-cakra in Hatha-Yoga, which is the birthplace of Indra and resides in the upper palate, where the tongue is placed to prevent energy from spilling out.[9] This would further enhance the idea that the pallavi relates to the head.

Finally, we might also note that the first word of the Sanskrit lyrics, "brahdisvaram" is broken up musically as "brhadi" and "svaram". The "brhadi" is interesting because it also contains the name of the tala, which is adi (which means "first"). The "svaram" are the notes/sounds of the song. One could speculate that Dīkshitar is beginning the song from the infinite Great Lord (Brahma) by opening up the sonic component of raga (referring to the svaram) and tala.

Anupallavi

In the anupallavi, we are presented with space that creates the immediate need for choice; enforced plurality through intimacy. There is a determined space that identifies nominal relationships. Space determines what you have and do not have/and what you can have or ca not have (hence the reference to Samsara). As E. Valentine Daniel asserts, quoting from Peirce, "Doubt is 'an uneasy and dissatisfied state from which we struggle to free ourselves'".[10] Simply put, the crisis of approach in the anupallavi (where the temple requires negotiation) is the crisis of samsara. The anupallavi text reads:

He, who delights the heart of BRHANNAYAKI, is adept at annihilating the fear of SAMSARA from the hearts of devotees.

The reference to Brhannayaki, the great crowd of captains and queens, may serve to reinforce the position of devotees' relationship to the temple. It is interesting that samsara is rooted in longing for something that you do not have (or having what you do not want) and in the anupallavi we are surrounding the temple, which serves as the "heart". The Nayaka were also members of the 17th century dynasty in Tanjevur who were known as patrons of the arts.[11] Could Dīkshitar be referring to them as well? For a patron of the arts is usually not the one performing the art; that is, they support what they can not have themselves.

Charanam

In the charanam portion of the kriti, Dīkshitar sings of the particular methods for apprehending the divine, but the last two lines narrow the methods into his own personal experience and culminate in the Viddudhi cakra, the cakra of purification and of sound, and finally explode back outward towards a re-conception of the universe as manifestation of the play of Siva. Here we find space as experience (object) and space that is experienced (subject) simultaneously; all things are only one particular thing that converges to this point. The charanam text reads:

He, whose feet are worshipped by distinguished sages, is adroit in bestowing on His devotees both ephemeral and eternal joy simultaneously; worshipped by the King of the serpents; is established in the NAGADHVANI, NADA, BINDU and KALA. He is adored and attended on by GURUGUHA and is the creator of this marvelous, diversified world of names and forms. He, established in the VISUDDHICAKRA, is the primordial cause of the creation of the universe. The whole of creation is a sport for him.

We can contemplate the conspicuous elements of "Nagadhvani, Nada, Bindu, and Kala" by considering their meanings first and then discerning what relationships they denote. Nagadhvani is the name of the raga used in the kriti (dhvani means "poetic expression", and naga may refer to the "ascetic").[13] This raga is rare and is appropriate with this unique (rare) temple. Nada is the source (from the word for "river") and can be seen as the central channel. The Bindu is the seed, the point from where it all begins. Kala, time, may refer to the manifestation of action outward. What the text seems to be saying is that the Great Lord manifests in the raga through the channel leading from the bindu, set forth by the plurality created by Sakti. All of this manifestation comes about through the Vissudhi Cakra, and thus it becomes personalized, for it is Dīkshitar's (as Siva) Vissudhi Cakra which holds this property and then propels it outward through song.

Theological Implications: Tantric Creations

What is our justification for declaring that Dīkshitar's compositions create a form of "sonic architecture", a manifestation of the divine just as the temple architecture is a manifestation of Śiva's desire to experience himself? Our defense rests on the composer's religious path, which above all informs his musical compositions: the path of the SriVidya tantrika.

Although a simple definition of Tantra, and of SriVidya, would be beyond the scope of this paper as well, we can provide some general statements so we have somewhere to start. It should be noted that our elaboration of the theological implications of this particular kriti will provide a much better way to understand the Tantric perspective.

Tantra comes from the verbal root "tan", meaning "to expand." In the words of Georg Feuerstein, the word "Tantra" can be seen as "the expansive, all-emcompassing Reality revealed by wisdom" and denotes the pulsation between the revealed and the concealed as a play by which liberation is experienced.[12] Tantras will argue for a God that is exactly your experiences. After "realization," you still wake up to the world! Epiphanies are forgotten; tantrikas want to be able to keep the awakened insight.

What is SriVidya? It is the living embodiment of Goddess-oriented Tantra otherwise referred to as Sakta Tantra. According to Douglas Brooks, "Sakta Tantrism can be imagined as a river of concepts and practices formed by the convergence of tributaries."[13a] SriVidya is one of those junctures where Kashmir Shaivism and Sakta Tantrism converge. It is a Goddess-worship tradition known to be the cult of Lalitha Tripurasundari ("Beautiful Goddess of the Three Cities", a reference from the Puranas). The "Tripura" has deeper, triadic implications that are evident in Dīkshitar's compositions. There is a focus around the image of the Sri Cakra as representing the expansion and contraction of the cosmos. SriVidya, which places paramount emphasis on mantra, relies on initiation to continue its lineage and its deeper wisdom. It is considered the living link to the Trika traditions, whose practice and philosophy is to be found in Abhninavagupta's "Tantraloka", written around 1000 c.e.

The central realization of this philosophy is the experience of freedom that is the play of Śiva. Śiva

Below: Access to the Brhadīśvara Temple. **Right**: The Linga gallery of the Brhadīśvara Temple.

is completely free from bondage and is all "One." But without bondage, there can be no way for Śiva to experience freedom and so he manifests through Sakti, who is his consort and the energy of all creation and action. Thus, music is Sakti and the keepers of music (i.e., gurus) are mimetic to Śiva; the dynamism of Sakti gives them their re-cognition as Śiva. We will return to this concept later.

The compositions of Dīkshitar, an initiative of SriVidya, cannot help but present themselves as both inviting and yet full of obstacles, for paradox creates power and this is crucial to Tantra. We should be cautious about dismissing his music as unnecessarily complex or overly religious, for the theological principles are not compromised by his Tantric approach, nor is the aesthetic beauty of his music compromised by its apparent subservience to Tantric philosophy. The complexity of his kritis is a complexity of particularities, for the music always appears at the precise point where secrets are revealed only to conceal more secrets. Suspension of paradox is the hallmark of Tantric philosophy. Further, even in the form of the kriti itself and in the purely musical virtuosity of his compositions Dīkshitar manages to re-create the experience of the divine. As Harold S. Powers states in his article on "Musical Art and Esoteric Theism:"

…however much one's attention may focus on the outward meaning of the songs, thence on to their esoteric allusions, soon enough the semantics begins to be overshadowed by the pure sound of the words, which take on independent life as carriers of musical rhythms and shapes…[14]

In a Tantric sense, Dīkshitar is also aware of the mimetic relationship of the kriti and such cosmological ideas as the phases of the moon, the qualities of the planets, and even on the more practical nature of meditative breathing and contemplation. Therefore, it is probable that he would also be aware of the mimetic nature of the kriti and the architectural form of the temple for which he composed.

The temple is a physical embodiment of outward dualism and inward monism; in other words, Brahman (in other words, Śiva). Although the space can be experienced in many ways, the tantrika can experience the temple as himself. To re-experience himself as the temple, he must draw on his inner shakti energy into manifestation. Dīkshitar's kriti places Dīkshitar in the fluctuating role of temple and Śiva, or music and Shakti. This hints at why his mudra-signature is Guruguha, for there is the subversion of roles within himself as tantrika; he is the guru, but he is also the manifestation. But we might also take this further, for if Dīkshitar is Śiva, then he is causing his own recognition. In a sense, then, the temple causes Dīkshitar to create the music to experience itself! The temple was always there (satkaryavada)—its physical manifestation just "happened" as a matter of play by Śiva. The matter only formed around it to re-cognize Śiva's freedom.

We can therefore say that Dīkshitar creates the temple sonically because as we have ascertained, the temple is not the physical matter but the spaces and relationships it creates. Dīkshitar recreates these in his composition. The kriti manifests as an impression of the temple itself, through metaphor, and we have to wonder whether the temple we see and the temple we hear are actually any different from each other.

Finally, it might also be said that experiencing the divine through music requires the participant to become instrumental; that is, to become an instrument. Dīkshitar's "Brhadīśvaram" recreates an experience of the particular divine through metaphor, though the metaphor is the particular construction of his virtuosity as a tantrika and as a performer.

Conclusion: "The whole creation is a sport for Him."

As we reconsider the dimensionality conveyed and present in this one particular kriti, it is important that we remember why and how this kriti exists. It is not simply a musical creation, though it is the metaphoric qualities of music that create signification. It is not simply a prayer by a devotee to the divine, though it is the establishment of relationship with the Divine that provides its theological implications. It is not simply an esoteric instruction by a Tantrika for achieving transcendence, though such depth could not have been achieved unless the composer had been a Tantrika. It is not simply a composition about a temple, though the temple becomes the metaphor to realize that our experience is a pulsation between the inner and outer temple that exists in the micro- and macrocosms. Finally, it is not simply a song about the Great Lord; it is the very way in which the Great Lord experiences freedom. The "Brhadīśvaram" kriti exists for no other reason that as "sport for Him." The models we utilize to appreciate this process, whether architecture, literature, or theology, are the reverberations from the pulsation of Śiva and Sakti and, as such, are no less manifestations of the divine.

Alfred Vitale will graduate in Spring 2004 with a B.A. in Religious Studies. This article is the result of a seminar in Carnatic music and religion, attended at the Eastman School in Spring 2002. His current Honors Research, supported in part by a Barth-Crapsey award in 2003, focuses on the evolution of modern esoteric religious traditions in the west and methodologies for their study. In June 2004, he will be presenting a paper at an academic conference on American Esoteric Religions at Michigan State University. Alfred will be applying to graduate programs in Religious Studies to continue his research.

Suzuki Violin School's Vol. 1 holds the songs used in this study and was the score during certain trials. The song "Andantino" was one of six songs the students sang.

Absolute Memory of Learned Melodies

Children Trained by the Suzuki Violin Method

Victoria Saah, 2003

Advised by Elizabeth Marvin, Ph.D. and Jennifer Williams Brown, Ph.D.

RC Department of Music; Eastman School of Music

The field of music cognition examines many interactions the brain has with music, from interpreting to producing to memorizing music. Within this field, a subset of studies looks at the acquisition and manifestation of an ability called absolute pitch. Absolute pitch (AP) is defined as "the ability to identify the frequency or musical name of a specific tone, or, conversely, the ability to produce some designated frequency, frequency level, or musical pitch without comparing the note with any objective reference tone (ie., without using relative pitch [RP])".[19] In other words, an AP possessor can either name a given pitch (called *pitch labeling*) or produce a named pitch without comparing it to any other pitch (called *pitch memory*).[9] Relative pitch (RP) possessors can also label and produce pitches, but only if they are given a named pitch with which to compare, called a reference tone. However, many adult RP listeners exhibit *pitch memory* without this reference tone, but only in the context of a learned melody. For example, many adults can begin a familiar song on the correct pitch without a reference tone.[9] This study examines *pitch memory* for songs in RP listeners who are young, and who are near the age at which AP is hypothesized to develop.

Current theories explaining the acquisition of AP disagree over whether AP is inherited or learned. Much empirical evidence supports the theory that AP is learned,[19] and some researchers suggest that the learning takes place during a critical period in childhood. As early as 1916, Copp introduced this theory and claimed that 80% of children could learn to produce and recognize middle C on command, but only if they began learning when they were young. Sergeant found a strong negative correlation between the age of beginning music lessons and the incidence of AP later.[14] Miyazaki also suggests an early training period for acquisition of AP as compared to RP, pointing out that the development of AP at a young age may impede the subsequent development of RP. Finally, Takeuchi and Hulse proposed a theory that identifies a period of learning in childhood, during which the acquisition of AP is possible.[15] They found that, compared to non-AP listeners, AP possessors were much slower and less accurate at identifying black-key pitches than white-key pitches. Since the white-key pitches are typically learned first in childhood, and since songs learned in early music lessons are usually in keys with mostly white notes (for simplicity), the researchers hypothesized that the children acquire AP before they have integrated the black keys into their vocabulary. This hypothesis suggests that a window of time exists until approximately age 6 or 7 during which acquiring AP for familiar notes is possible, and after which acquiring AP for newly learned notes is more difficult and less stable.

There is some evidence that even among RP possessors at adulthood, one aspect of AP may be present: the ability to produce a learned melody at a consistent pitch level, without a reference tone. Several studies demonstrate that adult RP possessors are able to access a stable *pitch memory* similar to that of AP listeners when singing familiar songs. For example, Halpern found that when adults with RP were asked to sing familiar songs multiple times in multiple sessions, each adult started on almost the same pitch within each song (with a mean standard deviation of 1.28 semitones).[8] Bergeson and Trehub found similar results in a study of mothers who sang to their infants: the mothers began each song within a mean of 0.82 semitones of where they started the previous time.[4] Clearly, there exists a *pitch memory* in RP possessors that allows each adult to reproduce a pitch within a melodic context consistently over time.

There is also evidence that this pitch memory is stable and accurate between (as opposed to within) adults as well. Levitin asked forty-six undergraduate and graduate psychology students, both with and without musical backgrounds, to hear two favorite popular songs (chosen by each participant from a selection of CDs) in their minds during two trials and then to sing them into a microphone.[9] Levitin hypothesized that these songs would have been heard in the same key every time the listener heard them, and that "repeated exposure to a song creates a memory representation that preserves the actual pitches of the song."[9] This preservation of pitch and key across all exposures to the song is what he called an objective standard. This standard allowed Levitin to identify a *pitch memory* (an "absolute" memory because the correct pitch was often preserved) across the adult population. He found that 40% of these participants sang the correct pitch on at least one of two trials. Furthermore, 81% of the participants began the song within two semitones of the correct pitch on at least one of two trials. These results suggest that the general adult population, which includes primarily RP listeners, can produce pitches from a *pitch memory* that is relatively accurate and stable over time.

It is not known, however, whether this type of memory for pitch exists in children. One recent study showed that infants tended to track absolute pitch patterns (and are disinterested in transposed pitch patterns) in nonmusical contexts,[13] but little focus has been directed toward pitch memory in children. The purpose of our study was to investigate whether a stable *pitch memory* might exist in children who are repeatedly exposed to an objective standard. Since children hear and sing many folk tunes and nursery rhymes in many different keys, the task of locating an objective standard is more difficult than it appears. Thus, to find a group of songs consistently in the same keys, we chose to study the specific population of children learning violin under the Suzuki method. This method of teaching an instrument emphasizes the repeated listening to and practice of beginner-level songs, as well as the singing of these songs to auralize ("hear" in their minds) the melodies. As a result, Suzuki students are repeatedly exposed to songs that are consistently in the same keys. We hypothesized that, if they were asked to reproduce these learned songs, they would be able to access a *pitch memory* (similar to AP possessors and the general adult population) and sing the songs in or close to the correct keys.

In two sessions, participants were asked to sing the same six songs from Suzuki's Violin Book Volume I, which began on varied pitches: C#5, A4, D4, F#4, D5, G4. We recorded the sung performances to measure participants' accuracy against the objective standard. The design of this study was influenced strongly by that of previous researchers. Like Levitin, we asked the students first to hear the songs in their heads before singing them.[9] Finally, we used distractor tones[3,10,15] between each song and in the AP post-test so that no RP could be used between trials.

Within this population of young Suzuki students, we also looked for an effect of age on the accuracy of pitch production. To test the early-learning theory of AP acquisition, we solicited two groups of children who began learning violin using the Suzuki method at different ages: one began at 3-5 years, the other 6-8 years. At the time we solicited participants, we did not know whether or not the children had AP. We hypothesized that if there were a period in early childhood during which children acquired AP, the group who began at a younger age might contain some AP listeners and would perform better. As for the question of whether any of these children were developmentally ready to sing in a stable key, Davidson, McKernon, and Gardner showed that by the age of 5 years children could begin a song and maintain its key all the way through.[6] To control for differences in vocal development between the age groups, we gave each participant a pitch-matching test to ensure that both young and old students could physically hear, internalize, and repeat a given pitch pattern accurately.

Finally, in addition to our *pitch memory* hypothesis and testing for the effect of age, we examined the effect of viewing the musical score on their performance. We hypothesized that if the children saw the score before singing the song, that the visual stimulus might trigger the hearing of the starting pitch in their minds and thus improve their accuracy. Previous studies involving the score's effect on auralization (this inner hearing of the pitch) in adults were done by Terhardt and Seewann in 1983, showing that non-AP participants could discriminate between correct and transposed excerpts when the score was visible (though the study did not compare performances without the score). Ward, on the other hand, had two conditions: one in which participants viewed the score while listening and one where they listened with no score visible.[18] He investigated whether having the notation visible (i.e., having the opportunity to auralize the pitches from the score) helped participants recognize these familiar pieces when they were transposed. His results revealed that listeners could identify small transpositions only with the score in front of them and not without it. Larger transpositions could be identified with or without the score. Thus, the score is not required for auralization, but it helps. In this same vein, the children learning the Suzuki songs may not need the score in order to begin on the correct pitch, but auralizing the starting pitch from the score might improve their accuracy.

In summary, we tested one primary prediction and two secondary predictions: our primary prediction was that the majority of these children will begin learned songs within a semitone of the correct pitch. Secondarily, the children who began musical training at a younger age will achieve higher accuracy. Finally, also secondary to the first hypothesis, we predicted that on the trials where the score is visible, the children will achieve higher accuracy.

The participants were 10 students (4 male, 6 female) at the Kanack School in Rochester, NY. They ranged in age from 7 to 10 years (mean, 8.30; SD, 1.25) and were all right-handed. All of the students studied violin under the Suzuki method and had been playing for approximately 3 years, within a range of 2.5 to 3.5 years. Since they had played for the same amount of time, they all learned the songs from Volume 1 and had played them consistently since their first year. These students had one of two teachers, both of whom taught a standardized version of the Suzuki method. This method requires the children to sing the songs they learn on the violin and continue to sing and play those songs for the duration of those lessons.

During the study, nine out of ten participants returned surveys filled out by their parents. These data showed that two (~20%) of the participants study another instrument besides the violin, and five (~50%) have immediate family members who play at least one instrument. About half of the students showed an interest in music before they started playing the violin, some around ages 1-2 years and even younger. Seven of the students (~80%) were sung to by their parents when they were younger; interestingly, the two without the experience of parental singing were the two who performed the least well on the pitch-matching test.

The 10 students were split into 2 age groups of 5 students

	Session 1	Session 2
	Song #	Song #
Without Notation	1	4
	2	5
	3	6
With Notation	4	1
	5	2
	6	3
	Pitch-Matching post-test	AP post-test

Table 1: Experimental Design

each: the younger was from 6 to 8 years (mean, 7.2; SD, 0.45) and the older was from 9 to 11 years (mean, 9.4; SD, 0.55). Each student participated in two separate sessions, spaced apart by one week.

In Session 1, the participant heard the first set of distractor tones to prevent the use of RP to find the starting pitch of the first song. The student was then asked to sing song #1, and only the first few notes were recorded. The investigator stopped the student after these few notes and repeated this process twice for two additional songs (songs 1, 2, 3). Then this process was repeated for three different songs (songs 4, 5, 6). During songs 4 through 6, students viewed the score for each song individually.

At the end of Session 1, the student was given the pitch-matching test; it was given at the end so that no tones from this test interfered with participants' recall ability or choice of starting pitch when singing the songs. The participant sang back each pair of notes after hearing them, and the investigator digitally recorded the singing. The purpose of this post-test was to control for vocal development and ensure that the students were able to sing the notes they auralized in their minds.

Session 2 began with the same process as the first session, but with the last three songs first (songs 4, 5, 6), this time without the score. Then the student was asked to sing the first three songs from Session 1 (songs 1, 2, 3) with the scores visible, and those responses were recorded.

At the end of Session 2, the AP post-test was performed. As before, it was given at the end of the session so that no bias was created by hearing other experimental tones. Before each pitch sounded, one of the three sets of distractor tones was played. After hearing the distractor tone sequence and then the pitch, each participant named the pitch aloud and the investigator wrote down each answer (in order to eliminate any motor skill interference by the children and also to avoid any answering bias caused by having participants see previous responses). The purpose of this post-test was to identify the cognitive mechanism used to find the starting pitches. Since we were examining pitch memory in RP listeners, this test identified any AP listeners who presumably would have used a different cognitive process.

The distractor tones, the pitch-matching post-test and the AP post-test were composed once and used for all participants. The distractor tones were made up of the 12 equally tempered pitches between A3 and G#4 in a sequence that lasted a total of 1 second. We randomized their order using a random-number table and their timbre was that of a synthesized violin in the musical notation program Finale. No pitch was repeated and there were no rests in the sequence. We used quarter, eighth, and sixteenth note values, the order of which was also determined using a random-number table. A total of three distinct distractor tone sequences were derived with the same rhythmic values and pitches, but randomized differently so that the students did not grow accustomed to the final pitch of any one sequence.

The pitch-matching post-test was made of 12 pairs of equally tempered pitches, whose starting notes were drawn from D4, E4, F4, G4 or A4; this is a comfortable singing range for children and encompasses most of the starting notes of their violin pieces. These starting pitches were ordered using a random-number table. Participants heard a two-note pattern (an interval). Each note lasted 1 second and had a synthesized violin timbre. The 6 intervals used were minor seconds, major seconds, minor thirds, major thirds, perfect fourths and perfect fifths. Each was heard twice, once ascending and once descending. The order of these intervals was also determined using a random-number table.

Finally, the AP post-test was made of the 18 equally tempered pitches between G3 and C5; their order was determined by a random-number table, and their timbre was a synthesized violin timbre in Finale. Each pitch was heard only once and lasted 1 second in duration. One of the three distractor tone sequences was played before each pitch sounded.

After both sessions of the experiment were complete for each subject, we scored their responses. In scoring the participants' singing, only the starting pitch was evaluated. Octave errors were considered to be correct, consistent with Levitin's analysis, since singers often change the octave to accommodate their vocal tessitura.[9] Since octave errors were permitted, only pitch class (chroma) was examined (not pitch height), so that the furthest a response could be from the correct pitch was 6 semitones. Errors by one semitone were also considered correct in both the violin songs and the pitch-matching test, since even AP possessors sometimes produce discrepant responses by one semitone.[17] This method of scoring also provided some leniency for premature vocal control. The actual spectral analysis provided the frequency (in Hz) and decibel (dB) level of each frequency within the sung tone. Only the steady state portion of the tone was analyzed. From the dB level, we discerned the intended sung pitch from its overtone series because the fundamental frequency was the loudest and most prevalent.

The analysis of the pitch-matching post-test was identical to that of the experiment proper: only the first note was analyzed. A score of 10/12 (83%) had to be obtained to pass the pitch-matching post-test. Semitone errors were considered correct for the same reasons as they were in the singing of the violin songs.

Scoring the AP post-test was simpler, as it only involved comparing the pitches named by the participants (and written down by the investigator) with the actual pitches of the notes heard. Participants were determined to have AP if they scored 16/18 (approx. 90%) or higher on the test.[12] No semitone errors were permitted on the AP post-test.[10, 12] Two participants were honest enough to state that when they heard each given pitch, they were not familiar enough with the letter names of the notes to assign a letter to that pitch. Instead, they offered to name the fingering on the violin of where that note would be. This unfamiliarity was probably the case with some of the other students as well, who did not protest the naming task; many of them were hesitant with the letter names, and some only used the letters E, A, D, B and G as responses, instead of all twelve possible note names.

In the two cases where the participants named fingering instead, there was an ambiguity in the scoring: when one says "finger 2 on the D string," that note can either be F or F#, depending on where one places finger 2 on the string. Only once did this become an issue (all other times, neither note was correct), and in this case the experimenter assumed that the participant meant the correct note and not the incorrect one. Even with this assumption, the participant scored well below 90%, so ultimately this ambiguity proved to have no effect on our results: none of the participants had absolute pitch, according to our post-test. Their mean percent correct was only 10.00%, with a standard deviation of 6.56%. Their range of accuracy was 0% to 22% of the pitches correctly identified.

The results of the pitch-matching post-test showed that most of the children could hear and correctly sing back a two-note pitch pattern. Nine of the ten participants passed the pitch-matching post-test with a mean percentage correct of 98.11% and a standard deviation of 5.67%. The tenth subject sang only 25% of the pitches back correctly. His data was therefore omitted from all data analyses. Since most of the students passed the test, we confirmed that they were not hearing correct pitches and singing incorrect ones; they were actually singing the notes they auralized in their minds.

The results of the experiment overall were not what we anticipated. Contrary to our primary hypothesis, the children did not perform significantly better than chance in singing learned songs in the correct key. Chance was 0.25 (25%) because semitone errors were permitted; three possible correct answers of 12 possibilities yield a 3/12 or 0.25 probability of a correct answer by chance. The participants' mean percentage correct was 31.44%, with SD of 17.12%. The results of a two-tailed, one-sample t-test of their performance against 0.25 revealed that they did not perform significantly better than chance would have predicted [$t(9) = 1.129$, $p<0.292$]. In fact, the participants sang the correct pitch almost as often as all of the other pitch classes, so the correct pitch or the semitones either side of it were not favored.

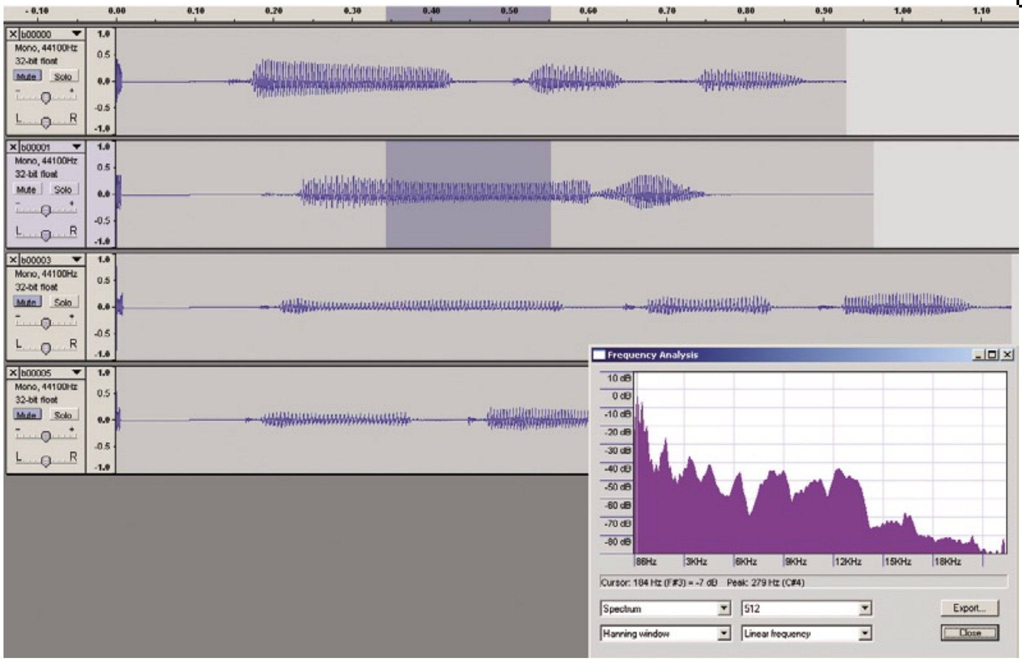

Audacity graph of the overtone series of each pitch the student sang.

The hypothesized age effect yielded similar findings: the younger group of students did not perform significantly better than the older group. Levene's test for equal variances found that the variances of the two age groups did not significantly differ [F=0.329, p<0.584]. Five participants were in the younger group (mean % accuracy=0.33, SD=0.157) and four participants were in the older group (mean % accuracy=0.29, SD=0.210), since one participant's data was omitted due to the pitch-matching problem. A two-tailed, independent samples t-test found that the younger group did not identify the pitches more accurately than the older group [$t(9)$=0.337, p<0.746].

Similarly, no effect of having the musical notation visible was found. Levene's test for equal variances again found no difference in the variances between the participants' performance with notation and their performance without notation [F=0.471, p<0.494]. Comparing the means of the two groups (M=2.85 semitones away without the score and M=2.74 semitones with the score) revealed that having the score in front of them had no effect on what note they chose to sing [$t(108)$=0.297, p(2-tailed)<0.767].

Following the between-subjects analysis, we examined consistencies in pitch production within subjects, from the first session to the second.[4, 8] We found that participants were remarkably consistent in choosing a particular starting pitch for each song, regardless of whether that pitch was near or far from the objective standard. In other words, even if the participant chose to start a song in Session 1 on a pitch that was 6 semitones away from the correct pitch, the participant usually sang that song in Session 2 near that incorrect pitch, still about 6 semitones away. On average on the second day, the participants never chose a pitch more than an average of 1.56 semitones away from the pitch they chose on the first day.

To further test the relationship between the two sessions, we averaged each participant's performance in each session, resulting in two scores per participant (one per session). A two-tailed Spearman correlation between these averages showed significant covariance in pitch production from

A group of children perform "Andantino" at an outdoor concert from memory.

the first to the second session [r=.803, p<.01]. In other words, most of the variance between sessions could be explained by which participant was singing, as each participant was quite consistent.

Thus the pitch representations that these children had in their minds were, whether correct or incorrect, relatively stable over time. The very high correlation between sessions within each student's performance is consistent with the results obtained by Halpern and Bergeson and Trehub, both of whom examined repeated pitch production within participants.[8,4] Halpern attributed her participants' consistent performance to an aspect of absolute pitch in adults. This conclusion cannot apply to the current study because the children's responses varied so widely from the objective standard. They sang both correct and various incorrect pitches consistently, so their pitch representations were stable but did not preserve the actual key of each piece.

On the other hand, Bergeson and Trehub attributed the consistent performances to a few factors that could apply to the current study. The first two are memory enhancers termed "positive affect" and "congruent mood."[4] These two factors set up a similar environment at the time of each pitch production, which aids memory and therefore decreases variability. This factor can be extended to the rest of the environment as well: the sessions were held in the Suzuki school where the students learned, sang and played the violin songs used in this study. This reproduction of environment may have contributed to their stable performance as well. The other factor is motor memory, which is not as directly applicable as the environmental factors because these students sang the Suzuki songs many times in the correct keys while learning and practicing them. If motor memory were involved, then they would continue to sing each song in its correct key during the experiment, by judging where to start the piece based on their vocal range. Motor memory could be involved, however, if these students sang or hummed the song to themselves outside of their lessons, and never bothered to match their starting pitch to the actual song. If this were the case, the student would pick a vocally comfortable starting pitch (an arbitrary one compared to the actual key of the song) and would repeatedly sing the song starting on that pitch, thus strengthening an inaccurate pitch memory.

This inaccurate pitch memory is contrary to our primary hypothesis that the children would sing the songs in the correct keys. It is, however, interesting in light of Levitin's study.[9] Levitin concluded that the general adult population has a somewhat stable absolute memory for pitches in the context of learned melodies. One would intuitively think that children with musical training would exhibit a similar, if not more secure, absolute memory for pitch. Since these children passed the pitch-matching test, we know they could produce the pitches they auralized, just as the adults could. Since they could sing the pitches they internalized, it is clear that they must not have had the correct pitches of the objective standard internalized (i.e., they did not exhibit an accurate memory for pitch like the adults did).

These children may not have developed an absolute memory by this point in their childhood, but it is likely that at least some of these children will acquire one by young adulthood like Levitin's participants. It is therefore possible that the explanation is one of cognitive development; perhaps the ability to form stable pitch representations for melodies at their objective standard develops during puberty. Another possibility is that these children have pitch memories, but not for these songs. Perhaps there exists a threshold number of times that one (adult or child) must listen to or sing a song within a certain time period before the correct pitch level is retained in and retrievable from the mind. The adults in Levitin's study sang a favorite song by a pop artist and most likely sang this song almost daily, while these Suzuki students sang these violin songs while learning them, but ultimately more often played the songs on the violin.

The result that the younger children did not perform significantly better than the older children carries implications for research on AP acquisition. Most research done to support the AP early-learning hypothesis has examined adult participants with AP and compared their early musical training to that of participants without AP. In contrast, this study examined children and did not find that the children who started earlier showed any more instance of AP. Since this correlation between AP and early training was not supported by this study, one of three reasons may be responsible. First, perhaps the age around which the early-learning window begins to close is not 6 years as this study's design specifies, but rather is age 7 or older. Second, perhaps there is a threshold age at which AP is expressed. The children who began music training early may indeed develop AP in the future (while the ones who did not may not develop AP), but they will not exhibit signs of AP or pass AP tests until a certain age or ability level. Third, the early-learning hypothesis may be relevant only for those children who do eventually exhibit AP. In other words, early training (vs. later training) may have no effect on children who do not have a genetic predisposition toward AP acquisition.[1, 2, 7]

Finally, the reason we found no effect of reading musical notation may lie in the particular pedagogy of the Suzuki method. One subject, when the score was placed before her, replied calmly that she did not need the score since she could not read music; she stated that she only looks at the music for dynamic markings. Critics of the Suzuki method have long stressed this point: since Suzuki students learn through repetition and imitation, a few of them manage to learn to play an instrument using their ear and not ever reading the notes. Most of the other participants glanced cursorily at the score before turning away to think of how the song's melody sounded. Only one participant studied the first note on the score to try and remember it. While one cannot deduce the effect notation had on the subjects by their physical reactions alone, it is clear by their equal performance with and without notation and by their nonchalant attitudes that having the score in front of them did not help them to remember the starting pitches because the participants did not use the score or retrieve any visual cues from it.

The results of the pitch-matching post-test and the AP post-test also deserve some attention. On the pitch-matching post-test, participants either sang almost all of the starting pitches correctly or they performed at chance level. Most of these children demonstrated the ability to internalize and sing back pitches. We therefore know that the children could sing each violin song in its correct key if given the starting pitch (i.e., if given a reference tone). However, they have not acquired the absolute memory for pitch that Levitin has shown many adults possess. Since the children performed at chance level in singing the violin songs in the correct key, perhaps the *pitch memory* skill is similarly dichotomous; like the ability to match pitch, they either have the ability to remember pitches or they do not (at least at this age).

The AP post-test results also hovered around chance level, which confirmed that no participants used AP to find a starting pitch. However, the fact that some of the students did not know the note names makes these results more difficult to interpret. By our definition of AP, if a student cannot match a note with its appropriate letter name, then the student does not have AP. But perhaps the student could have an idea of what that note is and just not have the *pitch labeling* capability because the notes have not yet been paired with their names in the students' minds. It is difficult to determine the extent to which children have cognitive representations of pitches in their minds because many of them have not yet acquired the tools with which to express and identify them. Since the conventional definition of AP specifies both *pitch labeling* and *pitch memory* abilities, then these children cannot be said to have AP at this time if the cognitive link between a pitch's sound and its name has not formed. As mentioned previously, however, these children may not yet have reached the developmental threshold at which students with early musical training begin to show signs of having AP.

It became clear from interpreting these results that the small sample size made a significant impact on the findings. Future research should sample a larger number of participants so that a single particularly good or particularly poor performance does not strongly affect the outcome. For example, one of the younger participants began zero out of twelve songs within a semitone of the correct pitch. This performance brought down the entire accuracy average for the younger group, and because the sample size is so small (n_{young}=5), the average changed dramatically due to one participant's performance. A larger sample size may allow significant effects of age or notation to surface and not be so drastically altered by individual performances.

It is difficult, however, to increase the sample size of children and still maintain an objective standard. Children are exposed to many melodies in many different keys; finding one or more pieces heard or sung consistently in the same key is difficult, and it is extremely difficult to find more than a few children exposed to that objective standard enough that they can reproduce it. Nonetheless, future studies attempting to find an effect of age on pitch production and absolute memory should sample more participants. Perhaps a design like Levitin's would allow the children to choose a favorite popular song from a rack of CD's, which would ensure that they knew the song quite well.[9]

Future studies should also examine the assumptions of the pitch-matching test. It is altogether possible that in short-term memory, pitch is internalized and reproduced by different cognitive processes than in long-term memory (for example, for a familiar song). Therefore, it could be erroneous to assume that if participants can create an accurate short-term representation of the pitch, they can also create an accurate long-term one. In fact, this point may provide an alternate explanation for the current findings: perhaps these children have not developed the cognitive ability to create long-term pitch memories or, if they have the ability, have not developed a way to access these memories.

Overall, these results raise new questions about the cognitive development of AP and RP in children. Is there perhaps a chronological or developmental threshold past which these skills are exhibited? If so, will the aspects of absolute memory for learned melodies that exist in the adult population begin to exist in the younger population after a certain developmental threshold is reached? Also, will the stable *pitch memory* within each participant become more accurate (i.e., closer to the correct pitch) if the participant sings the song in the correct key more often? These findings encourage more rigorous examinations of aspects of absolute pitch in children and more empirical tests of the early-learning theory. Only by learning more about how we develop cognitively through childhood can we begin to foster that growth.

Victoria Saah received a B.A. in Music from the University of Rochester in May, 2003. This article is an abridged version of her senior honors thesis. Victoria is currently working in Erie, Pennsylvania, while applying to medical school. She plans to begin medical school in the fall of 2004.

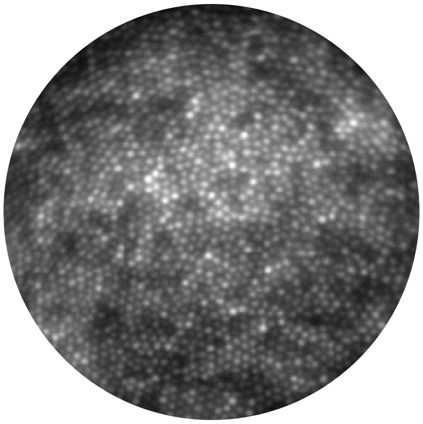

The Fixational Stability of the Human Eye Measured by Imaging the Cone Mosaic

Nicole M. Putnam, Heidi J. Hofer, Nathan Doble, Joseph Carroll, Li Chen

Advised by David R. Williams, Ph.D.

Institute of Optics

The human eye has many imperfections in the cornea and lens, referred to as aberrations, which distort the light entering the eye, resulting in a distorted image on the retina which is perceived by the person. Conventional glasses and contact lens only correct some of these imperfections. Some glasses and contact lenses correct astigmatism, but higher order aberrations in the eye are not measured or corrected. Adaptive optics techniques allow these other imperfections in the eye to be measured by a wavefront sensor. In turn, most aberrations may be corrected by a deformable mirror. This means that the subject will perceive a much clearer view of the world and also the retina may be imaged more clearly with adaptive optics than ever before, with individual cone photoreceptors being clearly resolved. Figure 3 demonstrates the basic idea behind adaptive optics.

Figures 1 and 2 show the difference between a retinal image without adaptive optics and then with adaptive optics for one subject, JW, one degree from the fovea. It is clear that the use of adaptive optics allows individual cone photoreceptors to be seen more clearly.

By correcting the subject's higher order aberrations, they are presented with a more clear view of the world. In the future, these same methods will be used to make customized contact lenses and customized refractive surgery available to consumers, making these improvements in vision available to everyone.

With the use of adaptive optics retinal imaging, we have developed an extremely accurate technique that directly monitors the movement of the cone mosiac in order to measure the fixational stability of the human eye. The relative position of the retina was determined with an accuracy that was at least 6 times smaller than the diameter of the smallest foveal cone photoreceptors. Fixation stability was measured in six subjects and the standard deviation of successive fixations was found to range from 1.5 to 5.0 minutes of arc in the horizontal and vertical directions. These results are in agreement with previous studies that have reported standard deviations ranging from 1.4 to 5 minutes of arc.[2,10] However, there are factors that set an upper limit on the accuracy with which fixation direction can be accurately determined.

One such factor is that under normal conditions, the image formed on the retina is never stationary. The eye is continually moving, but never shifts far from its mean position during maintained fixation. The desire to track and understand eye movements has led to an extensive history of experiments spanning over a century of research.[1,2,5] However, most previous methods of tracking fixation have done so indirectly by monitoring the front surfaces of the eye, and consequently have had limited accuracy. Due to the precision required by their fields, many psychophysical and medical applications would benefit from a more accurate method of measuring fixation. For example, if the position of a retinal stimulus could be determined with an uncertainty smaller than the size of an individual photoreceptor, it would be possible to determine the individual contributions that each retinal receptor makes to visual perception. A highly accurate method of tracking the retina could also be used to increase precision in laser retinal surgeries (such as photocoagulation surgery used to treat diabetic retinopathy), which would result in less damage to healthy retinal tissue and better outcomes for the patient. Here we present a method utilizing adaptive optics and retinal imaging that is able to directly track retinal movement with accuracy better than a sixth of the size of the smallest foveal cones. This method required us to determine the change in the relative position of a stimulus on the retina during successive fixations. We accomplished this task by monitoring the positions of individual cone photoreceptors in retinal images acquired simultaneously with a stimulus presentation. This technique, however, is contingent on acquiring clear images of cones in the retina. Consequently, an adaptive optics (AO) system is used to remove the blur in retinal images caused by the eye's spatially and temporally varying aberrations.

Figure 4 shows a schematic diagram of the Rochester AO ophthalmoscope.[4,11] The subjects' heads were stabilized with a dental impression. The subjects were then told to fixate on a target projected through the AO system with a Texas

Instruments digital micromirror device (DMD Display). Their optical aberration was then measured with a Shack-Hartmann wavefront sensor while the conjugate wavefront was sent to a 97 channel Xinetics deformable mirror (DM). Throughout this process, the wave aberration is constantly measured while the DM is updated.

Once a desired rms. wavefront aberration has been achieved (~0.1μm over a 6.8mm pupil), the flashlamp is triggered. A 4msec pulse at 550nm (0.6μJ) is reflected off the retina and follows the corrected path off of the DM. The process ends with the image being taken with the retinal imaging camera (Princeton Instruments). The simultaneous imaging and stimulus presentation was enabled through inserting a high quality pellicle beamsplitter into the science arm before the retinal imaging camera and DMD display. It was critical for this experiment that the DMD display and the retinal imaging camera were exactly centered and conjugate with respect to each other. This was achieved by firstly placing a plane mirror in the pupil plane of the science arm, the plane mirror also being optically conjugate to the retinal imaging camera. The axial position of the DMD display was then adjusted and the plane mirror tilted. The axial position of the DMD display was re-adjusted until no movement of the retro-reflected image was observed on the retinal-imaging camera. Only once this was achieved could the retinal imaging camera and the DMD display be deemed conjugate. The two devices were then centered by adjusting the lateral position of the retinal-imaging camera until the center of the fixation target fell on the center pixel of the camera. As a result, the positions of the fixation target and the retinal imaging camera are always locked with respect to the subject's retina, even if the subject's pupil moved relative to the AO system.

The target was a Maltese cross with a diameter of one degree displayed at 550 nm (luminance 8 cd/m^2) - the same wavelength used for retinal imaging. The alignment steps previously described ensured that both the target and CCD camera were at the same vergence, and that both used the same 6 mm artificial pupil. Target size, color, and luminance have no statistically reliable effects on drift magnitude, and as a result, were not considered a factor in this part of the experiment.[8] Mean fixation position and saccade frequency also had no effect in the current study because images were taken before a saccade occurred.

To allow fixation of the retina at 1.25 degrees, the subjects were told to fixate at the center of the Maltese cross that was placed at 1.25 degrees along the horizontal axis. To avoid difficulties in fixation due to the after-image of the retinal imaging flash, the subjects were also instructed to wait until the after- image had disappeared. Since points on the retina are fixed with respect to one another, the movement of the fovea could be reconstructed from the movement of photoreceptors in these retinal images at 1.25 degrees. Changes in fixation position from trial to trial were determined from the relative displacement of the cone photoreceptors in retinal images. From these relative displacements, the standard deviation of successive fixations was calculated. However, the absolute fixation position and any systematic errors could not be measured.

The cone mosaics of six subjects who were attempting

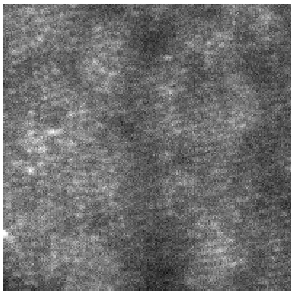

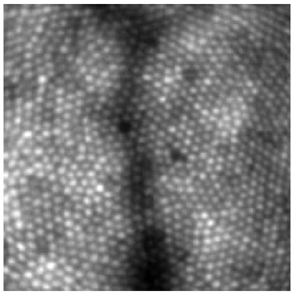

Far Left, Figure 1: A sum of retinal images without adaptive optics. **Left, Figure 2**: A sum of retinal images with adaptive optics. **Below, Figure 3**: Schematic of adaptive optics.

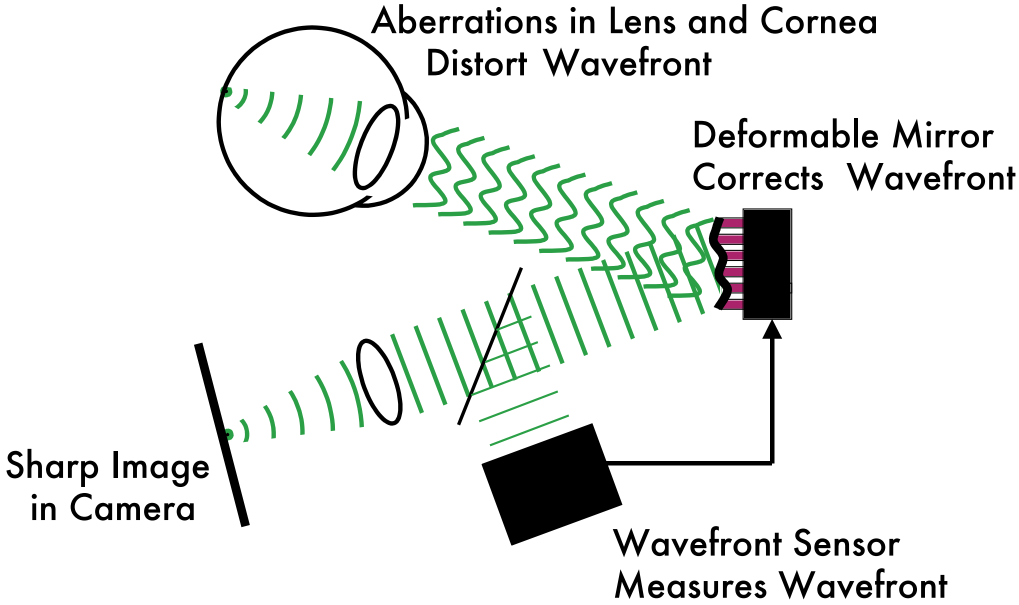

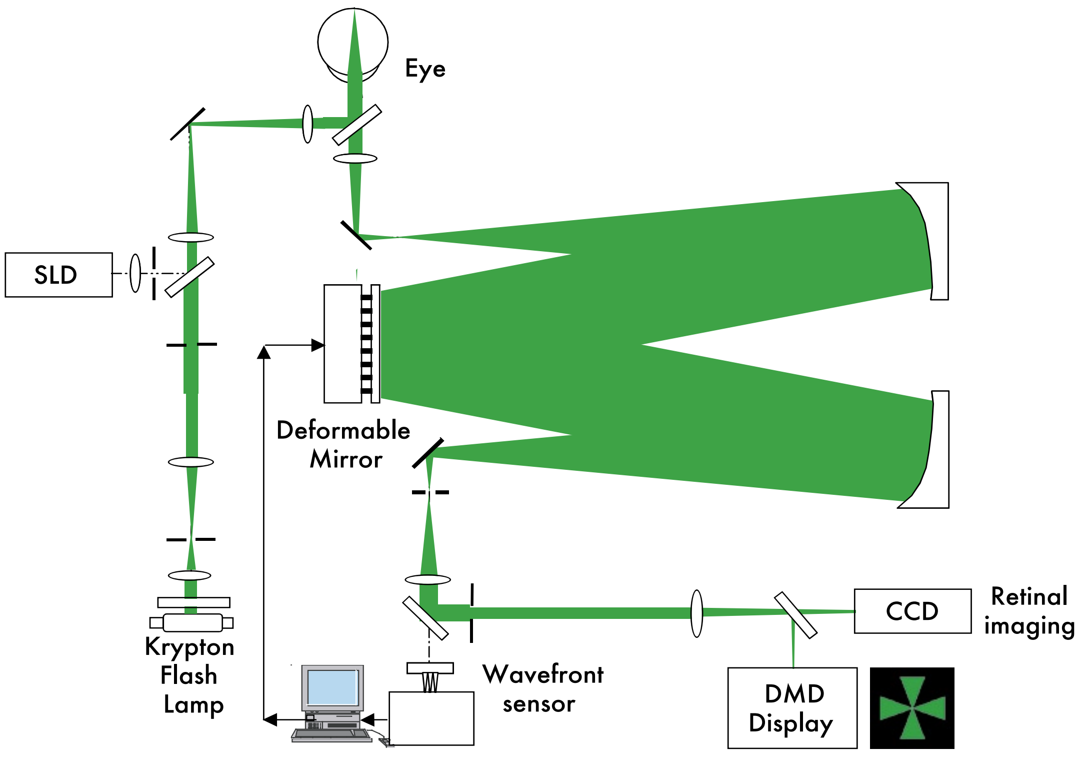

Figure 4: Schematic diagram of the Rochester adaptive optics ophthalmoscope.

to accurately fixate on the center of a fixation target were imaged with the Rochester adaptive optics (AO) ophthalmoscope. The subjects ranged in age from 21-35 years and had normal vision. All subjects were aware of the purpose of the research and were given time to practice. Two of the subjects also had considerable previous experience being imaged in the Rochester AO system. The subject initiated the AO correction procedure, which took 0.25 to 0.5 sec, and was immediately followed by image acquisition.

A second experiment was also conducted in which the subject initiated the adaptive optics correction procedure and subsequently initiated the image acquisition when he felt that he was fixating accurately, all other constraints remaining unchanged.

Changes in fixation position from trial to trial were determined from the relative displacement of the cone photoreceptors in retinal images. From these relative displacements, the standard deviation of successive fixations was calculated. However, the absolute fixation position and any systematic errors could not be measured.

To determine the movement of cones in retinal images, images were interpolated by a factor of 4 and then cross-correlated. The location of the peak of the cross-correlation function between any two images gives the retinal movement (and thus the change in fixation) between those two trials. Retinal image shifts were determined by taking the peak pixel of the correlation function. This limits the precision with which we can determine the retinal image shifts to one-fourth of an image pixel. This means that the accuracy with which the retinal image shifts could be determined was better than 0.03 minutes of arc, or 1/16th the size of a single foveal cone photoreceptor. The registration process could be made even more accurate if a better sub-pixel fit of the correlation function was used, such as a Gaussian curve fit, which returns sub-pixel estimates of the correlation peak. Determining the shift from the cross-correlation in this manner could reduce the uncertainty due to registration error to as little as 0.0008 minutes of arc, or 1/1000 the size of a single cone. However, other sources of error had a greater influence on measurement accuracy, so there was little motivation to use a more precise method of determining the shifts of retinal images.

The subject was fixating on a target that was placed at 1.25 degrees along the horizontal axis, and as a result, any eye rotation introduced an artificial movement in the registration process that was larger in the vertical direction than in the horizontal direction.

A process analogous to cross-correlation was performed on the registered retinal images to determine the angle of rotation between images. On average, subjects experienced a standard deviation of torsional rotation of 22.1 minutes of arc. The precision with which torsional eye movements could be determined was 6 minutes of arc. The error estimate due to rotation in the vertical direction with the current setup is 0.13 minutes of arc or 1/4 the size of a singe foveal cone photoreceptor. The results reported account for the horizontal and vertical angular error. Imaging in the fovea would eliminate the error due to rotation in future experiments.

The value of the retinal image shifts was then converted to arc minutes and the standard deviation was calculated. The overall fixational stability was characterized by taking the standard deviation of retinal image position across at least 67 trials per subject.

The standard deviation ranged from 1.5 to 5.0 minutes of arc in the horizontal and vertical directions with an average of 2.84 minutes of arc in the horizontal direction and 2.14 minutes of arc in the vertical direction. Table 1 summarizes these results.

One previous study indicated that measured values of the standard deviation ranges from 1.4 to 3.2 minutes of arc for subjects fixating at a distant target.[2] Another indicated that the standard deviations ranges from approximately 2 to 5 minutes of arc during maintained fixation.[10] The results from the current study support these previous results.

For experiments in which the subject initiated the adaptive optics correction and subsequently initiated the imaging flash when he felt he was fixating most accurately, the results

Subject	Horizontal Std (arcmin)	Vertical Std (arcmin)	Std of Torsion (arcmin)
JP	1.52	2.39	18.4
ND	1.74	2.11	19.4
NP	1.98	5.02	21.4
BP	2.52	2.16	21.8
GYY	2.69	2.65	28.5
AP	2.21	2.71	23.2
Average	2.11	2.84	22.1

Subject	Horizontal Std (arcmin)	Vertical Std (arcmin)	Std of Rotation (arcmin)
JP	1.19	2.13	19.8
NP	2.2	4.08	20.0
ND	2.38	2.37	22.4
Average	1.92	2.86	20.7
Average from Exp1	2.11	2.84	22.1

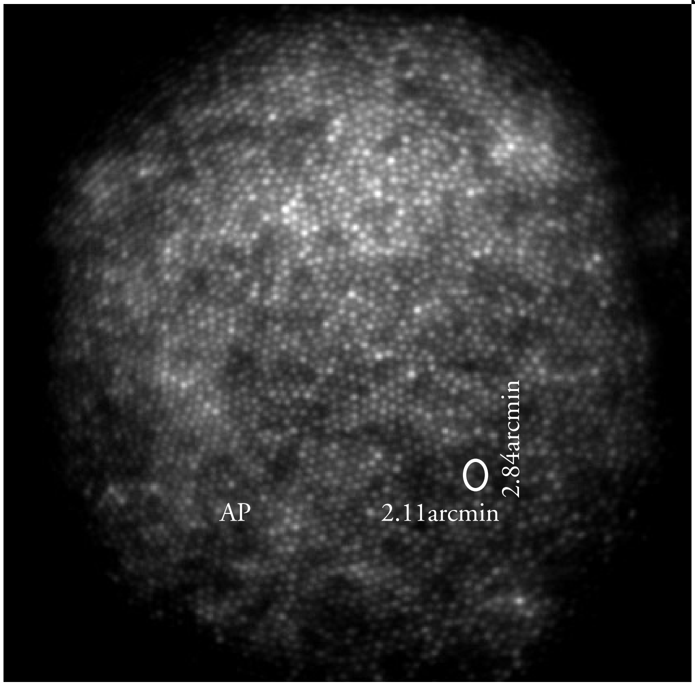

Top Left, Table 1: Self-initiated image acquistion. **Bottom Left, Table 2:** Results from the second experiment where the adaptive optics correction and the image acquisition were both self-presented by the subject. **Above, Figure 5**: The cone mosaic at 1.25° eccentricity for one subject with the average standard deviation of fixation shown.

were similar. Over three subjects with at least 70 trials per subject, the standard deviation of fixation ranged from 1.2 to 4.0 minutes of arc with a mean across subjects of 3.18 minutes of arc in the horizontal direction and 1.78 minutes of arc in the vertical direction. The average standard deviation of torsional rotation was 19.7 minutes of arc. Overall, there was no significant change from the previous procedure, indicating that in the first experiment subjects were ready for the flash. Table 2 summarizes the results from the second experiment.

Focus error in the system setup also introduced uncertainty into the measurement of the fixational stability. This error corresponds to an uncertainty in the final measurement of 0.07 minutes of arc, or 1/7 the size of a single cone photoreceptor.

Considering the focus error along with the error introduced from the image processing from computing the image shift, the total error in the measurement of the fixational stability is at most 0.08 minutes of arc, or 1/6 the size of a single foveal cone photoreceptor. This is the error using our specific setup, which would be increased if any of the precautions to decrease relative movement of the imaging camera and fixation target were not made. In addition, all measurements were made using monochromatic light, and any differences between target wavelength and imaging wavelength would introduce an additional uncertainty in the measurement of the fixational stability as a result of chromatic aberration in the eye.

Due to the optical arrangement and the use of adaptive optics, the method developed for this study is much more accurate than all previous methods for measuring the fixational stability of the human eye. The relative position of the retina was determined with an accuracy that was at least 6 times smaller than a single foveal cone photoreceptor. These measurements also determine the maximum accuracy with which a stimulus can be delivered to a specific location on a subject's retina without tracking the eye.

Nicole Putnam graduated from the University of Rochester in 2003 with a B.S. in Optics. This article is a summary of work done in her undergraduate independent study in David Williams' Lab at the Center for Visual Science. Nicole is currently at the Unversity of Arizona studying Optical Science.

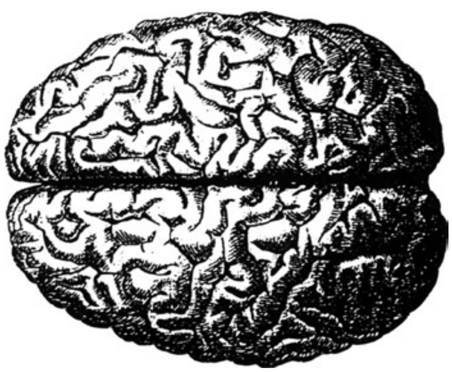

Short-Term Memory In Second Language Learners of ASL

The Competing Effects of Modality and Mother Tongue

Matt Hall, 2003

Advised by Elissa Newport, Ph.D.

Department of Brain & Cognitive Sciences

Mechanisms behind auditory linguistic short-term memory (STM) have been the focus of much research over the past several decades. Studies of memory span began with Miller's famous finding that for English speakers, STM has a limited capacity of about seven items.[1] That hypothesis has since been modified several times as evidence has been uncovered suggesting that memory span is limited not by the number of items to be recalled, but by the time needed to mentally rehearse these items. That is, words that take longer to mentally rehearse result in a decreased span.[2] Cross-linguistically, languages in which names for digits are shorter show longer digit spans, and opposite is true for languages with longer digit names.[3,4] When Miller's hypothesis is recast in terms of time, auditory linguistic STM is observed to hold roughly 2 seconds worth of information. Since English digits have an articulation time somewhere in the neighborhood of 310msec to 321msec, basic arithmetic explains that English digit span should be around 7 items.[3,5]

General Structure of Short-Term Memory

The current preferred model for describing STM comes from consists of a three part structure: a "visuospatial sketchpad," a "phonological loop," and a "central executive."[6,7]

The phonological loop is itself divided into a phonological storage buffer and an articulatory rehearsal loop. Evidence for this division comes from two well-established phenomena: the phonological similarity effect and the effect of articulatory suppression. The phonological similarity effect refers to the fact that STM span decreases when the items to be recalled are too phonologically similar, which suggests that the phonological information is used for encoding and rehearsal.[8] Articulatory suppression occurs when subjects are instructed to continually repeat a certain utterance during a test of STM. This also results in a lower span, suggesting that mental representation of phonological information also draws on articulatory mechanisms.[9] It is important to note, however, that these studies almost universally use linguistic stimuli, and thus it has been difficult to parse out effects that are due to the linguistic nature of the stimuli which may differ from those due to the simple auditory nature of the stimuli.

The visuospatial sketchpad is subdivided into visual and spatial STM, which each exhibit a certain degree of independence with a certain degree of interaction, the exact nature of which is not yet understood. Tests of visuospatial STM are known to show a smaller span than the phonological STM, but it has once again proven difficult to establish whether this span of only about five items is due to the visual nature of the stimuli or to the non-linguistic nature of the stimuli.[10]

Signed Languages and their Contributions to STM Research

More recently, researchers have begun to parse out the effects of modality and language on STM by studying signed languages. Most of the research to date has been conducted using American Sign Language (ASL), a natural human language which exhibits all the properties of spoken languages with comparable complexity.[11] By studying signed languages such as ASL, researchers can ask questions about how STM is affected differently by linguistic and non-linguistic stimuli, and how the modality of input (visual vs. auditory) impacts STM span, because ASL provides a way to dissociate between auditory input and linguistic input.

However, before researchers could use ASL to inform the study of STM in spoken languages, it was important to verify that analogous components of spoken language STM were present for signed languages. The earliest evidence for sign-based parallels to spoken language phenomena came from Bellugi & Siple (1974), who were the first to demonstrate a sign-based phonological similarity effect.[12] More recently, evidence has come from Wilson and Emmorey, who demonstrated an interaction between phonological similarity, manual articulatory suppression, and a word length effect as evidence for an ASL phonological loop with a phonological storage buffer and an articulatory rehearsal mechanism.[13,14] This impressive correspondence between spoken and signed language phenomena showed that visuospatial materials can be processed linguistically under certain circumstances.

However, one important difference between spoken language and signed language STM is found in span length. English span is most commonly measured by a WAIS-style digit span task in which subjects are presented with lists of numbers in increasing lengths.[15] The longest length at which they can correctly recall at least one of two lists is considered their span. For English, this number is generally in the neighborhood of 7 items.[1] When this same procedure is applied to ASL, however, span is significantly shorter, around 4-5.[16,17] These results have been replicated several times in both Deaf and hearing native signers.[18]

The usual interpretation of this finding has been that, since ASL digits take longer to "pronounce" than English digits,[19] the reduction in span is simply an instance of the well-established word length effect. However, Boutla et al. present the results of ASL span measurements that carefully manipulated stimulus duration, presentation rate, and other timing factors to see if span increased when timing was more similar to English. No alterations of the temporal processes of the stimuli resulted in an increased span, either for Deaf native signers or for hearing native signers (also called CODAs: [hearing] children of Deaf adults) who learned ASL as a first language, even though these hearing signers showed a typical English span.[18]

Therefore, it seems that ASL span for native signers has a limited capacity that is not explained by the timing constraints to which spoken languages are subject. According to Boutla et al., the most likely mechanism behind this phenomenon is a modality effect: an advantage for auditory linguistic information that ASL does not share because of its visual nature.[18]

The performance of native signers, both Deaf and hearing, was unexpected. Since ASL has demonstrated all the other phenomena associated with linguistic processing, it is surprising to witness such a strong modality effect regardless of the linguistic nature of the stimuli. A question in the present research concerns whether this only holds true for native signers; second language learners of ASL might not be subject to this modality effect. No studies of STM in second language learners of ASL have previously been conducted. Our original hypothesis predicted that ASL digit span for native English speakers who were learning ASL would resemble the English span of about 7 items, on the assumption that adult learners of ASL would likely encode these materials in English.

Research on STM Processes in Auditory Second Languages

Before studying the behavior of second language learners of ASL, we must place their performance in the context of what is known about STM processes in second language learners of spoken languages. Research on STM span in multilingual subjects is a relatively new field, and the existing research focuses almost exclusively on "bilinguals," but does not always specify whether these are native speakers of both languages. One early study claimed to establish a difference in digit span for bilingual speakers of Welsh and English, crediting timing-related factors rather than proficiency-related factors as the cause.[3] However, methodological concerns and difficulty replicating these findings cast doubt on these claims.[20,21] Researchers who followed have asked whether STM span differences observed between a bilingual speaker's two languages are due to differences in sub-vocal rehearsal rate, or to proficiency-related factors.[22] Evidence seems to support the so-called "mother-tongue superiority effect," in which subjects show a span characteristic of their mother tongue even for stimuli in which subvocal rehearsal rates or other time-related factors would predict a different span.[23] That is, subjects who are bilingual in two languages with different characteristic spans (in monolinguals) show a digit span typical of the language in which they consider themselves dominant.

Seen in this light, the performance of CODAs (hearing children of Deaf adults) on the ASL span tasks is quite surprising! The mother-tongue superiority effect would predict that since these hearing native signers reported themselves to be English-dominant, they would show the English digit span of about 7 items. However, recent results show a span in the non-dominant language which is significantly lower than these subjects' own (dominant) English digit span of 7.05 and not significantly different from the native monolingual span of 4.91.[18]

The present two experiments tested subjects of even lower proficiency than these CODAs, reasoning that the mother-tongue superiority effect would be expected to have an even more significant impact for non-native speakers. However, even if these second language learners do show a shorter (more native-like) span, it will still be necessary to rule out proficiency concerns before convincingly attributing the results to a modality effect.

The first experiment set out to measure the basic ASL spans of second language learners who were advanced students but still learning the language. If these subjects show English-like spans, it would seem that they are still mentally representing information in their native language, despite having stimulus presentation and recall in ASL, thus following the predictions of the mother-tongue superiority effect. If, on the other hand, these subjects show native-like ASL spans, we will need further analysis to determine whether the subjects are indeed showing evidence of a modality effect similar to that experienced by native signers (hearing and Deaf), or whether they are performing mental translations between ASL and English, thus occupying STM resources which might otherwise contribute to increasing span length. In addition to the determining the basic spans, a manual suppression condition was added to assess whether subjects were using a manually based mental representation system. If manual suppression is found to affect STM span, this would constitute evidence that second language learners are mentally representing information in their second language and could imply a linguistic modality effect.

Experiment 1: ASL Spans of Second Language Learners

The participants were 11 hearing undergraduate students from the University of Rochester who were advanced ASL students. These students had all finished four semesters of ASL language instruction and completed or were currently enrolled in at least one upper-level ASL course taught in sign without voice interpretation. The video-recorded stimuli for the English and ASL digit spans consisted of the digits 1-9 and were ordered randomly to form lists of increasing length, starting at 2 items and proceeding up to either 8 or 9 items. There were two lists at each length. Lists of letters were generated by arbitrarily pairing selected ASL letters to digits, and replicating the lists from the digit span. To avoid chunking, lists were modified when this process produced meaningful acronyms. The particular letters used in the ASL task were chosen to avoid phonological similarity in ASL, although no similar control for their English equivalents was made. The stimuli were produced by a native signer at a rate of one item per second, paced by a visual or auditory metronome (as appropriate), and verified by frame-by-frame analysis. The signs were produced with a neutral facial expression, and only

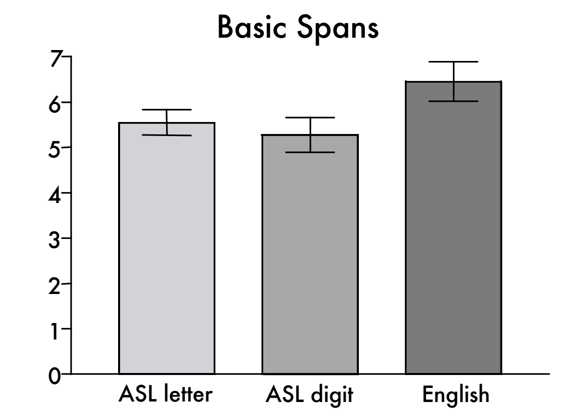

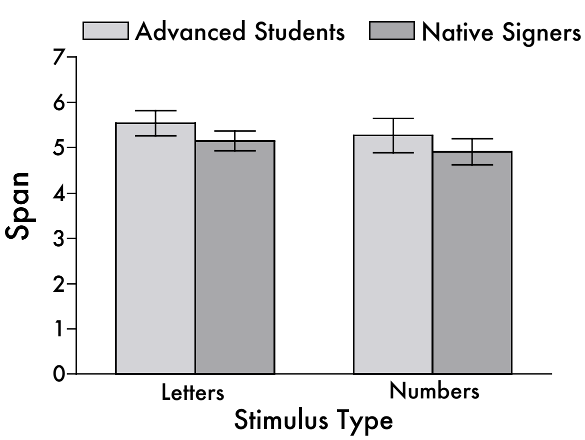

Top, Figure 1: Mean forward serial recall span of advanced ASL students. Error bars in this and all subsequent figures represent SEM. **Bottom, Figure 2:** Comparison of mean STM spans of advanced students and native signers. Native signer data from Boutla et al. (under review).

those mouth movements that would be natural for native ASL signing.

The above process was used to create five WAIS-style tests: two of ASL letters, two of ASL numbers, and one of English numbers.[15] All subjects received all tasks. Subjects were first given the basic WAIS tests in ASL for both letters and numbers, in counterbalanced order. Then they were given these tests again using a manual suppression task (see below), also in counterbalanced order. Finally, they were given the English digit span task. Using the WAIS definition, span was defined as the longest list length at which at least one of the two lists was recalled correctly.[15]

The manual suppression task was modeled after Wilson & Emmorey (1997), and consisted of the subjects opening and closing their hands in an alternating fashion (roughly from an ASL "5" handshape to an ASL "S" handshape") throughout stimulus presentation. This particular task was chosen because it had been demonstrated to be effective in the published literature and was a good analog to the types of interference tasks used in spoken language research in terms of obeying phonological and phonotactic constraints of the language.[9,13]

Mean basic spans for ASL letters, ASL digits, and English digits are shown in Figure 1. As can be seen, the ASL span in both tasks was lower than the English span. These spans were computed using a repeated measures one-way analysis of variance (ANOVA), which showed that ASL letter and digit spans were not significantly different from one another, but they were significantly lower than English span. Figure 2 shows how the ASL spans for the advanced second language learners compare with those of the native signers tested in a separate experiment.[18] Overall then, the relatively short ASL spans for second language learners resemble those of native ASL monolinguals, while their longer English spans resemble native English monolinguals. This is the same pattern observed in CODAs.[18]

These data indicate that advanced ASL students do not behave according to the mother-tongue superiority effect, which would predict that they should show the longer, English-like span for all tasks. The reasons for this, however, are not immediately clear. The subjects may simply be having difficulty in the ASL tasks because of their lower proficiency with the language. Alternatively, they might be mentally representing the information in English, with the observed decrement due to the mental resources needed to perform a two-way translation from ASL stimuli to English mental representation and then back to ASL for recall. As a third possibility, they may be showing a modality effect, with visual ASL stimuli producing a shorter span than spoken English stimuli. This difference could be a difference between modalities for linguistic stimuli, as has been suggested for native signers, or could be a nonlinguistic modality effect characteristic of other studies of visuospatial memory.

Performance in ASL span tasks using manual suppression was examined to determine whether participants were encoding the material in ASL. It was determined that manual suppression had no impact on span. This noneffect suggests either that subjects are not mentally representing information in ASL or that the suppression task itself was not robust enough to show an effect. Unfortunately, we have reason to believe that the latter may be true. Therefore, we cannot empirically discount the possibility that advanced students show low spans because they are mentally representing information in ASL. However, anecdotal evidence including unsolicited self-reports from several participants indicates that they may have been using English-based strategies for encoding and recall. In light of these two concerns, we recommend that the above experiment be replicated with a more robust suppression task.

The results of this experiment show a shorter span for ASL materials than for comparable materials presented and recalled in English, even for second language learners of ASL. However, it is not clear from these results whether this difference is due to a modality effect, as has been seen in native users of ASL, or to other factors particular to second language learners that could produce the same patterns.

Because results from Experiment

1 could not yield conclusive answers to these questions, Experiment 2 was designed with two modifications. First, subjects of lower proficiency in ASL were used to demonstrate that the performance of the advanced students was not likely to be due to difficulty in correctly apprehending the ASL-English correspondences. Secondly, and more interestingly, Experiment 2 allowed us to investigate the hypothesis that the observed decrement in span might be due to mental translation between the two languages. These hypotheses were investigated by including conditions in which the subjects were instructed to make a one-way translation from ASL to English or vice versa.

Experiment 2: Exploring Mental Translation

The subjects were 12 hearing undergraduate students from the University of Rochester who were considered to be intermediate-level ASL students. They were recruited from the population of students currently enrolled in their fourth semester of college-level ASL language classes, and had received minimal to no ASL exposure prior to college. The apparatus and stimuli were the same as Experiment 1, with one additional WAIS-style digit task in English.

The present experiment tested basic ASL letter and digit spans and English digit span as in Experiment 1, and added three new conditions: ASL stimuli with English recall (for both letters and numbers) and English stimuli with ASL recall (numbers only). All subjects received all conditions, beginning with the basic span measures for ASL letters and numbers (in counterbalanced order), followed by English numbers. Then subjects were presented with ASL stimuli but instructed to recall them in English (with letter and numbers run with order counterbalanced), and finally were presented with English numbers and instructed to recall them in ASL. The nature of the trials was the same as in Experiment 1, but run without manual suppression.

Mean basic spans for ASL letters, ASL numbers, and English numbers for these intermediate ASL students are shown in Figure 4. ASL spans were significantly lower than English, but there was no significant difference between the two ASL conditions. This finding parallels the trends observed in CODAs and in the advanced students from Experiment 1.

Figure 5 shows the basic spans of the intermediate students compared with those of advanced students. An unexpected trend seemed to emerge in which intermediate students consistently outperformed advanced students. Although there was no main effect of proficiency level, the trend observed on the graph indicates that it might be worth controlling for basic memory span when comparing span between groups. If a difference in basic memory span were observed using English digit span as a baseline, we might suspect that the subject pools simply differed in their basic STM capacity. For the present data, this control was performed by an analysis of difference scores. Since we were most interested in the difference between the subject's native language (English) and their second language (ASL), difference scores for each condition were computed (English digits minus ASL digits; English digits minus ASL letters), and a two-way ANOVA was computed as before. This analysis indicates that there are no significant differences in basic ASL or English spans between intermediate and advanced ASL students.

Most importantly, as previously noted, the intermediate ASL students show lower ASL spans than English spans, as do advanced ASL students.

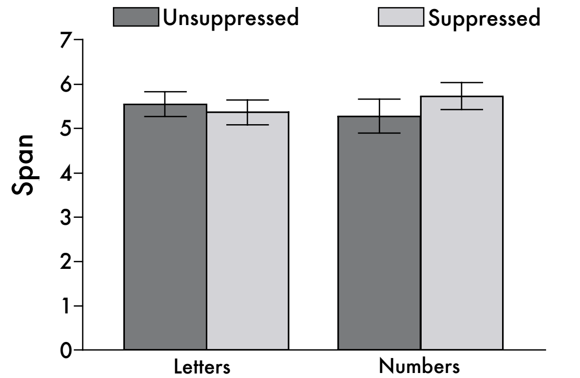

Figure 3: The effect of manual suppression on mean STM span for advanced ASL students.

Although this could suggest that the subjects are mentally representing the materials in ASL, two other explanations remain. First, it is possible that both groups of subjects simply have difficulty with the ASL–to-English correspondences because of their lower proficiency in ASL. Second, they may be making mental translations that deplete STM resources. To test this first explanation, span measurements from the one-way translation conditions were examined under the reasoning that if subjects show a longer span in these conditions, we can be sure that they have sufficient knowledge of the ASL-English correspondences. Data from these translation conditions also allow us to test our predictions about mental translation by comparing subjects' performance on the basic ASL conditions against the translation conditions. If subjects are already performing a two-way translation in the basic ASL conditions, we would

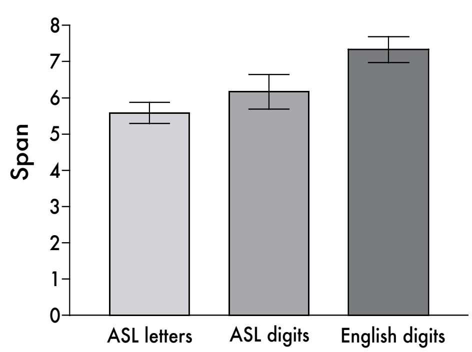

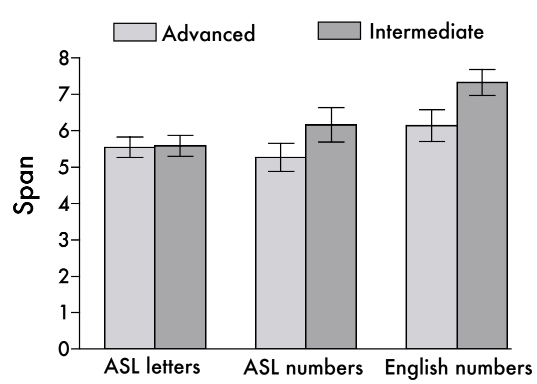

Above Left, Figure 4: Mean forward serial recall span for intermediate ASL students. **Above right, Figure 5:** Comparison of mean span for advanced and intermediate ASL students. Note unexpected superior (although non-significant) performance of intermediate students on both ASL numbers and English numbers.

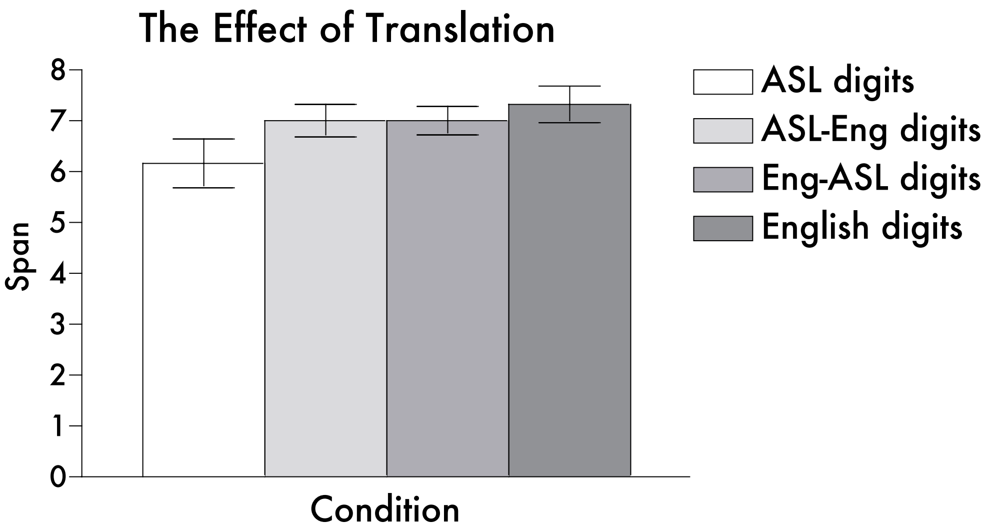

Figure 6: Mean forward serial recall digit spans for intermediate students for the basic ASL, basic English, and the one-way translation conditions

not expect the one-way translation conditions to be very different.

In comparing the ASL-to-ASL letters condition with the ASL-to-English letters condition in order to determine the effect of a one-way translation for letters, it was determined that the spans are not significantly different. This finding ran contrary to our expectations and likely results from the uncontrolled variable of phonological similarity. Although the letter stimuli were controlled for phonological similarity in ASL, many of the letters were phonologically similar in English (B, C, D, G, S, F), which is known to reduce English letter span.[8] For a more informative perspective, we turned to the data from the digit span tasks, where these concerns did not play a role.

Figure 6 presents the spans of intermediate students for the ASL-to-ASL, ASL-to-English, English-to-ASL, and English-to-English conditions. It is clear that the only span that is lower than the basic English condition is the ASL-ASL digit span. Two conclusions emerge from this set of data. First, it is clear that the intermediate students do not have problems making the correct ASL-English correspondences under one-way translation. Therefore, it is reasonable that the smaller span of the advanced students for the basic ASL conditions was also not due to difficulty understanding the input or correctly expressing the material in recall. Second, and more importantly, it appears that when subjects perform a mental translation where English is involved either as input or output, they show an English-like span, even in response to ASL stimuli. Since mental translation was an option available to the subjects in the basic conditions (and perhaps encouraged under manual suppression), the fact that they made no apparent use of it suggests that the observed difference between basic English and ASL spans is due to the modality effect rather than translation. However, it is possible that although a one-way translation for either encoding or recall may not have produced a significant decrement in span on its own, the aggregate effect of both translations might.

Therefore, we further investigated the hypothesis that the subjects were showing a low span in the pure ASL conditions because they were performing two mental translations: once from ASL stimuli to an English mental representation, and then back to ASL again for recall. According to this hypothesis, ASL span should be able to be predicted from the basic English span, minus the sum of the decrements for encoding and recall translations. We therefore compared the subjects' actual ASL digit span against the span predicted by the translation decrement hypothesis. This predicted value was determined by taking the mean basic English digit span, subtracting the difference between that and the ASL-to-English condition (the "encoding" translation) and then further subtracting the decrement between the basic English span and the mean from the English-to-ASL condition (the "recall" translation). This resulted in a predicted ASL span of 6.67, which was not significantly different from their observed ASL span of 6.167. For the sake of completeness, we then performed the analogous operations to compare observed English span against the English span predicted by the translation decrement hypothesis, given ASL span. The resulting prediction of 7.833 was not significantly different from the observed span of 7.33. Therefore, we cannot rule out the possibility that for second language learners, the observed decrement in span may not be due to a modality effect, but to two-way mental translation.

Finally, the intermediate subjects' ASL spans were then compared with those of native signers from Boutla et al. (under review).[18] It was discovered that the intermediate students show a longer span than native signers. This difference is primarily due to the intermediate students' high digit span.

Initially, this result seems to suggest that the intermediate students are behaving according to the mother tongue superiority effect, at least for digits. In other words, they are showing a longer digit span than native signers, which some might argue could be due to mental representation in English. However, it is important to remember that no control of basic STM capacity was made between these two groups. When controlling for basic STM capacity, no differences were found between advanced and intermediate students. If it were possible to control for basic STM capacity between intermediate students and native signers, we might not find any significant difference. (For a side by side comparison of all three proficiency levels, see Figure 7.) While this is far from a definite conclusion, it is helpful to consider this possibility, especially since we have no direct evidence that subjects are mentally representing the materials in English unless instructed to do so. Conversely, it is also important to remember that it has not yet been empirically demonstrated that any of the non-native subjects are mentally representing information in ASL even when their span is native-like in length.

Conclusions

Two hypotheses can accurately account for the pattern of data described in Experiments 1 and 2. One of these is the modality effect hypothesis, which postulates that STM for any visual information, even if it is linguistic in nature, is limited to a capacity of 3-5 items.[10,18] Support for this idea from the present experiment includes the observation that STM span for second language learners fails to obey the mother-tongue superiority effect that has been observed for spoken language bilinguals who are dominant in one of their languages.[23] Since the only significant difference between those studies and the present experiment is the modality in which the second language operates, modality is the clearest cause of the results. However, one puzzling piece of evidence is that monolingual English speakers perform equally well on a digit span task when stimuli were presented in the auditory and visual modality, suggesting that the mechanism behind the present result is more complicated than a simple modality effect.[24] Further research manipulating the nature of the visual stimuli for second language learners (signs vs. digit words vs. Arabic numerals) should help to disambiguate these results. If second language learners show a shortened span only for ASL conditions, the modality effect would appear to be linked specifically to ASL as a visuospatial language, and not to the purely visual or purely linguistic nature of the stimuli.

The second viable hypothesis is the translation decrement hypothesis, which postulates that subjects are mentally representing everything in their first language. According to this hypothesis, the observed decrement for ASL spans is claimed to be due to the mental translations necessary between encoding and recall. Our experimental manipulations reveal that this explanation has not been discounted. Note that this prediction runs counter to the mother-tongue superiority effect, but this effect was found in studies of "bilinguals," many or most of whom may have been native speakers of both languages. It is possible that the mother tongue superiority effect does not appear so directly in adult learners, who may show greater decrements when mentally translating to or from the non-dominant language. There is good reason to suspect that language processes used by second languages learners may differ significantly from those used by native speakers.[25] Research on second language learners (i.e. non-native bilinguals) of spoken languages with differing characteristic STM spans should provide useful insights into the mental behavior of second language learners.

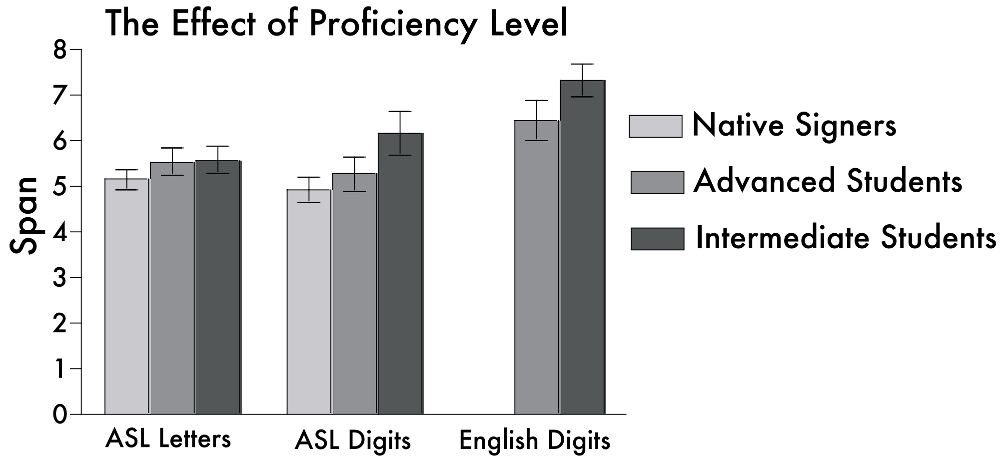

Figure 7: Comparison of mean basic spans for all proficiency levels. Native signer data from Boutla et al. (under review). Note the intermediate students' consistently higher performance on both ASL and English digits.

Matt Hall was a magna cum laude graduate of the University in May 2003 with a Bachelor of Arts in Brain and Cognitive Sciences (Highest Distinction and Honors in Research) and a Bachelor of Arts in American Sign Language (High Distinction). This article is the result of the independent research conducted throughout his senior year which earned him Honors in Research for BCS. Matt is still at the University as a research assistant to Daphne Bavelier in the department of Brain & Cognitive Sciences, where he continues to study the organization of spoken and signed languages in the brain through behavioral and functional magnetic resonance imaging techniques.

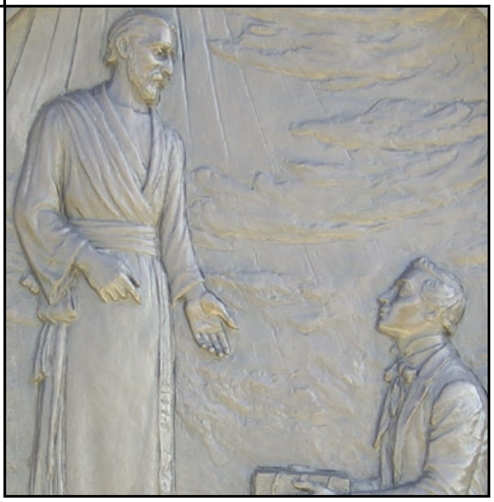

One of the faceplates on the monument atop Hill Cumorah.

Joseph Smith and the Mormon Church

A Case Study of a Modern Covenantal Prophet

Erin Zahradnik, 2003

Advised by Kenneth Gross, Ph.D.

Department of English

The term *prophet* is a subjective title in today's world—it refers to the bedraggled man on a New York City street corner, warning of an impending apocalypse; the ascetic who walks the streets of India, gathering followers with his messages of unconditional love; the mystic who reports visions of angels; and the ancient voices buried in the pages of the Bible. How does one distinguish a true, divinely-sanctioned messenger from a false prophet? In one sense, it is impossible to designate a true prophet as someone whom God spoke to or spoke through, because there is no way to definitively prove such an event occurred. Neither science nor religion has yet found a way to quantify God. The belief that someone was a true prophet lies completely in faith, an indefinable *knowing* of one's own definition of absolute truth and the resonance of that definition in the teachings of someone else, the 'true prophet.' Such faith is subjective; each individual experiences it in slightly different ways. As such, true prophets are in fact defined by each individual person; it is one's faith in a prophet that makes him a true prophet, and it is this ability to uncover faith in people that defines a *successful* prophet.

One of the major requirements for a successful prophet is sincerity. The prophet must truly believe himself to be called by God. He must not doubt the veracity of the visions, prophecies, covenants, and doctrines that he receives. Although he might doubt his own ability to carry out his duty, he must *never* doubt God, because God is his reason for being.

A prophet has a "dual battle of sympathy," in that he is imbued with the divine ability to see man from God's point of view, on top of his innate ability to see God from man's perspective.[1] A prophet embodies all of mankind in himself during his interactions with God. Fright, struggle, guilt, and desperation; we see these emotions in the Biblical major and minor prophets, in Jesus of Nazareth, in the Prophet Muhammad, and in the Prophet Joseph Smith.

Such feelings are not indications of weakness; they are instead signs of truth and sincerity. A prophet is, despite his divine calling, a flawed human being, not perfect as only God can be. There is no one more aware of this fact than a prophet. Indeed, he is tormented by this because he feels his inadequacy all the more potently through constant interaction with God. Even so, the prophet must speak for God, not only with his voice, but with his entire being; he must speak as though his "life and soul are at stake in what he says and in what is going to happen to what he says".[1a] And they *are* at stake! What happens to mankind also happens to the prophet.

Such prophetic anguish manifests itself as a palpable, desperate faith. When a prophet is called to be a messenger of God, this faith and the ability to inspire such faith in others is what creates believers. Therein lies the core commonality of the true religious prophets of the world, the successful prophets who delivered their messages, who had such deep faith that others could not help but listen to the prophecies with body, mind, and soul. Such an ability in a prophet is far beyond simple charisma, it is that undefinable empathic ability of embodying all of mankind in a single person, and of reflecting back to mankind the prophet's desperate struggle. To feel the entire world in oneself is both glorious and torturous, but that is what a true prophet must do.

The Act Of Prophecy

The actual act of prophecy encompasses both "inspiration and experience".[1b] This inspiration comes at a single dynamic moment; it is "an act of giving; not an eternal word, but a word spoken".[1c] "In the beginning was the Word…and the Word was God...In him was life; and the life was the light of men".[2] There was a word, a message of God that was God himself, a giving of himself to mankind in the most intimate way possible. When mankind blots out that light, they no longer see life in its ultimate glory, they no longer see life as God, as a transcendental experience in every moment, in every touch and sound and sight. A prophet is called on to restore that light to mankind, to show them what life truly is—not a timespan of simple existence, but a state of being *in* God, *part* of God, of being the light of absolute transcendental divine existence in every moment.

This existence is what the prophet experiences during his visions and his communions with God. It is something that comes "upon [the prophet], not out of him"; something that is much greater than a single moment, that happens not to a single man, but to the collective man, because the prophet is acting as all of mankind.[1d] It is the difference between simply hearing the word and having the word become part of oneself: "Thy words were found, and I did eat them".[2a] The prophet *becomes* that light, the prophet is inside of God and God is inside of the prophet. Barriers have ceased to be and the prophet becomes the revelation; he understands, in the intangible manner of unshakeable *knowing* and faith, what 'dwelling with God' *really* is.

Covenantal Prophets And Prophetic Personality

The Hebrew prophets dealt with the present, historical concerns of God's people, but it was the duty of another subset of prophets, heretofore referred to as *covenantal prophets,* to introduce, on God's behalf, entirely new covenants to mankind. While the Hebrew prophets strove to understand God's relationship with man, covenantal prophets sought to *change* mankind's relationship with God, to add new aspects to it, to bring man closer to God and the Word.

Covenantal prophets are the messengers through which God offered covenants to humanity, covenants which were much more than laws or acts of negotiation. They were experiences of God giving himself to mankind in return for mankind doing the same for God. They were events of surrender where God and man become not Creator and Created, but simply *Creation*. The covenantal prophet is a living symbol of the covenant, of man's struggle to reconcile himself with the idea of a personal God, a God embodied inside himself.

Moses, Jesus of Nazareth, Muhammad, Joseph Smith, Baha'u'llah: all were covenantal prophets, each delivering the *logos spermatikos*, the embryonic essence of the Word that exists everywhere.[1e] Despite their individual existence in widely disparate historical time periods, they all embodied the essences of a true prophet: the intense prophetic personality that consumed the Hebrew prophets, and the precarious dual self-identity as man's representative to God and God's representative to man.

Covenantal prophets fulfill four major criteria: (1) they receive direct divine delivery of a "new" covenant; (2) they have a major hand in founding the new religion that comes out of the new covenant; (3) they display human personality—struggle, anger, sadness, perhaps even transgression, all associated with their enormous divine calling; and (4) they have unshakeable inner faith in God and his message, as evidenced by their influence on humanity during the time of their living prophecy.

Decoding prophetic personality is perhaps the most challenging aspect of identifying covenantal prophets. The majority of the prophets lived and died before literacy was widespread, and thus there are no personal journals by the prophet, only snippets of personality culled from Scriptures and revelations attributed to the particular prophet. But when a prophet is called to duty in the modern world of printing presses, literacy, and newspapers, his life is much more readily examined. Putting aside the question of veracity in favor of the prophet's success at delivering his message, one name springs to the forefront: Joseph Smith, the American covenantal prophet of the Church of Jesus Christ of Latter-day Saints (LDS), also known as the Mormon Church.

Joseph Smith: A Case Study Of A Modern Covenantal Prophet

Joseph Smith was a poor farmboy, a truant at times, who spent his teenaged years searching for buried treasure on his family's farm in upstate New York. Located in the middle of 'the Burned-over District,' so-named because the region was 'burning with the fires of the spirit,' Joseph grew up surrounded by constant religious fervor.[3] A combination of "Yankee…moral intensity," industrious immigrants, and an upsurgence of evangelistic preachers led to an environment where "emotional religion was…a congenital characteristic, present at birth".[3a]

> *A typical young man in the Burned-over District awaited the day when the Holy Spirit would marvelously elect him to church membership… He might…change [religious] affiliation several times as one sect or another held services nearby… He could scarcely evade a religious experience in New York.*[3b]

This was an environment where it was common for people to have divine visions, visits from the Holy Spirit, and ecstatic religious communions with God in the midst of revivals. It was the perfect place for a prophet to appear.

In 1820, 14-year-old Joseph Smith sought refuge in the woods near his home in Palmyra, NY, in order to pray for an answer to his question of which church he should join. In response, he had a vision in which God and Jesus appeared to him and told him that he should not join any church, because, as the light-enshrouded figure of God said, "they draw near to me with their lips, but their hearts are far from me, they teach for doctrines the commandments of men, having a form of godliness, but they deny the power thereof".[4] Despite God's appearance, Joseph had not yet been informed of his impending prophetic duty, but he *had* received personal revelation, an experience that would play an important doctrinal role in the future LDS Church.

Although such a vision had to have a deep psychological impact on Joseph, he was relatively reticent about it, merely telling his mother that "I have learned for myself that Presbyterianism is not true".[4a] The friends with whom he did share details were antagonistic; according to Joseph, a Methodist preacher friend "treated my communication not only lightly, but with great contempt, saying it was all of the devil".[4b]

In itself, a vision was an exciting event to an individual, but one must remember the socio-religious environment of the Burned-over District—visions and other ecstatic religious experiences were not uncommon. It was the *nature* of his vision that caused Joseph such persecution. This was no simple case of an answer to a prayer—this was *God* saying that all of the churches were *wrong*. To disbelievers of Joesph's vision, it was heresy of the greatest degree, but Joseph could not doubt the veracity of his personal revelation.

On September 21, 1823, Joseph had another divine visit, this time from the angel Moroni, who informed Joseph that "God had a work for [you] to do; and that

[your] name should be had for good and evil among all nations".[4c] Moroni then explained that there was a book written on golden plates, the lost gospel that Christ had delivered to early American inhabitants, buried on a nearby hillside along with two stones that "God had prepared…for the purpose of translating the book".[4d] Joseph's call to covenantal prophecy had begun.

Revelation In The Book Of Mormon

Joseph's translation of the golden plates (the Book of Mormon) was revelation in itself, both because heretofore hidden truths were being revealed, and because Joseph was given divine assistance while translating. The angel Moroni provided Joseph with "two transparent stones, clear as crystal," spectacles of a sort, called the Urim and Thummim[1,5] Using these divine stones, Joseph "translated [the Book of Mormon] by the gift and power of god, by [the] means of the urim and thummim".[6] In order to receive revelations, "[t]hose using [the stones] offered their prayers unto the Lord and the answer appeared written with letters of light on the Urim and Thummim, but disappeared again very soon after".[5] This was divine revelation in the most literal sense, almost magical, and certainly miraculous to witnesses, which furthered the veracity and reality of Joseph's prophethood among his followers.

The lingual structure of the Book of Mormon is very similar to that of the King James Version of the Bible (KJV); the stories it relates, however, are radically different. The relationship between man (mainly the prophets) and God in the KJV Bible is a two-sided dialogue where God and man have the power to influence each other.[1f] The Book of Mormon takes this relationship even further into the realm of human terms—by incorporating the petitionary prayers of individuals (1 Nephi 2:19-24; 1 Nephi 11:2-36; Ether 3:9) and the answers that God gives them, God becomes more human, a definable character, and almost finite.[7]

A finite God has enormous implications for a theology—by defining God's revelations in terms of personal, literal dialogue, one runs the risk of assigning a certain changeless *personality* to God, which in turn denotes that you, the individual, and the collective believers, know the limits of God.[7a] This is a God who people can *know* in such a way that they are comfortable engaging in dialogue with him, asking for and then fully expecting his help in their daily lives. This is an act of *knowing* God in a much deeper and psychological sense than the traditional Christian definition of *knowing*, which typically implies only that one accepts the existence of God and his influence on people. This Mormon *knowing* is on a much more personal and intimate level, and stems from a unique aspect of the Book of Mormon—the immense importance of personal divine revelation.

This personal revelation in the Book of Mormon is given not just to prophets but to almost every personage in the book.[7b] While divine revelation in the Bible is relegated to the prophets, the most important aspect of revelation in the Book of Mormon is that it is widespread, available to everyone; in fact, it is every person's *right* to receive divine revelation.[7c] Of course, this revelation is not on the order of a covenant, but that of more ordinary, everyday matters—where the best place to hunt is, where the enemy camp is, further explanations of 'God's mysteries' as manifested on Earth.[7d] Covenantal revelation, however, involves much greater truths that have either not yet been told to man, or that have been so lost through time that God's direct and concerted effort is required to unveil them. In the case of Joseph Smith, this effort was manifested in the production of the Book of Mormon, which once again revealed certain truths that had long been absent from religious practices.

It was imperative that Joseph complete the translation as soon as possible, not simply because God wanted it, but out of necessity for the spreading of this new covenant. While Joseph's oral revelations, whether from the Holy Spirit, the angel Moroni, or from God himself, were used to build the theological and sociological structure of the Church, they were not enough to spread his latter-day message. A book was symbolic of something greater than a mere man; it was a foundation to the Prophet's message, it gave a sort of divine sanction to his words. The early Mormons saw it as a solid, unchanging testament of the veracity of their Prophet. "From eternity to eternity [God] is the same," Mormon scripture asserts.[8] The Book of Mormon was needed as proof of that sameness, to show that the truths of the Bible were also evident within it, to prove that the same deity who inspired the Bible also inspired the Book of Mormon. To Mormons, both books are revelation, both God's words, both truth, but one is more accurate truth because it was translated by God himself, not by man. It was Joseph's proof, the living evidence of the veracity of his visions and revelations. One did not have to be a prophet to truly experience this new revelation.

In and of itself, the Book of Mormon corroborated with the Bible; many passages are essentially reworked from the Bible (21 chapters of Isaiah contained in 1, 2, and 3 Nephi and Mosiah 12, 14; Malachi 3-4 in 3 Nephi 24-5; 1 Corinthians 12:1-2 in Moroni 10; Acts 3:22-6 in 3 Nephi 20).[9] However, the 'hidden truths' it revealed clashed with some aspects of the Bible. This is understandable if, as Joseph said, translation had marred the original meaning and text of the Bible. With such focus on the subjectivity of translation, it is not surprising that God eventually commanded Joseph to commence an "inspired translation" of the King James Version of the Bible.[10]

The Nature Of Revelation In The JST Translation Of The Bible

Translation has historically been a major topic of contention in the area of religious writings. This is understandable—it is a task in which man must be able to interpret and understand the Word of God in its original written language well enough to translate the very spirit of the Word into another language. This is truly an impossible task theologically, because it requires that man understand the underlying reasons and goals of God's actions, and to presume to know what God was trying to convey with each word, each letter.

It is not shocking then, to purport that "no human being is able to reproduce God's word as revealed to man in its entirety unless a new revelation takes place in which God makes known His word to the translator in a new language".[11] According to this viewpoint, every translation of the Bible is wrong. Because of its divine origin, every word and every letter of Scriptural text "contains God's mystery".[11a] Any change at all to the original text takes away some aspect of God's revelation, rendering it hidden from all readers of that particular translation.

While the major reason for Biblical translation is to put the word of God into the hands of non-Hebrew, non-Aramaic, and non-Greek readers, it is also an "opportunity to give divine sanction to one's ideas".[9a] Because translation is really interpretation, one's particular religious beliefs will inevitably be incorporated into one's translation. They will shape interpretation and thus shape the end result—at the very best, imbuing a passage with meaning that was not originally intended, at the worst, inadvertently *changing* the original, intended meaning.

Truly, only one type of person can justify major addendums to Scripture, and that is a covenantal prophet. Because they act on God's direct instructions and new revelations, changes in laws, wording, even theology is technically permissible, as long as one agrees that the particular prophet actually *is* acting under God's directives rather than on his own inclinations.

The 8th Article of Faith of the LDS Church states: "We believe the Bible to be the word of God as far as it is translated correctly."[8] Joseph's reasoning, as he explained to his followers, was that the history contained within the Book of Mormon highlighted the fact that many truths in the Bible had been lost to years of mistranslation and misinterpretation.[10a] In a way, this is undeniably true—some aspect of the original work is always lost in translation. Not only are there copyist mistakes to take into account, one must also consider that languages are not mirror images of each other, and many verbal nuances are simply untranslatable.

But Joseph did not simply correct sentences or word choices; he changed sentence structures, he added entirely new sections and scenes to various books, and he inserted new aspects of theology. In doing so, he changed the nature of the man-God relationship as contained in the KJV Bible, added references to Jesus in the Old Testament, and inserted information on the age of accountability of children (age eight), the afterlife and God's ultimate plan, the nature of premortal existence, the roles of the priesthood, and the priesthoods of Enoch and Melchizedek.[17]

Small sentence structure changes can have large effects on theological doctrine—several passages in the KJV (Jer. 26:13, Amos 7:3) include the text "the Lord repented." The Joseph Smith Translation (JST) redirected the repenting to a human, rather than God, for the Prophet Joseph decreed that God does not repent because he is perfect; only man repents.[9b] Similarly, instead of God hardening Pharaoh's heart in the KJV Exodus, the JST version has Pharaoh hardening his own heart. The responsibility of human failing is removed from God and transferred to imperfect man. While such revision does strengthen the idealistic 'perfection' of God, it also adds to the 'changelessness' of God—God cannot change his mind, because he does not make mistakes or reconsider actions; everything that he does is perfect and complete the first time. Combined with the personal, dialogical revelation of God in the Book of Mormon, the JST changes are almost paradoxical about the nature of God. God is at once like man, yet far removed; perfect, yet once imperfect; changeless and eternal, yet he was not always God. The JST explains God, while also making him even more of a mystery, a paradoxical being that is both God and man in one essence.

A being who is often considered in Christian theology to be both God and man at once, Jesus Christ, is subject to a new theory of origin in the JST. The JST version of the early chapters of Genesis is included in the *Pearl of Great Price* as the "Book of Moses." The presence of the *premortal* Jesus is evident from the beginning of mankind's existence—an angel appears to Adam in the Garden of Eden and tells him: "thou shalt do all that thou doest in the name of the Son".[1h] The coming of Jesus is also revealed to Noah (Moses 8:24), Moses (Moses 1:6), Enoch (Moses 6:52), and to the Hebrew patriarchs (JST Bible). This addition of prophetic prediction of the Messiah is the reason behind the extreme emphasis of Jesus in the LDS faith, but this is not the typical Christian version of Jesus. The LDS Jesus is both a redeemer, a prophet, and 'spirit-child' of God. This Jesus has always existed alongside God; the prophets spoke of him from the beginning—he was a covenantal prophet, but a unique one who brought a crucial covenant to mankind, a covenant of which the full import was lost from history until Joseph was called to restore it.

View from the top of Hill Cumorah, where Joseph dug up the plates.

The presence of a premortal Christ introduces an entirely new view of the realm of God—the spirit world. Everyone is "born as a spirit to Heavenly Parents before being born to mortal parents on the earth…[and] lives with God as his spirit children before coming to the earth as mortal beings".[15] While KJV-based Christian doctrine recognizes a distinctive 'soul' or 'spirit' versus the physical body, Mormonism takes this a step further and separates existence into 'spiritual' and 'physical.' Thus the passage in Jeremiah, "Before I formed thee in the belly I knew thee" takes on a whole new meaning.[2b] This is a literal *knowing*, not so much fore-ordained in the sense that God transcends linear time and can see the past, present and future equally well, but as God tells the JST-version of Abraham, "thou wast chosen before thou wast born".[12] The prophets were spirits with God in the spirit world, and were chosen at *that* point to be mortal prophets. It is no longer a matter of the measure of man in life that makes a prophet, but the measure of the man's spirit in his *premortal* life. A prophet is a prophet before he is born, not just when his call to prophecy begins; a prophet is a prophet before he even *knows* that he is one.

On top of the new revelation of the existence of a spirit world, the entire issue of creation underwent revision in the JST. In fact, creation becomes a relative term—instead of the KJV "God created the heaven and the earth," the JST version says: "the Gods organized and formed the heavens and the earth".[12a] The theological import of this passage is immense; all of Judeo-Christian-Islamic theology holds monotheism and the concept of "there is no god but God" at its very core. Joseph Smith's revelation in the Book of Abraham breaks with this core in such a way that one *cannot* classify Mormonism as a Christian sect; it is an entirely different religion, albeit heavily based on Christianity. Although Mormons recognize one god as 'the heavenly Father,' doctrinally they acknowledge the existence of other gods.

Such acknowledgement is implicitly required when one takes the Mormon nature of God into account. Modern Christianity agrees that while man was made in God's image, man can never *become* God, he can only hope to dwell with him. God's nature is unreachable, untouchable for mere mortal men. But the ultimate state of existence in Mormon theology is godhood—Joseph revealed that in the highest kingdom of heaven (the celestial realm), very devout and pious Mormons will dwell with God *as equals*, as gods themselves. In the celestial realm of heaven, God will make believers "equal in power, and in might, and in dominion" with himself.[8a] This is the crux of the new covenant that was revealed to Joseph Smith—that man might become like God if he followed the laws contained in the revelations that his latter-day prophet received. To follow God's commandments is to emanate God and the state of godhood that a person can attain after physical death.

If man is able to become a god himself, then what *is* God? Or more importantly, what *was* God? He may be an "unchangeable God" now, but, according to the revelations of his prophet Joseph, "God himself was once as we are now, and is an exalted man, and sits enthroned in yonder heavens!...[H]e was once a man like us…God himself, the Father of us all, dwelt on an earth".[8b,13] With this revelation, Joseph tore away a good bit of the mystery of God. No longer was he the ethereal, indefinable being of Judeo-Christian theology—he was something definitively finite, something palpable, something within the reach of human experience.

In fact, the Biblical prophet of sorts, Enoch, is an incredible focus in Mormonism, precisely because of this 'man as God' theology. Even in the KJV, Enoch does not die a normal death; "Enoch walked with God: and he was not; for God took him".[2c] Judaic mysticism explicates this passage in the phrase "Enoch is Metatron," meaning that Enoch encompasses man, God, and angel into one figure.[14] In Mormon theology, Enoch represents the ideal of man and man's purpose—to create a Zion, the Mormon New Jerusalem, and then to dwell with God in one's original highest spirit state.[14a] The Melchizedek priesthood that Enoch represents is the divinely sanctioned position where one "hold[s] the keys of all the spiritual blessings of the church".[8c] This is the state to which man must aspire, to act as a prophet of God, to embody God on earth through this priesthood. This is to walk the path that God once walked when he was a man.

Because the almost unanimous agreement among theological religions that the nature of man is *not* akin to the nature of God, Joseph Smith had to have had tremendous faith in his own revelations to go against millenniums of theology. The concern at hand is not whether Joseph's visions and revelations were true, but the manner in which the revelations of God's new covenant changed the relationship between man and God. Joseph had endless faith in the divine sanction of his revelations, and that belief is not to be taken lightly; such deep-rooted faith in revelation translated itself to his followers and is at the core of Mormonism.

Continuing Revelations

Somewhat at odds with the exalted position of written Scripture in modern religion, Orson F. Whitney, an ordained Apostle of the LDS Church, asserted in 1916 that "no book presides over this Church and no books lie at its foundation".[7e] Because "the First Presidency and the Quorum of the Twelve Apostles are prophets, seers, and revelators to the Church and to the world," such a statement was meant as revelation and thus, fact.[15]

Canonical evolution is true of most major religions, of course, especially in American religions. But Mormons readily *admit* this; in fact, continuous canonical change is a major part of the LDS religion. The ease of this admittance is probably due to the presence of a living prophet—in this way, it is always God who decrees changes that follow the flow of societal change. In other forms of Christianity, it is often the people who demand a change of Church doctrine and policy; God is left out altogether. The only major Christian branch that parallels Mormonism in this sense is the Catholic Church, with the Pope acting as a prophet of sorts. But the Pope, while said to be the closest living man to God, is acknowledged by Catholics to act as a man, to issue decrees from a human standpoint, and not to act as merely a mouthpiece of God. When the LDS Prophet (the President of the Church) speaks on doctrinal issues, he is assumed to be 'inspired' by the Spirit, by God. "Today's inspired utterances may become part of tomorrow's standard works"—this is how easily LDS canon can change when the Prophet speaks.[7e]

Continuing revelation is an LDS doctrinal standard: "God … is not silent. The Church of Jesus Christ of Latter-day Saints is led by revelation from God to a living prophet, who receives God's word just as Abraham, Moses, Peter, and other ancient prophets and apostles did".[16] While much of Christianity characterizes truth as something unchanging and eternal like God himself, in Mormonism, today's truth might not be tomorrow's—truth is a fluid, living entity. As Joseph Smith defined it, "truth is knowledge of things as they are, and as they were, and as they are to come".[8d] God is the "Spirit of truth," continually being revealed.[8e] Such revelation is an eternal, endless process, because God himself is eternal and endless.

Conclusions

The depths of prophetic identity cannot be understood unless one truly tries to see through the eyes of a prophet. Who *was* the Prophet Joseph Smith? Imagine what it would be like to want an answer so desperately, to throw oneself entirely into a prayer so that *self* dissolves and there is only the question. In this moment, God the Father and Jesus Christ appear and reply to the prayer and the petitioner. There are no disembodied voices, no whispers on the breeze, no subtle signs from the heavens; this is *God*, appearing as he did for Abraham and for Moses, now appearing to Joseph Smith. And what was Joseph? A poor farmboy, a younger son? A man who wanted more than anything to make his mark on the world? Who had potential and the talent but did not understand how to channel it? Out of everyone in the world, God came to *Joseph* and prepared him to be a prophet. And not just an ordinary prophet, a *covenantal* prophet, the rarest kind.

Why Joseph? And yet, why Jeremiah, why Ezekiel, why Moses and Muhammad? In a sense, these are unanswerable questions; they presume to understand the mind of God, and ordinary theology says that such understanding is impossible. But there must be some key aspect of chosen prophets, some spark that calls God, that God perhaps *identifies* with. If man was made in the image of God, what *is* that divine image? It is quite far beyond anything physical; God is not about form, God is about the mind, about transcendence, faith, deep connection. *That* is the realm in which one may find the image of God, and that is where the covenantal prophets innately dwell. They embody God without realizing it.

Covenantal prophecy is an evolving calling, depending on the state of society for the type of revelations that God bestows on his chosen prophet. Early societies appeared to respond best to impassioned speaking, physical metaphors, and parables. More modern societies increasingly required literary revelations in the form of books--God's divine words translated into written language, something that everyone could have and hold, a form of personal revelation. While the method of prophetic delivery to mankind may be a changing entity, the true covenantal prophets are always characterized by distinctive psychological qualities—their faith and their triumph through suffering.

While the revelations of Biblical prophets were tinged with divine pathos, God's sorrow for the mess that mankind had made of his perfect world, Joseph Smith's revelations rang with divine exaltation.[1g] The time of latter-day prophets had arrived; mankind could receive the lost scriptures and truths and understand God's plan more fully. They could dwell with God, become gods, achieve transcendence. This Mormon God is not a jealous god, angry about idolatry and transgressions; he desperately wants his creations, his 'spirit children' to understand their true divine potential, to embrace the godhood that dwells inside of them, to *become* an equal, and to beget their own spirit children on their own Earths.

Such revelation is the Prophet Joseph Smith's greatest gift to mankind, whether one chooses to believe its truth or not. Joseph took Christian theology far beyond the rigid bounds that society had given it, building an entire religion based entirely on faith, rather than rooted in dogma. The basis of Mormonism is personal revelation, and one can never receive the Spirit, which confers such revelation, without deep, unyielding faith. Joseph's great miracle as a prophet was to re-introduce the value of faith to mankind, to be a living demonstration of how far faith can carry a person, to show that God is not far removed, but is with every person. All one had to do, Joseph wordlessly demonstrated, is listen.

Hill Cumorah monument with angel Moroni on top.

Erin Zahradnik graduated from the University of Rochester in 2003 with a BA in English and a Take Five program in Astrobiology. This paper is the cumulation of an Independent Study on the Literature of Prophecy taken with Professor Ken Gross in Spring 2003. Erin is currently working in the Rochester area and will soon begin applying to medical schools.

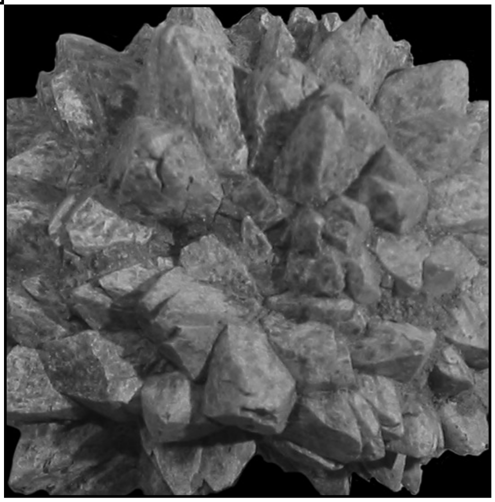

A representative glendonite from the Deer Bay Formation, Blackwelder Mountains, Ellesmere Island.

A cold start to the Cretaceous

Defining climate history near the North Pole

Peter Lippert '03

Advised by John A. Tarduno, Ph.D.

Department of Earth and Environmental Sciences

Contemporary political summits and policy decisions, like those at Johannesburg and Rio de Janeiro and the Montreal and Kyoto Protocols, are clear evidence that world leaders and their communities are increasingly concerned about the Earth's climate and its future. Policy makers frequently turn to scientists for definitive answers and pragmatic solutions to climate problems and too often focus on only anthropomorphic influences. Although human impact on the environment certainly deserves the attention it receives, all of us must recognize that the Earth's climate is a complicated and multi-faceted system that responds to inputs on both human and geologic time scales. The correct interface of science and policy should consider not only anthropogenic stimulants to Earth's climate, but also the natural responses to and controls of this system. Scientists who try to explain, model, and predict future climate change examine Earth's former climates in the geologic record to test their models and predictions. Past climates let us not only gauge our current understanding of Earth's dynamic climate, but also evaluate how climate varies over time periods that span innumerable human generations.

Most geologists regard the Cretaceous (a period of geologic time 145 to 65 Ma, where Ma denotes million years ago) as the model "greenhouse" Earth. The planet was ice-free, sea levels were very high, and carbon dioxide levels were four-to-eight times the present level.[5,14,29,32,66] Cretaceous climate was characterized by warm temperatures and low global temperature gradients.[3,36] Fossil plant and animal data are also suggestive of a very equable climate: there were forests and ectothermic reptiles (*e.g.*, "cold blooded" champsosaurs and turtles) at polar latitudes.[34,66]

Recent research, however, suggests that this anomalously high, globally distributed warmth did not characterize the entire Cretaceous. Several authors report evidence compatible with the presence of polar ice and strong global temperature gradients at the beginning of the Early Cretaceous.[17,31,32,60,61,62,70] For example, Abreu *et al.* correlated Cretaceous and Cenozoic (the last 65 Myr) stable oxygen isotope data and eustasy (global sea level) and concluded that there is a positive correlation between the two data sets. The ratio of ^{18}O to ^{16}O in ocean and meteoric water is temperature dependant, which means that it can be used to estimate paleotemperatures. The positive correlation between stable oxygen isotopes and sea level suggest that there may have been continental ice in the Albian (late Early Cretaceous, 110 Ma).[1] These results suggest that climate may have been dramatically different during short (< a few million years) intervals when compared to the elevated mean temperature that characterized most of Cretaceous time.

Stable carbon isotope data from low latitude Tethyan (the Tethys was an ancient global, subtropical sea) sites also suggest that there were short but significant perturbations in the Earth's carbon cycle during Valanginian (137-132 Ma) and Albian-Aptian (121-99 Ma) time and Cenomanian-Turonian (99-89 Ma) time.[10,71] Organisms prefer to utilize ^{12}C in their metabolic reactions, so organic carbon is enriched in ^{12}C; conversely, inorganic carbon, such as carbonate, is often enriched in ^{13}C. Dramatic changes in the $^{13}C/^{12}C$ ratio (usually displayed as $\delta^{13}C$ values, where positive $\delta^{13}C$ values indicate an increase in inorganic carbon relative to organic carbon) indicate a change in how carbon is partitioned in the global carbon budget; these changes are referred to as carbon isotope events. There are many characteristic C-isotope events during the Cretaceous, and they may be related to widespread volcanic outgassing that elevated atmospheric CO_2. CO_2 is an effective greenhouse gas, so more CO_2 would amplify greenhouse conditions and spur climatically induced ecological changes, which in turn would have significant effects on the global carbon cycle.[73] One possible feedback follows: as the Earth became more humid, increased chemical and mechanical weathering and increased continental runoff added significant quantities of nutrients, particularly phosphate, to the oceans. High nutrient levels could have stimulated biologic activity and indirectly created an ecological crisis in which increased marine productivity and subsequent eutrophication led to widespread anoxia in intermediate and deep waters. These ecological crises are referred to as oceanic anoxic events (OAEs) and are often coeval with extensive reduced black shales in the marine record, especially during the Cretaceous.

Most research on Cretaceous climate has been conducted on low latitude rocks and results are typically generalized to represent the global climate. Recent results from high latitude areas, however, suggest that Cretaceous temperatures were episodically cooler than the Tethyan record indicates.[17,24,60] Typical paleoclimate studies of the Cretaceous have relied on climate proxies such as stable oxygen isotope data, which are affected by diagenetic processes and are subject to various interpretations. Although some authors have reported more direct climate proxies— flora and ectothermic reptiles— there are still few direct Cretaceous climate indicators from high latitudes.[60,66] For these reasons and others, polar regions are particularly important for understanding past climate change, especially with a limited supply of available geologic data.

Kemper's and Schmitz's description of glendonites in the Deer Bay Formation of the Sverdrup Basin in the High Canadian Arctic is an important but often overlooked development in Cretaceous climate studies.[48] Glendontites are characterized by their unique aggregate euhedral shape [see title figure] and form when the mineral calcite replaces ikaite ($CaCO_3 \bullet 6H_2O$), the high pressure, low temperature species of hydrated calcite. Although rare, glendonites are found throughout the geologic record and are sometimes associated with dropstones and rafted wood. Natural ikaite of Recent age is found only near the sediment-water interface in near-freezing marine and continental marine environments (*e.g.*, Ikka Fiord, Greenland; Bransfield Straight, Antarctica) and has never been found in continental marine waters above 4°C.[55,63] De Lurio and Frakes determined that near-freezing temperatures coupled with high alkalinity and elevated orthophosphate levels are necessary for ikaite precipitation in the normal marine environment.[15] Ikaite precipitation is not favored under such conditions, however; it is only facilitated. These associations and its restrictive precipitation controls suggest that ikaite (and therefore glendonites) may be an unique and unambiguous cold-water climate indicator.

In a recent review, Price attempted to correlate the occurrence of high latitude glendonites with eustasy and show that glendonites were a proxy for extensive polar ice.[56] His study was done only on a gross scale, however, and it cannot be used to rigorously test the proposed correlation. Rigorous correlations require high-resolution stratigraphy of glendonite-bearing sediments to determine not only their precise age, but also their duration in the geologic record.

Our work is aimed at deriving this high-resolution stratigraphy utilizing a technique called magnetostratigraphy (see below). We supplement the magnetostratigraphy with ammonite and *Buchia* biostratrigraphy to calibrate our paleomagnetic data to the geologic time scale. Such calibration allows us to investigate the duration of glendonite-bearing horizons, the relationship between glendonites, OAEs, and eustasy, and the potential mechanisms that may have caused Early Cretaceous cooling events.

From Field to Lab: Sample Collection

Our field work in 1999 and 2000 concentrated on the sediments of the Deer Bay Formation, a relatively homogenous sediment package of dark silty shallow marine shales and siderite concretionary beds deposited and precipitated during the Early Cretaceous.[45] The Deer Bay Formation is bound by two sandstone units: the Jurassic Awingak Formation below and the Isachsen Formation above.[21] These formations are part of a continuous sequence of Mesozoic (250-65 Ma) and Cenozoic rocks of the Sverdrup Basin of Arctic Canada [Fig. 1a-b], and were brought to the surface by tectonic movements between North America and Greenland.[16,33,50,65,69]

Samples were collected from the steeply dipping (~65°) south limb of a synform (a trough-shaped structure) in the Blackwelder Mountains of west-central Ellesmere Island (80° 38.85' N, 85° 00.79' W) [Fig. 1b-c]. Approximately 200 samples from 36 sites spanning 230 meters of the Deer Bay Formation were collected as either one-inch field-drilled cores or, where weathering or lithology precluded coring, hand samples. Because the goal of our paleomagnetic study is to measure the direction of the magnetic field preserved in the rocks (hopefully at the time of deposition), we measure the unique orientation of each sample before removing it from the outcrop.

Orienting high latitude samples is not straightforward, however. Our field area is very near the present north magnetic pole (78.5° N, 103.4° W) [Fig. 1b], so the lines of force that describe the direction of Earth's magnetic field are nearly vertical there.[53] Consequently, any magnetic disturbances due to external forces (*e.g.*, solar wind) can cause large errors in magnetic compass readings. To account for these complications, we oriented our samples with a sun compass and compared any measurements made with a magnetic compass to daily records of the present magnetic field measured at various high latitude magnetic observatories.[65]

In addition to drilling core samples for paleomagnetic studies, we described the physical characteristics of the rock at each core site and from at least one site between two successive core sites. Representative fossils, shale samples, and geologic specimens (*e.g.*, fossil wood, glendonites) were collected from each sample site for age control and geochemical analysis. Fossil animals were found throughout the section and included gastropods (snails), bivalves (clams), a few echinoderms (sea lilies), belemnites, and ammonites (hard-shelled cephalopods). Although there are numerous forms of life represented, these forms are not diverse. Jeletzky comments that, in general, North American Boreal fauna lack diversity and are dominated by a few species with abundant representatives.[39] Our data, in addition to data collected by previous researchers, support Jeletzky's observations and are suggestive of a cold marine environment at the time the Deer Bay Formation was deposited.

Glendonites were preserved in three distinct horizons in the upper half of the section. The first horizon is the thickest (~15m) and occurs directly above a thin (~5m) band of reddish shale. The other two horizons are considerably thinner and are separated from the initial horizon by approximately 17 meters. A large slump covers the section immediately above the last horizon. The entire section spanned by the glendonite horizons is generally devoid of fossils. The same stratigraphic, lithological, and taphonomic (fossil preservation) patterns are characteristic at Expedition Fiord on western Axel Heiberg Island, several hundred kilometers to the west [Fig. 1b]. The only significant difference between these two sites is that the glendonites at Expedition Fiord are much smaller and more weathered than the glendonites at the Blackwelder Mountains.

Fossils and age control

William Smith, the now famous English civil engineer and surveyor, observed that fossils, especially fossil assemblages, succeed one another in a regular and determinable order. The principle of faunal succession was very controversial when Smith proposed it, but it is now a fundamental part of historical geology and the geologic timescale: a species originates, persists for some time, and then becomes extinct, never to reappear in the geologic record. By carefully documenting the distribution of fossils preserved in layers of rock, these layers can be subdivided into time units based on the unique fossil assemblages they contain, a technique called biostratigraphy. Although geologic ages determined by biostratigraphy were originally relative, developments in geochronology, such as radiometric age-dating, supplement the fossil data, and biostratigraphy is now a very useful tool for determining both the relative and absolute ages of rocks.

We collected numerous invertebrate fossil organisms for age control to calibrate our paleomagnetic data to the geologic time scale. Belemnites are extinct straight-shelled cephalopods that were abundant in the Jurassic and Early Cretaceous. Several poorly preserved specimens were collected throughout the section, and although the poor preservation made identification difficult and more speculative than definitive, the specimens appear to be forms characteristic of the Valanginian in the Early Cretaceous.

William Elder, a paleontologist with the National Park Service, graciously identified 22 representative *Buchia* specimens collected throughout the section [Fig. 3a-b]. *Buchia*, robust bivalves similar in gross morphology to clams, are by far the most abundant macrofossil in the Deer Bay Formation. Elder identified several forms that are characteristic of the Boreal Valanginian.[74] He remarks that the distribution of specific forms in the uppermost part of the section suggest that this part of the Deer Bay Formation records the gradational boundary from the Lower Valanginian to the Upper Valanginian.

Ammonites, extinct cephalopods similar in appearance to the Modern Nautilus, are by far the most useful index fossil for Mesozoic biostratigraphy. In fact, the Mesozoic is divided into ever-shorter time units based in part on ammonite speciation patterns. We were very fortunate to find several well-preserved ammonite specimens [Fig. 3c], all of which correlate to the uppermost Lower Valanginian.[27,30,32]

Magneto-what?

Although biostratigraphy is an invaluable first to second order technique for calibrating an outcrop to the geologic timescale, its time resolution can be limited. For example, we can deduce from ammonite biostratigraphy that the Deer Bay Formation was deposited during the Valanginian stage of the Early Cretaceous, but we cannot determine more precisely when glendonites appeared in the Deer Bay Formation and for how long they persisted. Thus, we would like to supplement the biostratigraphy with an independent dating method that yields higher time resolution. Although it probably is not immediately obvious, paleomagnetism— the study of Earth's ancient magnetic field as preserved in rocks and sediments— can (ideally) establish the high resolution stratigraphy we need to better understand the nature of glendonite-bearing horizons in the Deer Bay Formation.

The time-averaged (over 1000s of years) magnetic field of Earth (geomagnetic field) resembles a magnetic dipole oriented along the Earth's spin axis, an observation referred to as the Geocentric Axial Dipole (GAD) hypothesis. The magnetic field is vectorial and is described by both its direction and magnitude, which vary with latitude. The vertical component of the direction of the magnetic field is called *inclination* and is a fundamental direction in paleomagnetic studies. Inclination is vertical and up (that is, pointing out from Earth) at the south magnetic pole, parallel to the surface of the Earth at the magnetic equator, and vertical and down (pointing into the Earth) at the north magnetic pole. The current configuration of the geomagnetic field, which we call *normal polarity*, is such that the north magnetic pole is in the northern hemisphere (and the south magnetic pole is in the southern hemisphere); hence, the north-seeking arrow of a magnetic compass points geographically north. Another fundamental observation from paleomagnetism is that the direction (and magnitude) of the Earth's magnetic field has changed randomly with time such that at some points in geologic time the north magnetic pole was actually in the southern hemisphere. We call this configuration of Earth's magnetic field *reversed polarity*.

The random pattern of field reversals through geologic time became apparent when geophysicists measured the magnetic directions preserved in young, well-dated volcanic rocks on land and later from volcanic rocks on the ocean floor. Oceanic crust is produced at spreading ridges where hot magma erupts and cools to form a rock called basalt, which often records very well the direction of the magnetic field at the time it cools. The creation of oceanic crust at spreading ridges is a continuous process, and since crust moves away from spreading ridges, measuring the preserved directions of the magnetic field in oceanic basalts from a spreading ridge to another tectonic plate boundary (such as a subduction zone, where crust is destroyed as it sinks deep into the Earth) provides a continuous record of magnetic field reversals from the present back to approximately 160 Ma (the age of the oldest oceanic crust). These series of measurements through geologic time are precisely dated (usually using radioactive age

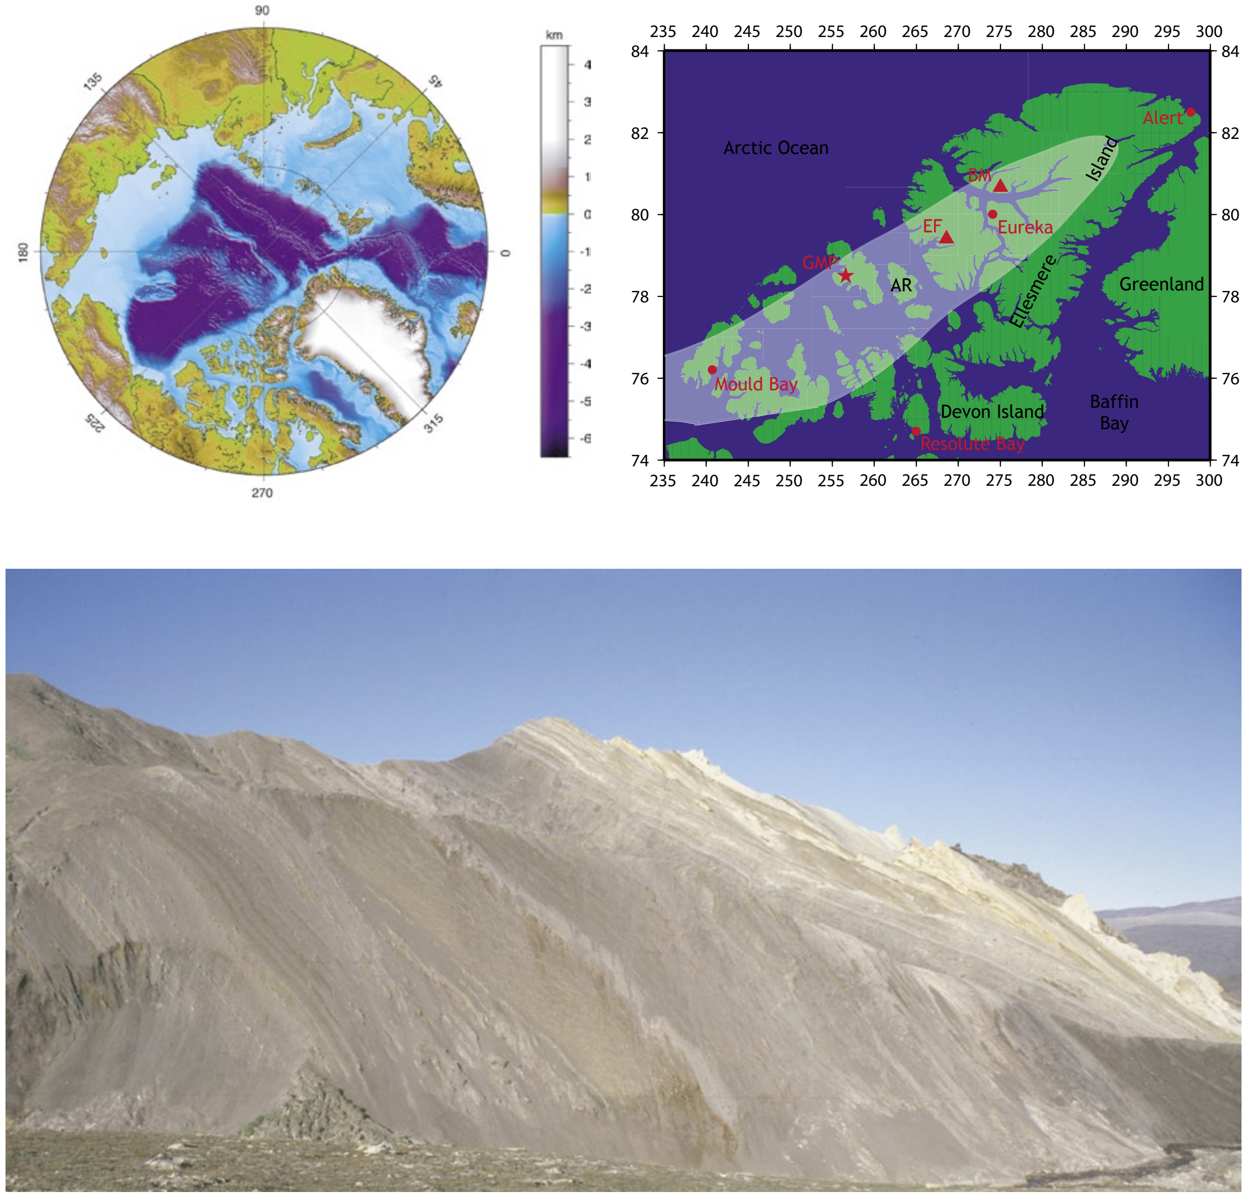

Figure 1: a. Arctic geogrpahy and bathymetry from Cherkis *et al.*[11] b. The Queen Elizabeth Islands, Arctic Canada; GMP: geomagnetic pole; BM: Blackwelder Mountains, the site of this study; EF: Expedition Fiord; AR: Amund Ringes Island; area shaded in gray is approximate extent of the Sverdrup Basin. c. Stratigraphic section discussed in this paper; glendonite-bearing horizons are immediately right of the yellow and red horizons in the middle of the photograph. North is toward the right.

dating techniques) and compiled into what is referred to as the Geomagnetic Polarity Timescale (GPTS). A period of reversed polarity and its immediately proceeding normal period is called a 'chron,' which is characteristic to a point geologic time, much like a fossil can be specific to a point in geologic time. Because the magnetic field reverses randomly, the *pattern* of chrons is unique throughout time. Furthermore, because the magnetic field encompasses the entire planet, all rocks will record the same pattern of reversals at a specific point in time. Geologists can measure the magnetic directions preserved in their sampled rocks and, using age tie points (*e.g.*, fossils, radioactive age dates), can calibrate their measured pattern of field reversals to the GPTS; this dating technique is called magnetostratigraphy. Magnetostratigraphy is often more precise and often offers far higher resolution than other types of stratigraphy (*e.g.*, biostratigraphy, lithostratigraphy), enabling geologists to have better age control over the rocks they are studying.

Extracting the magnetic direction preserved in a rock sample is non-trivial, however. Samples are gradually demagnetized in a step-wise fashion using a variety of techniques, such as applying a known magnetic field to a sample or heating the sample to a specified temperature, all of which try to isolate the most stable, characteristic magnetization preserved in the rock. Between each demagnetization step, samples are measured in a magnetometer that can measure very magnetic fields. The demagnetization data from each sample is then compiled in graphical form, from which the vector of the magnetic field preserved in each sample (and therefore the direction of the magnetic field in that sample) can be determined using a variety of statistical techniques.

Many different factors can complicate or overprint the primary magnetization of a sample some time after the rock formed, such that

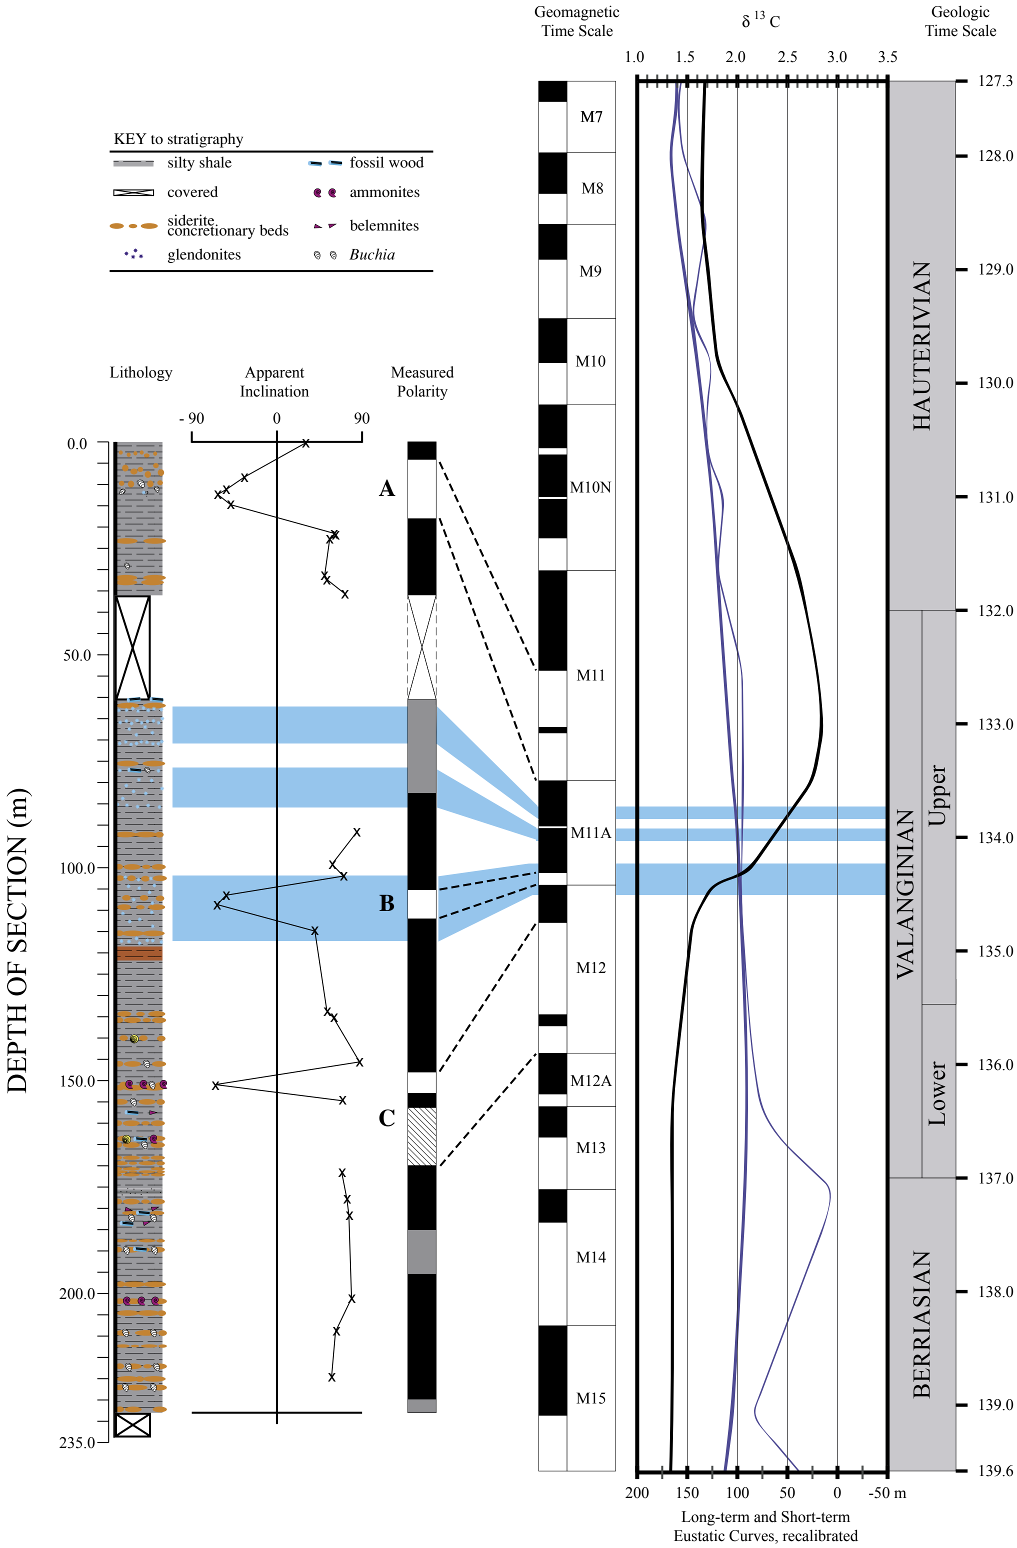

Figure 2: Calibration of the Deer Bay Formation, Blackwelder Mountains, to the geomagnetic time scale,[54] geologic time scale,[27,32] the global sea level curve (blue curves; thicker line is long-term trend; thinner line is the short-term trend),[31] and the $\delta^{13}C$ curve (black curve)[10]. Black areas in the measured polarity column correspond to sites with normal polarity, while white areas correspond to sites with reversed polarity; gray areas correspond to undetermined site-mean directions. The glendonite-bearing horizons (light blue) appear to represent short periods of time within a span of one million years; furthermore, they appear to be coeval with the onset of the C-isotope excursion.

the characteristic (*i.e.*, most stable) magnetization may not be the primary magnetization. Fortunately, we can use standard paleomagnetic field tests to determine the relative time at which a rock was magnetized. The most straightforward field test is a tilt-test. The beds from which we collected our samples are presently steeply dipping, but were horizontal when they were deposited. If the magnetization is Recent, then the inclination will be shallow when we restore the beds to horizontal. Conversely, if the magnetization is ancient, inclination will steepen once the beds are restored to horizontal. Most samples have steep inclinations after tectonic correction, which signify that ancient primary magnetizations have been isolated and preserved. These results indicate that we can construct a magnetic polarity column.

Several sites showed clear normal polarity directions. There were a few examples of clear reversed polarity, evidenced by shallow *in situ* directions that became steep after tectonic correction. The sites from which these reversed directions were observed are stratigraphically distinct, which suggests that these sites record reversed polarity zones. Rarely did all samples from such sites yield clear reversed directions. More commonly, a few samples showed clear reversed polarity directions after overprints had been removed, while other samples appeared to show a progression of increasingly greater degrees of overprinting. The majority of the overprints seemed to record a Recent field direction, which is consistent with a phenomena known as viscous remanent magnetization.[9] In some cases, however, the overprint appeared to have a steep, pre-tilting direction similar to the normal polarity characteristic magnetization isolated at most of our sites. In one zone, the mixing of reversed and normal polarity pre-tilting directions is sufficiently complex that we can only assign the interval as mixed polarity.

To summarize our directions and construct a polarity column, we first compiled all sites with normal polarity directions and calculated site-mean directions using spherical statistical methods.[23] Site-mean data are compared to various expected

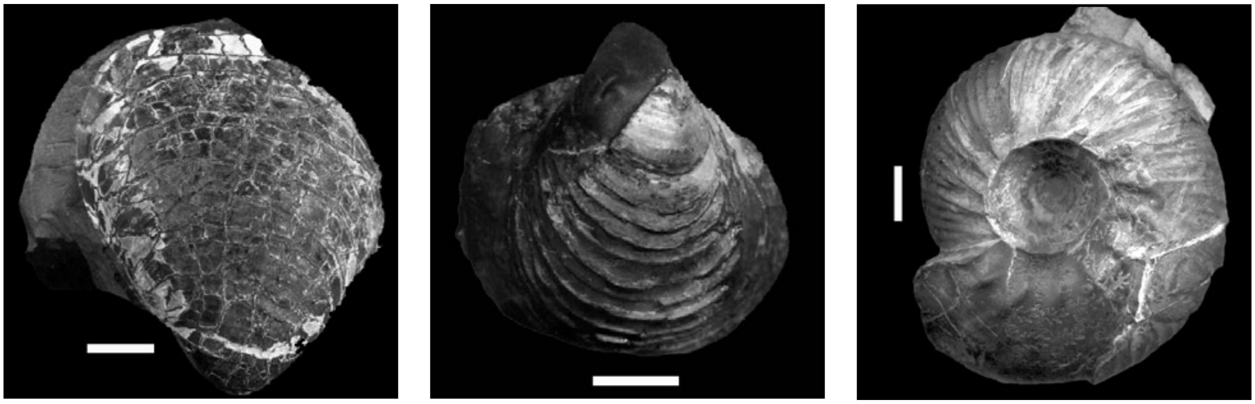

Figure 3: Some of the key fossils from the Deer Bay Formation used as tie-points to the geologic timescale. a. *Buchia keyserlingi* (Lahusen), a bivalve characteristic to the Lower Valanginian (approximately 137 to 135.5 Ma); b. *Buchia sublaevis-inflata*, characteristic to the Upper Valanginian (approximately 135.5 to 132 Ma); c. a *Polyptychites keyserlingi* ammonite, distinct to the uppermost Lower Valanginian. See Figure 2 for timescale.

directions for the Early Cretaceous, which provides another test for primary magnetizations.[6,64] Normal polarity paleomagnetic results from the Deer Bay Formation are consistent with expected directions for this region of the Arctic in the Early Cretaceous, which supports the results from our field tests.[68]

Getting good dates

Our preliminary polarity column contains three reversed polarity intervals (called A, B, and C here). Zones A and B are defined by clear reversed polarity samples. Zone C, however, has only a single site with straightforward reversed polarity directions; all other sites in this zone have both normal and reversed polarity directions. Our polarity interpretation is based on the following tie points. First, our biostratigraphy places the base of our section in the Lower Valanginian, close to the Upper/Lower Valanginian boundary (~136 Ma) [Fig. 2].[27,30,32] Second, the top of the section is at the base of the Isachsen Formation, which is thought to represent the Hautervian/Valanginian boundary.[22,48] Using these tie points, we assign the uppermost reversed polarity interval (Zone A) to reversed polarity Chron M11 of the GPTS [Fig. 2]. Thus Interval B could be a record of one of the reversed polarity subchrons in M11A, or a record of Chron M12. However, the paleontological identifications suggest that the mixed polarity zone C corresponds to chron M12 (further suggesting that Interval B is a subchron in M11A).

Putting it all together

Preliminary calibration of the magnetostratigraphy of the Deer Bay Formation to the GPTS provides new information about the timing and duration of fine-grained sedimentation and glendonite-bearing horizons in the High Canadian Arctic. Biostratigraphic correlation unambiguously places the Deer Bay Formation in the Valanginian Stage; magnetostratigraphic correlation suggests that the sampled section is greatly expanded (*i.e.*, more sediment per unit time) compared to typical low-latitude successions. Previous authors have only speculated on the duration of glendonite-bearing horizons.[48] We suggest that all three discrete zones of glendonites occur during the upper-most interval of normal polarity in Chron M12 and throughout Chron M11A (a span of less than one million years) [Fig. 2]. Individual glendonite zones, therefore, probably represent short periods of geologic time: *e.g.*, ~100 kyr for the upper two horizons, and ~400 kyr for the lowest horizon. The implication is that polar climate crossed a temperature threshold, below which ikaite formation was facilitated, at least three times in a period as short as one million years.

Sea level and climate change

Price, in a review of previous work on glendonites, emphasized that Valanginian glendonites represent a global climate signal because they temporally correspond to low eustasy.[56] Tethyan and Lower Boreal stratigraphic records indicate that global sea level was at a Late Mesozoic minimum during the Valanginian [Fig. 2], after which sea level reached a global high throughout the Cretaceous and Early Tertiary. Eustasy does not reach similarly low levels until the Middle Miocene, 120 million years later.[31,32]

Glendonite-bearing horizons do temporally correlate well to stratigraphic records [Fig. 2], indicating that polar glendonites may record a global climate signal. Terrestrial ice is one possible mechanism for low eustasy, especially for rapid and short-term (<3 myr) eustatic changes. We make a cautionary note in regard to relying heavily on rapid and short-term sea level change to invoke climate changes: several authors have questioned the reliable recognition of third-order and lower sea level changes.[13,26]

The biologic response

We can use our magnetostratigraphic column to correlate the Arctic Sverdrup Basin to lower latitude sections from which Early Cretaceous stable isotope stratigraphies have been measured. In particular, we note that Channell, Erba, and Lini have constrained a Late Valanginian C-isotope excursion with magnetostratigraphy in Tethyan limestone sections [Fig. 2].[10] $\delta^{13}C$ values initially rise during Chron M12, peak during the normal period of Chron M11, and slowly decay to pre-excursion values from Chrons M10N to M8.[10] This positive excursion, and the occurrence of numerous black shale horizons, suggests high biological productivity in the oceans and a global oceanic anoxic event.[51,73] The global nature of this event has been bolstered by recent drilling results from the northwestern Pacific Ocean.[7]

The upper Deer Bay Formation corresponds in time (as defined by our magnetostratigraphy) to the Valanginian OAE defined by Lini *et al.*[51] Although organic carbon values from our section are not yet available, we note that the overall stratigraphy of the Deer Bay Formation is consistent with enhanced organic carbon burial.[48] Importantly, we note that, on a finer scale, the glendonite-bearing horizons from the Deer Bay Formation appear to be contemporaneous with the onset of the carbon isotope excursion [Fig. 2].

Weissert *et al.* related the $\delta^{13}C$ excursion to peak Paraná (Brazil) volcanic activity and rifting in the North Atlantic.[73] The Valanginian C-isotope excursion is increasingly considered to signify the onset of the overall trend toward greenhouse conditions during the Cretaceous.[51] In this sense, the correspondence with the glendonite events, which suggest cooling, at first seems paradoxical. We note, however, some similarities with our observations in the Valanginian with developing work on Cenomanian-Turonian sections (during another, better preserved OAE). In particular, a number of authors have noted short-term (100 kyr and less) variations within the positive $\delta^{13}C$ excursion that they relate to polar ice and/or sea level falls.[25,62,71,72] Although the evidence for polar ice is hotly debated, there does appear to be a sufficient stable isotope event (both in carbon and oxygen isotopes) to indicate relatively short-term cooling and significant changes in the global carbon cycle within the broader positive Cenomanian-Turonian C-isotope excursion.[37]

Cooling in the greenhouse?

The traditional view of positive carbon isotope excursions is that they represent major changes in the global carbon cycle during especially intense greenhouse climate conditions. Our correlations suggest that the Valanginian event was punctuated by rapid cooling events. Mild greenhouse conditions may have increased nutrient delivery to the oceans, which stimulated oceanic primary production. The consequent reduction in atmospheric CO_2 (from a vigorous carbon pump; *e.g.,* agal photosynthesis) could have produced a small, antagonistic response to the mild greenhouse conditions and either stabilized or even cooled surface atmosphere temperatures. This response, however, may not have been significant by itself to produce cooling; rather, two or more climate variables would probably need to work in concert to push the climate system to a critical threshold.

For example, the orientation of the Earth's orbit around the sun is not constant over time. Rather, the shape of the orbit (eccentricity), the inclination of the Earth's spin axis with respect to its plane of orbit (obliquity), and the wobble of the spin axis due to the rotation of the Earth (precession) change cyclically with 100 kyr, 41 kyr, and 23 kyr periodicities, respectively. These orbital cycles, called Milankovitch cycles, affect the amount of solar radiation that reaches Earth, and thus affect Earth's climate. Our temporal control on the glendonites, although a great improvement over previous results, is still limited. Nevertheless, there is a hint of orbital forcing because two of the glendonite horizons appear to represent 100 kyr intervals, and the third a 400 kyr interval. Thus, eccentricity may have amplified the effect of small biologically-induced reductions in atmospheric CO_2 and briefly pushed the climate system over a critical threshold.

Summary

The Deer Bay Formation in the Sverdrup Basin of Arctic Canada contains three discrete glendonite-bearing horizons suggestive of cold water environments. Previous studies and reviews have tried to correlate the occurrence of glendonite-bearing horizons to episodes of low global sea level, aiming to infer global climate signals from these correlations; but these reviews have been limited to gross temporal correlations. The primary goal of this study was to sample the Deer Bay Formation and to measure and determine a magnetostratigraphic polarity column with which we could make detailed temporal correlations. Primary magnetizations are preserved by the Deer Bay Formation and allow us to construct a preliminary magnetostratigraphic polarity column. The faunal assemblage constrains the lower portion of our measured section to the uppermost Lower Valanginian (136 Ma).[27] Our temporal correlations provide the first refined age control on the duration of these high latitude glendonite-bearing horizons. We suggest that the three distinct glendonite-bearing horizons in the Deer Bay Formation represent short (100-400 kyr) cooling events. Also for the first time, Arctic stratigraphy can be readily correlated to low latitude records. Our results suggest that glendonite-bearing horizons are coincident with a Late Mesozoic eustatic low. More surprisingly, however, is our discovery that glendonite-bearing horizons are contemporaneous with the onset of the Valanginian OAE as defined by C-isotope stratigraphy in Tethyan sections. Additional work to refine the magnetostratigraphy for the Deer Bay Formation, as well as efforts to measure directly a C-isotope stratigraphy in our section, will refine our correlations. Furthermore, our conclusions need to be tested against other hypotheses for the Early Valanginian C-isotope event and the validity of glendonites as paleoclimate indicators. The available data, however, are compatible with the idea that glendonite-bearing horizons in the Deer Bay formation represent short reversals of climate during what is generally recognized as the start of Cretaceous greenhouse conditions.

Peter Lippert graduated from the University in 2003 with an Honors B.S. in Geosciences and a Take 5 Fellowship in the anthropology and social history of modern science and medicine.

This article is a summary of his undergraduate thesis work with the Paleomagnetic Research Group, which is based on field work he did in the Arctic in 1999 and 2000. Peter is currently at the University of California, Santa Cruz, studying the uplift history of the Tibetan Plateau and associated Asian tectonics.

Members of the Class of 2003 who received distinction in research

Biological Sciences: Biochemistry
Chad A Galloway
Lisa Johns
Rebecca Kathryn Montange

Biological Sciences: Ecology And Evolutionary Biology
Katherine M. Lander

Biological Sciences: Microbiology
Andrew David Berti
Hetal Mahendra Patel
Stephen Joseph Salipante

Biological Sciences: Neuroscience
Matthew Aaron Tremblay

Brain And Cognitive Sciences
Matthew Lee Hall
David William Mccolgin

Geological Sciences
Kathleen Elizabeth Bailey
Peter C. Lippert **
Allyson O'kane **

Economics
Gary B. Deutsch
Laura Marie Kilday

English
Scott Douglas Challener
Charlotte Anne Cline
Elizabeth Lucy Conway
Candace Marie Curran
Suzanne Elizabeth Decker
Victoria Helen Goode
Timothy Matthew Green
David Chandler Howland
Paul Jerry Linczak
Sabrina Marie May
Katherine Marthe Mcgraw
James Anthony Messina
David Christopher Morris
Dianne Elia Patterson
Kali Michelle Quinn
Andrew J. Rosenthal
Alison Henstrand Schroth
Erez Abraham Solomon
Lindsay Reed Walton
Erin Kathryn Zahradnik
Arielle Myers Zibrak

Health And Society
Bridget Kathleen Cougevan
Michael Geoffrey Katz
Kathryn Lee Komaridis
Jonathan Michael Tan

History
Joshua Ross Blumenfeld
Corinne Marie Carpenter
Kathleen Bridget Casey
Michael Andrew Egolf
Emily Ruth Fletcher
Katherine Marie Kaminski
Letty Laskowski
Dave M. Rudy

Linguistics
Douglas L. Ball

Modern Languages and Cultures (MLC): MLC Comparative Literature
Jeffrey Clayton Dubois **

MLC French
Jessica Denise Burt **
Francine P. Chow **
Elizabeth Aphrodite Morphis **
Lindsey Anne Wood **

MLC German
Kristin Ann Dickinson **
Michele Ann Krugh **

Music
Daniel Mark Dipaolo
Victoria Marie Saah *

Political Science
Kara Weeks DeSantis **
Sudhir N. Shenoy *

Psychology
Michael Seidman Axler *
Suzanne Elizabeth Decker *
Dara Ilene Faden *
Susan Elise Klassen
Joan Elizabeth Knihnicki **
Grace Kong *
Rachel Elizabeth Morrisseyott *
Christopher Peter Niemiec **
Arathi Rajendran
Micheal Joseph Reiter *
Amy E. Rodrigues *
Kristyn Ann Saveliev
Flint Mark Edward Spitler *
Teresa Mary Tygert **
Donald Anthony Young

Religion
Scott Douglas Challener
Paul Marie Jacques Gelbard
William E. Moore
Katrina Susan O'Neil
Sarah Anne Otto
David Michael Schneider

Note: A single asterisk (*) denotes high honors in research; a double asterisk (**) denotes highest honors in research.

These names were provided as per the Registrar's office. Departments not listed acknowledge outstanding scholarship via alternate methods.

Undergraduate Research Programs at the University of Rochester

A selection of programs with comments by recent participants.

For complete program listings please visit our website at jur.rochester.edu.

The Camelot Project

The program is dedicated to the development of an online database of Arthurian texts, images, bibliographies, and basic information. Undergraduates may intern in the program for credit, performing their own research, writing papers, and developing their own websites as a culmination of their work.
For more information on the Camelot Project, contact Alan Lupack in the Robbins Library at alupack@library.rochester.edu.

Teresa Lopez:

"As a Camelot scholar, I learned a lot about research methods. This came in handy this year as I am in the honors program for English majors and it was much easier to complete my research with some knowledge of where to start. More so, the Camelot project increased my desire to study Arthurian Literature in graduate school.

The Camelot Project also sparked an interest in continuing research in English literature (mostly Arthurian). I had planned to attend graduate school for literature, but I had not decided on what to focus on with a master's degree or PhD. Now I plan on working towards becoming a professor and doing my own research in Arthurian Literature.

The most intriguing about the Camelot Project are the people involved in its creation and upkeep. Professor Alan Lupack, Professor Russell Peck, and Rosemary Paprocki are just a few of the many people that keep the project going and their enthusiasm is infectious. They truly made the project a positive addition to my academic career at the University of Rochester."

Susan Bauer
"Tom Thumb in the Arthurian Tradition"
Advisor: Alan Lupack

Teresa Lopez
"Uther and Igraine"
Advisor: Alan Lupack

Robyn Pollack
"Merlin and Vivien"
Advisor: Alan Lupack

de Kiewiet Summer Research

The de Kiewiet Summer Research Fellowships are designed to give University of Rochester students in the B.S. and B.A. Tracks in the Program in Biology and Medicine an opportunity to obtain substantial research experience in the laboratories of Program Faculty. Students in their junior year are invited to submit competitive applications. For information of de Kiewiet Summer Research, contact Doris Kist in the Department of Biology at djck@mail.rochester.edu.

Kelly Wentworth:

"The de Kiewiet program provided me with the opportunity to engage in my first research project. It has helped me to develop important and useful laboratory skills and techniques, and has provided me with the opportunity to apply the material learned in the classroom to the laboratory.

I am heading off to medical school next year, and I feel that the exposure that I have gained through the de Kiewiet program has increased my interest in research and heightened my awareness of the importance of research in the medical field.

It was a Great opportunity to participate in a lab of your choice and receive excellent support and guidance from the faculty."

Evan Kingsley
"Retinal pathology of the $Cln3^{-/-}$ mouse, a model for Batten Disease"
Advisor: David Pearce, Biochemistry and Biophysics

Matthew Maurer
"Protein factors involved in mismatch repair of the mitochondrial genome in Saccharomyces cerevisiae"
Advisor: Elaine Sia, Biology

Mark O'Hara
"Analysis of TWI4: a homologue of TWI1"
Advisor: Martin Gorovsky, Biology

Niraj Patel
"MTS1 binding research"
Advisor: Ravi Basavappa, Biochemistry and Biophysics

Rebecca Porter
"Genomic screening of Saccharomyces cerevisiae for pre-mRNA splicing factors"
Biochemistry and Biophysics
Advisor: Yi-Tao Yu

Anne Stey
"Excitotoxic intracellular and mitochondrial calcium regulation"
Advisor: Shey-Shing Sheu, Pharmacology and Physiology

Kelly Wentworth
"Identification of the interacting proteins of the protein kinase PKK by a yeast two-hybrid system"
Advisor: Luojing Chen, Center for Human Genetics and Pediatric Disease

Cornelia Zorca
"Attempts to generate a germline HHP1 knockout in Tetrahymena thermophila"
Advisor: Martin Gorovsky, Biology

The McNair Program

The two year program is comprised of a series of colloquia designed to help students identify appropriate graduate programs, gain admission to them, and secure financial support. Participants complete a series of assignments, including writing the personal statement; the program culminates with each scholar completing a semester-long mentored research project. Other program supports include intensive writing assistance, GRE Preparation, graduate school visits, study group sessions, and help locating summer research opportunities. A stipend of $1000 is given to each Scholar upon completion of his or her research project.

The Summer Research Scholars spend June and July at the University of Rochester conducting research under the guidance of a faculty mentor in fields that can lead to the Ph.D. Scholars prepare intensively for the Graduate Record Exam and take a course entitled "The Culture of the Academy". Students present the results of their research projects at a university-sponsored conference, and have the opportunity to present at disciplinary conferences. Each participant in the Summer Research Program receives a $2,800 stipend, free room and board on campus, and travel expenses. For more information contact Donna Simmons at mcnair@cc.rochester.edu.

Rene Herbert:

"The McNair program gave me the opportunity to further explore my career interests and the independence of working on a project that I developed. I not only gained research experience, but I also gain the opportunity to work with professors outside of the River Campus on my individual research. The support of the McNair program and the interest of my mentor during the program is a valuable tool of the McNair program.

This research experience has made me love what I wanted to do even more and gave me motivation to seek additional research opportunities for the future. I was able to find more research opportunities for the Spring 2004 semester as well as work with my mentor during the academic year. It has also given me an opportunity to present my research at various research programs and conferences, which is a good opportunity for graduate schools to recruit you based on your research.

The McNair program is a great opportunity for anyone interested in continuing or starting research in undergraduate education in a any field. The opportunities available with McNair has an impact on the decisions you make in the future and also give you research opportunities with you mentor and to present your research all over the United States."

Dan Chen
"Activity Level Measurements in Clinical Depression"
Advisor: Craig Harman, Biomedical Engineering

Ling Dong
"Efficient Modeling of a Knee Joint from MRI"
Advisor: Amy Lerner, Biomedical Engineering

LaShara Evans
"Controversy and Confusion; Hans Selye's Stress Model in Relation to Disease: 1945-1975"
Advisor: Ted Brown, Community & Preventive Medicine

Naomi Extra
"The Cultural Representation of Difference Amongst African American Women in Film"
Advisor: Sharon Willis, Film Studies

René Herbert
"Who Uses the Health Start Center and Why?"
Advisor: Ann Dozier, Community & Preventive Medicine Healthy Start Center Staff

Avril Little
"How Are Depressive Symptoms and Anxious Moods Related to Preparation for Future Care in Older Adults?"
Advisor: Silvia Sorenson, Psychiatry

Meng Lu
"Measuring and Analyzing the Wavefront Aberration in Cats"
Advisor: Krystel Huxlin, Opthamology

Emily Medina
"'I Can Hear You Talking But I'm not Listening:' Communication Between Physicians and Patients - A Biopsychosocial Approach"
Advisor: Ted Brown, Community & Preventive Medicine

Olga Medina
"Does the Purported Vital Role that Anti-Personnel Mines Play in U.S. Military Strategy Make it Acceptable to Violate Humanitarian Rights?"
Advisor: James Johnson, Political Science

Tarik Smith
"Behavior and Biochemistry Explorations in a Mouse Model of Battens Disease"
Advisors: David Pearce and Elizabeth Kriscenski-Perry, Developmental Biology

Richard Tipton
"The Michigan Cases: The Constitutionality of Race-Preference in College Admissions Programs"
Advisor: Tyll Van Geel, Warner School of Education & Human Development

Keron Twum
"Qualitative Research Study on Obesity and Exercise"
Advisor: Carolyn Tabak, Pediatrics

Brian Warsop
"The Analysis and Characterization of Nickel (II) Metal Complexes For Carbon-Oxygen Bond Activation"
Advisor: Bill Jones, Chemistry

Picture References

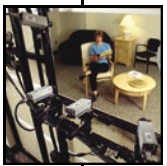

All images from interview credited to University of Rochester Public Relations

All images from "Sonic Architecture and The Creation of Space" credited to Hindu India: From Khajuraho to The Temple City of Madurai by Henri Stierlin, Photos by Anne and Henri Stierlin. Taschen Books, 1998

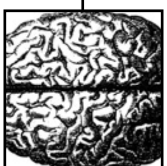

Opening image to "Short-Term Memory In Secondy Language Learners of ASL" credited to the cover of The Body Within by Esther M. Sternberg, MD.

Note: Citations for all articles are available online at http://jur.rochester.edu.

About the Journal

The Journal of Undergraduate Research (ISSN# 1547-9641), Volume 2, Issue 2, Spring 2004 was assembed on a Apple PowerMac G4 Quicksilver using Adobe InDesign 2.0. Graphics were produced using Adobe Photoshop 7.0 and Adobe Illustrator 10. Microsoft Office was used for text editing and review. Fonts used included Adobe Garamond, Futura, Times New Roman, and Symbol (for Greek characters). This 52-page journal was perfectly bound by Merlin Digital Solutions Group of Rochester, NY.