The goal of our presentations today is to give you information regarding the current computational resources available to UCSB researchers here on campus and at the National Super Computing Centers. After laying out the basics we hope that interested folks will engage us to improve the usefulness and accessibility of these and future resources. It's also quite probable that there are other resources available to you. I doubt I've investigated all of the options.

First off- We're here to help you. While I am busy, and I'm sure Stephanie and Burak are too, our groups will try to help you with your computational research needs where we can.

Second- This presentation is more about opening a dialog with interested researchers to identify appropriate computational systems. You may notice significant gaps between our areas of expertise and what you require. If you don't find what you're looking for here today, email us with those concerns if you don't know who else to contact. We might be able to put you in touch with groups that can help.

Third- Know your data and your research and what I'll call the interpersonal relationships your data is exploring. Please be aware of UCs Human Subject's policies, where to go for help, and use the Secure Compute Research Environment if necessary.

Fourth- The Center for Scientific Computing has some pretty neat computers and software.  I'll lay out a few of the things they can do and what they're strengths (and perhaps some weaknesses) are.

Finally I'll lay out some future projects which may be beneficial to your research.
After that Burak will briefly talk about UCSB's interaction with the National SuperComputing Centers and how you access those centers.
And then Stephanie will introduce the Library's data labs and licenses.

Who We Are

Let me begin with the least exciting portion of my talk - Human Subjects Data Considerations.  In researching ideas for this talk I contacted the Human Subjects department of the Office of Research and met with Melissa Warren and Dorin Donohoe.  While a complete review of data considerations is way beyond the scope of this talk, I hope you leave with a couple of points- the most important being- when in doubt, contact Melissa and Dorin.  They are a valuable resource to help mitigate university risk.
Reading the definition of Human Subjects from the Code of Federal Regulations I was struck by how it is so broad

as to include basically all social media data. Melissa informed me though that this does not mean consent is necessary in many cases. As an example she forwarded an email she sent to a professor interested in using Twitter data for a project last year.

She replied to the professor - 'think it's safe to say Twitter is public information and permissions should not be needed when gathering data'. Twitter's Terms and Conditions explicitly state "Most of the information you provide us is information you are asking us to make public." Cool- however then she goes on to list numerous articles regarding other ethical concerns. It's quite a list- enough to give any researcher pause.

She continues her email… with the main recommendations being "Anonymize" the data when possible and

determine the level of harm that may be caused by using the data if the Tweeter was identified. These follow directly from

the Belmont Report's 3 ethical principles - Recognition of the personal dignity and autonomy of individuals and special protection of persons with diminished autonomy,

Beneficence - obligation to protect persons from harm by maximizing anticipated benefits and minimizing possible risks of harm, and Justice - obligation to ensure that the benefits and burdens of research are distributed fairly.

Two quick anecdotes about the internet of today. Currently corporations are collecting tremendous amounts of personal data from users. Every one of these users have 'Agreed' to the Terms and Conditions of those corporations, often including clauses allowing research to be performed on their data. But this is a tense area at the moment. In 2012 an Uber Senior Vice President using Uber pickup and drop off location and time data showed the correlation between various dates and, as he put it, 'Rides of Glory' or 1 night stands- people returning to their pickup locations after 4-6 hours on a Friday or Saturday night. It took a while for folks to notice but in 2014 the Internet was not amused and Uber's data mining practices became an internet sensation of negative publicity.

Facebook, a company whose data UCSB researchers now have access to, and researchers from Cornell (one of whom subsequently moved to UCSF) manipulated approximately 700,000 users feeds to assess the effects on their emotions. From Forbes and the PNAS (Proceedings of the National Academy of Sciences) paper "The authors found that when they manipulated user timelines to reduce positive expressions displayed by others “people produced fewer positive posts and more negative posts; when negative expressions were reduced,

the opposite pattern occurred.”" Again once the Internet understood that a company and university researchers were manipulating them for research purposes, the outcry was substantial and Cornell (and then UCSF since one of the researchers was now PostDoc ing there) were treading on dangerous legal ground since as institutions receiving federal funding they are subject to the 'Common Rule for human subjects' that mandates informed consent.

Granted these are not the same as collecting Tweets, but being aware of risk and harm is important to the university.

So what should you do? Basically if you’re downloading data from publicly available internet sites, download the sites Terms and Conditions at the time of data collection. Terms and Conditions change constantly - having them available from the date of collection is good practice.
For sites that require logins or access such that users think they have an expectation of some privacy, you’ll definitely want to anonymize the data and I’d recommend contacting Melissa or Dorin.
And finally consider your research from the subject’s perspectives - minimize risk and harm.

If you are using sensitive data, undoubtedly you already know about the Secure Compute Research Environment service run by ETS, initiated by ISBER. Essentially the SCRE will create a virtualized workstation on a secured

private network with your data securely transferred in, and a set of commonly used analysis tools.  It uses multi-factor authentication and satisfies data security plan requirements of various funding agencies.

Enough scariness, on to the Center for Scientific Computing…

The Center for Scientific Computing came about through a collaboration between the Materials Research Laboratory and the California NanoSystems Institute in 2009.  When Elings Hall was built the CSC datacenter was the most up to date datacenter on campus, since surpassed by the North Hall Data Center.  By relocating compute clusters and other departmental research computing systems to the datacenter the CSC began freeing up departmental space, and it became a hub of activity for research computing.  However those departments without local specialized computing resources were left out of the ongoing collaboration of the local sysadmins.  Those departments that had computational resources benefitted from the interactions of the sysadmins towards common goals, while those not in the loop were completely left out. Is now the time to rectify that disparity?  I can't say- I'm just a High Performance Computing Systems administrator, but my thought is- yes, now is a great time to bring in social science researchers in to the realm of HPC utilizing "Big Data" in support of your research. Paul and I administer 5 clusters at the CSC, 2 of which are available to all campus researchers.  For those that don’t

know- clusters, originally called Beowulf Clusters, are groups of computers connected together with high-speed interconnects and high performance file-systems. The clusters we manage run in the 90-100% utilization and in a typical year over 20 million core-hours of compute time take place on our clusters. Here is a picture of Knot- with the 2U GPU nodes in the middle and the 1U compute nodes above. Typically Paul, myself, and many unlucky student workers put the spaghetti of cabling in place ourselves. Luckily with Knot, a collaboration with Hewlett Packard relieved us of that and so some unlucky folks at HP got stuck with it. I will say, they did a terrific job.

As mentioned, UCSB currently has 2 compute clusters available to researchers - QSR and Knot. QSR is more of a standard Beowulf Cluster while Knot has some interesting specializations that I'll get to momentarily. QSR is a much older cluster and Knot the result of an MRI grant from the NSF in 2011.
QSR is smaller- only 128 cores. Knot is much larger.

Let's begin with looking at how researchers use these clusters. The computations that run on the clusters basically fall in to 3 main categories- serial, node-based, and parallel MPI jobs.
Serial jobs are jobs that you could easily run on your laptop or local workstation. Where clusters become useful for serial jobs is when you want to run tons of jobs with sightly different variables across a domain, and perhaps

across several variable dimensions. With some relatively simple shell scripts you can submit tens or hundreds (even thousands, though please contact us first if you get that ambitious), of jobs that run across an entire parameter space on a given set of data. Whether the jobs take 1 hour, or 100 hours, our clusters are perfect for this sort of exploration.

MPICH (single-node) jobs are jobs that can take advantage of the numerous cores on a single compute node to talk to one another during the computation to make sure that boundary conditions, or other conditions, are consistent across the different core as they compute their tasks. The computations on the different cores talk to one another to make sure that the larger parameter space is consistent across the cores.

Multi-threaded (single-node) jobs take advantage of the various computational libraries available that can spread the work out across cores and then combine the results before proceeding. These computations are more of a single controller process that can dole out work to other cores when it encounters computational work that the multi-threaded libraries can distribute to worker cores.

And finally there are the traditional Parallel Message Passing Interface jobs that use the high-speed interconnects of the system to scale well beyond the number of cores on a node to crunch really large computations. Typically these require really complicated programming where many hours of work has gone in to determining where, exactly, the crunch points are in computations. I'll say that Burak knows a lot more about

this sort of programming than myself- I've exclusively dealt with vendor software that parallelizes as opposed to looking at parallelizing code myself.
As noted, the data you're working with has to be on the local high performance file system for the nodes to complete their jobs.

Knot has a couple of very interesting platforms for researchers. First there are the Fat nodes - nodes with much more RAM than the standard workstation. Once models or simulations exhaust the local RAM, performance drops by orders of magnitude (or they crash completely). With 1 TB of RAM, the size of arrays or the parameter space you wish to explore increases by perhaps even 2 orders of magnitude. If you're working with Agent Based Modeling, that's the differences between 10,000 agents and 1 million agents. And if those simulations can run for 10 or 30 or 60 days per run, you're exploring an enormous parameter space, whether you're spreading out threshold limits, agent numbers, or other factors in the agent's decision making.
Exploring computations from a different angle, NVIDIA GPUs offer a tremendous boost in compute power for simplified calculations. GPUs offer hundreds of processors with limited instruction sets and limited RAM. Code which takes advantage of GPUs can boost calculations by a factor of 20 easily. In addition to the MRL, I support the MC-CAM and they use GPUs for modeling the directed self-assembly of co-polymers for the

next-next-next generation of processors- in essence using Field theory to lay out polymers where traces between core interconnects will go.
Whoops- sorry- went off to left field there…
Do note though that the FLAME agent based modeling system is currently on Knot, and I am working on getting the GPU based FLAME modeling system active on the GPUs.
Lastly there are the Intel Phi nodes which are Intel's answer to the GPUs. They provide a much more standard instruction set than the GPUs so compilation of standard libraries should be simpler, but they are currently in their infancy and there is not much code out there to support them just yet.
The GPUs and the Phis have their own login node on Knot with a GPU and Phi on board to make compiling, testing, and debugging much simpler for the programmer.

To get access to the CSC clusters, visit our webpage and click on the Request User Account Request button. Knot and QSR are available to all campus researchers without charge.
We are a linux shop, and more specifically a CentOS Linux shop and our users predominantly interface with the clusters using the command line.
A quick aside- there are applications out there that can use web interfaces to create and submit jobs. At the moment the only one of these is in production on our clusters and it is used for BioInformatics (it's known as

Galaxy). We're not opposed to installing others if they will assist in your research and they're compatible with our clusters.
Jobs are a set of scripts that tell an application to process data, output results, and then terminate - all unattended.
Users transfer data from their local computers to the login node of the cluster using a secure file transfer agent like FileZilla, CyberDuck, scp, etc.
Here's an example of a simple job submission script that will run an R script using the Rmpi & Rsnow packages. It's asking for 16 cores - 4 nodes using 4 cores per node. It sets up the SNOW - simple network of workstations - using a single MPI directive. It outputs the nodes we're using to a file called nodes.

Here are the nodes I got for the calculation. 4 nodes, and the repeating node number signifies an additional core on that node.

Here is the R script that was submitted.

And here is the output file that was written for that job submission. If you want to output figures, graphs, or checkpoints, you'll want to explicitly send them to specific files and not to windows.

We use the Torque resource manager with the Maui scheduler. These two programs work in concert to prioritize your job, identify open cores on nodes, send your job to those open nodes and cores, track the amount of time used by the job, and finally gracefully end your job. The amount of time used by your job over the number of cores is then tracked over the past week and added to your usage to determine your Fairshare priority. All knot users have the same Fairshare priority, though we do make short-term exceptions when folks are in a crunch time. The Fairshare policy works such that the more you use the cluster, the lower your job's priority is. This generally means that you can hop on the cluster and if you haven't used it in a while your job will run pretty quickly after submission if you are not asking for a really large quantity of cores.

We have no restrictions on runtimes - probably because mother nature, So Cal Edison, and UCSB all tend to interrupt our systems with massive power outages. That said users do run jobs for days and weeks, and a few have gone up to 3 months for a single job. If you plan on a long-running job- it's highly recommended you checkpoint your job periodically so that in the case of interruption the job can pick back up near when the job was interrupted.

Serving the entire campus, knot has a lot of software installed on it. A few of the packages which might interest you….

If your research group or department is interested in having more dedicated compute power, the CSC runs 3 condo clusters where you buy the nodes and we provide the infrastructure and administration of those nodes. To maximize your compute power we supply the high performance file system and data backup, and the networking. We do limit what you can buy so that it can fit in nicely with our existing clusters - which means they'll be CentOS linux at the moment. And then we create specialized Fairshare policies for the amount of your buy-in plus some extra, since the CSC also purchases extra nodes. Condos are pretty popular with about 14 different groups buying in over the last 3 years.