

Summary

┆ **56 Pairs Passed**

┆ **10 Mismatches**

- ┆ Mismatches are African American samples that have not previously been seen.

┆ **12 Target Tumor Poor Quality**

Row Labels	Count of QC Status
Fail: Low Target Coverage Tumor	24
Fail: Mismatch	20
Fail: Poor Sample	2
Pass	112
(blank)	
Grand Total	158

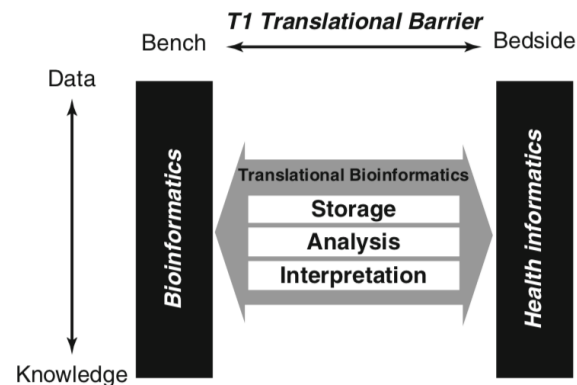
Foundations of Translational Biomedical Informatics

TRANSLATIONAL BIOINFORMATICS / BIOMEDICAL INFORMATICS

- Translation of biological (“bench”) discoveries into actual impact on clinical care (“bedside”) and ultimately on population health

Definitions

- The American Medical Informatics Association (AMIA) defines Translational Bioinformatics as “the development of storage, analytic, and interpretive methods to optimize the transformation of increasingly voluminous biomedical data, and genomic data, into proactive, predictive, preventive, and participatory health.”



Jessica D. Tenenbaum, Nigam H. Shah,
and Russ B. Altman

DATA SCIENCE

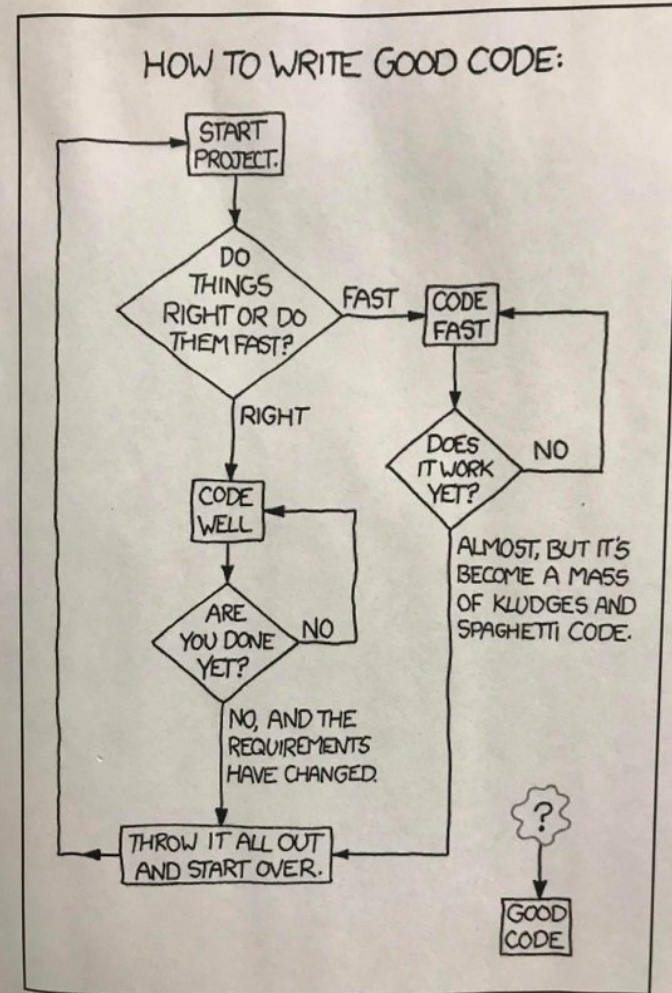
Data in the context of bioinformatics?

- Biomedical informatics involves big data. One of the biggest drivers lately is genomics and sequencing. Lets go with that example. Aligning a genome is a problem that is easy to break up. Aligning 1 read doesn't impact aligning another read. If I have 1 billion reads, and a single CPU can align 2500 reads per second it will take 72 hours. If you want the answer quicker you need to split the analysis up and send it off to different computers that ideally all share network storage so that we can keep track of data. This is solved by High-Performance-Computing.

At this point, we should highlight there are three major areas of HPC that bioinformatics sees:

- (1) Problems that can be easily broken up – divide and conquer. Analysis of genomes is a great example. Divide and conquer works well when the core calculation doesn't require knowing what the result is from other calculations. For example, aligning 1 read to a genome, frequently (except for assembly), aligns to the same position in the genome regardless of what everyother read aligns.
- (2) Problems that require lots of communication at each step & are less easy to break up.

KNOWING THE OBJECTIVE; BIOMEDICAL INFORMATICS IS OFTEN DRIVEN BY 1 OFFS





WHAT TO EXPECT

COURSE OVERVIEW

- The objective of this course is to train individuals with strong backgrounds in biological or medical fields the analytical and computational skills for analysis of biomedical data.
- It will introduce students to tools and concepts that will be instrumental throughout the program. Particular focus will be on applicability to the healthcare field and training students to effectively implement, develop, and design bioinformatic solutions within different healthcare applications from prototyping to production. They will be trained and have an understanding of modern molecular data with a major emphasis on data analysis and data processing associated with next-generation sequencing data.
- This course targets individuals who have laboratory experience generating biomedical data, and aims to provide them with the foundations, basic principles, and core concepts in scripting and computing that are necessary in biomedical informatics. The course will focus on teaching by example with the understanding that applied biomedical informatics frequently favors rapid and iteratively developed single-use analyses that must be both reproducible and documentable, for example how to work within a command-line based environment, basic scripting such as with R, bash, or javascript. They will learn versioning, unit testing, and prototyping focusing on being able to rapidly analyze and explore datasets. They will learn the fundamentals of web-applications, biomedical databases, and on-line resources and how to utilize and integrate these within one's own analyses using APIs and connectors. It will also include an overview of regex and web mining, data-types and data structures, program flow, versioning and best practices. R will be utilized throughout as part of the course to familiarize students with various programming tools. High performance computing will be introduced from a user perspective along with best practices.
- This course is an introductory level course and restricted for Masters of Science (MS) degree in Translational Biomedical Informatics. This course is not intended for those experienced in command-line tools, scripting, database, and web-based applications. This course and the timing of core concepts will complement companion courses provided at the same time.

LEARNING OBJECTIVES

- The goal of this introductory platform course is to teach core fundamentals that will allow someone trained in biology or medicine to use modern computing and bioinformatics tools to rapidly and reproducibly answer biological questions within an applied setting. The focus is not on teaching how to developing tools, modules, or frameworks for community distribution, and more focused on how researchers can use existing tools together to explore novel biomedical questions in ways that retain reproducibility (such as through versioning). Still students will learn how to interface and communicate with development teams, and gain a basic understanding of project management frameworks.
- While the course teaches using the statistical framework R, it does not teach fundamental of statistics and presumes students have had an undergraduate or graduate level biostatistics or biometrics course. Application of statistical approaches will be taught on how they can be deployed using defined software and data, though this course does not teach statistical interpretation or design of analyses as that is beyond the scope. Several electives are available should students wish to supplement and gain further expertise.
- Upon successful completion of this course, students will be able to: interface with the command line; utilize versioning tools following best practices; create basic scripts in R, create basic web-apps in R-Shiny; describe the basic concepts of data-types and data structures and when to use them; effectively use best practice high performance computing concepts.

REQUIREMENTS

- This course has specific hardware and software requirements as part of the Master's in Translational Biomedical Informatics Program. In order to optimize the ability for students to work together with uniformity there are specific computing hardware requirements. Students will be required to have a 2013+ MacBook, MacBook Air, MacBook Pro, iMac or Mac Pro with version 10.12+ (macOS) with a minimum of 4 Gbytes of RAM. The course will require a suite of open-source software that will be provide 2 weeks prior to the start of courses. Students will need to have video and audio capabilities that typically come with most MacOS computers.

MICROSOFT REPORT SCIENCE

Scientists rename human genes to stop Microsoft Excel from misreading them as dates


Sometimes it's easier to rewrite genetics than update Excel

By James Vincent | Aug 6, 2020, 8:44am EDT

f t SHARE



Illustration by Alex Castro / The Verge



VERGE D



DESCRIPTION OF ASSIGNMENTS

- Informatics conducted in lecture halls are inherently difficult, especially for those who do not have prior experience. This course forms the framework with other concurrent courses, and early participation will be essential. The work load for this course will complement other concurrent courses, and the work-load expectations will be front-loaded to insure the foundations are provided within the first half of the course.
- While content will be available on-line assignments and coursework requires timely iterative completion. It will be difficult to catch-up, and teamwork necessitate that deadlines cannot be individually altered.
- There will be a month long final project on a topic of student's choice either on their own or within a group aimed at mimicking a bioinformatics analysis, and communication to collaborators. The deliverables are a project scope of work, a written report of the data analysis in an R Markdown, a website, and a two-minute video communicating what the group learned. The project proposal described the motivation for the project, the project objectives, a description of the data, how to obtain the data, an overview of the computational methods proposed to analyze the data and a timeline for completing the project.

TIME COMMITMENT

Historically, many people spend 20 hours+ week.


- Most people 20+ hours, said that it was by their choice. A substantial portion have said that this course was one of the most valuable they took. That it was not because of the teacher. It was the content and love of data.

This is graduate course – there are points you will struggle

- I expect you to try to figure out a problem when you are stumped at least 4 hours. However, never more than 8. The problem is 95% typos.

The goal to train you how to solve problems

ALWAYS START WITH SOMETHING THAT WORKS




Answer · Recommended for you



Why should I hire a software engineer if I can just copy and paste code from Stack Overflow?


my code solving a specific problem



code found on stackoverflow solving the same problem



Programmer Morale



I don't really understand your code lol u should comment

Diogo

I don't understand it either

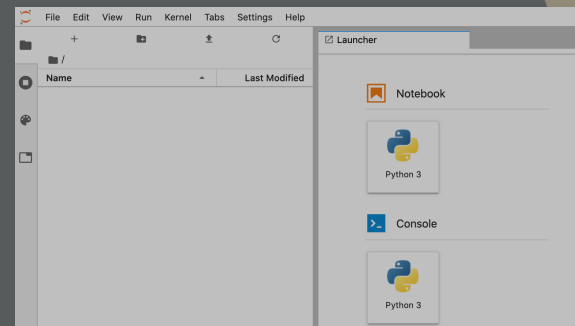
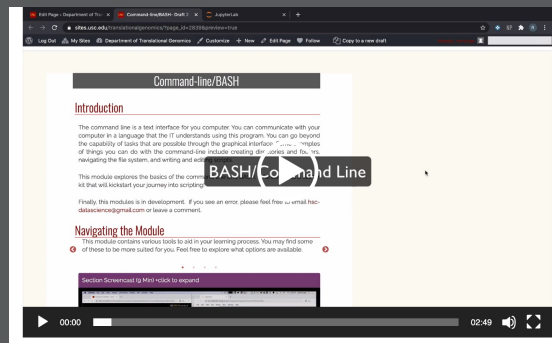
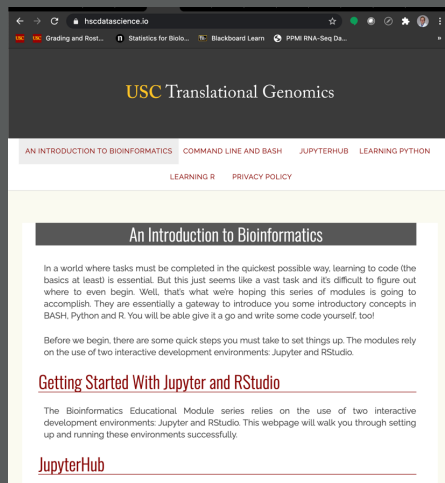
D

It was a lot of trial and error

Diving In

<https://www.hscdatascience.io/>

5+ Learning Modes

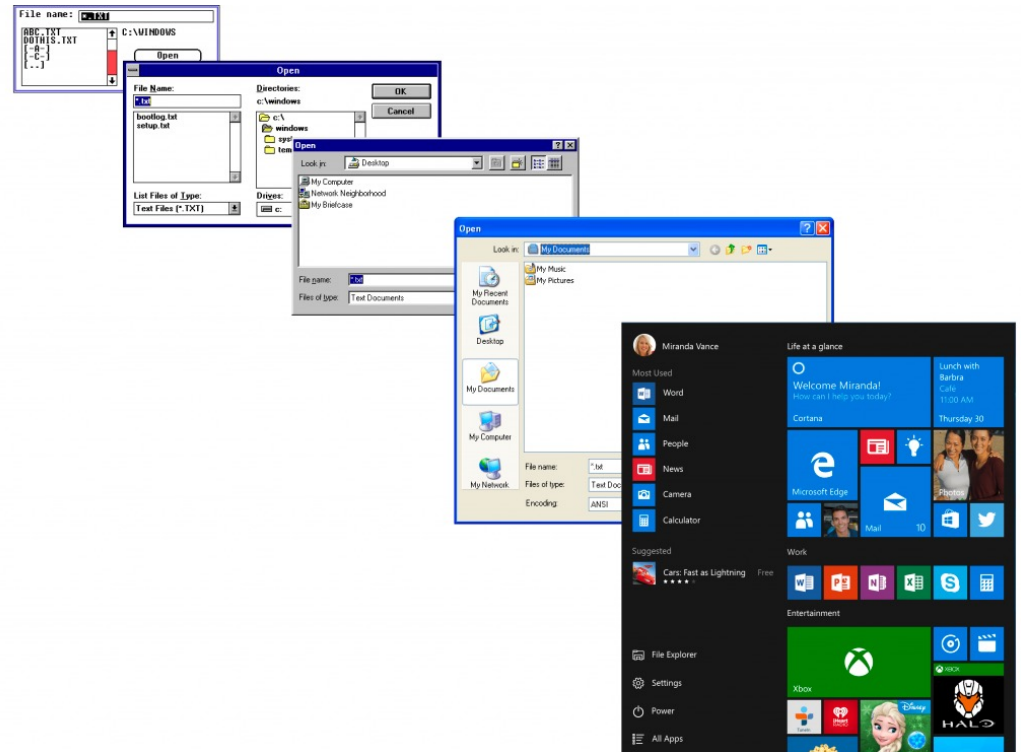




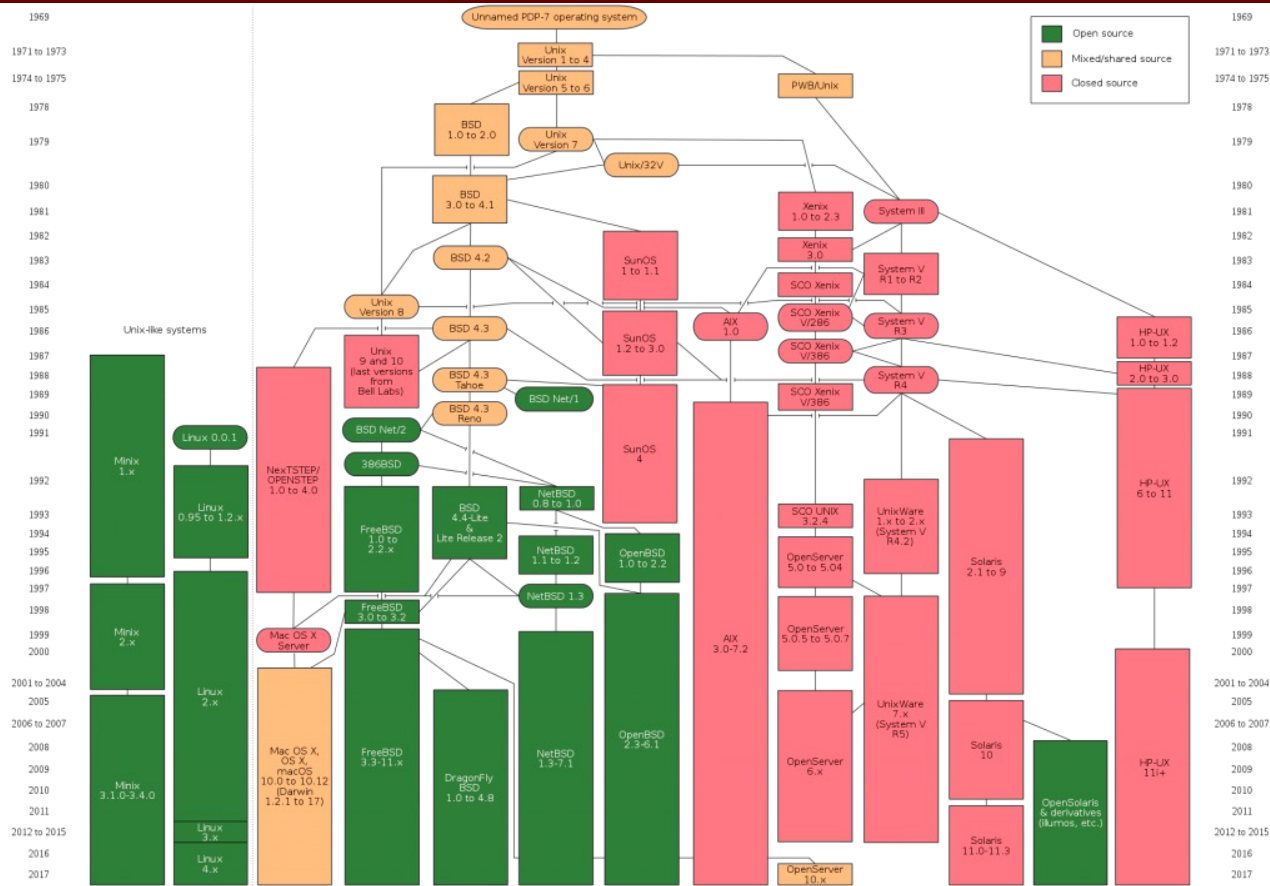
OPERATING SYSTEMS

Windows Over 20 Years

- Windows has generally seen many different Graphical User Interfaces (or GUI's) over the years. Generally, the advent of Windows and key innovation was to remove the need for shell level computing. However a lot of scripting and the ability to manipulate, wrangle, edit, and work with text was lost. Bioinformatics tends to need these tools, and these are generally found in the various flavors 'nix. There are ways to get to a 'nix environment within Windows such as through Cygwin or other VM devices but we generally don't discuss that in this material.




WHAT IS UNIX/LINUX/BASH/COMMAND-LINE/TERMINAL?



Jargon Is
Just That


SOFTWARE TERMINOLOGY

Then	Now
application	app
program	app
operating system	app
script	app
shell	app
batch file	app
compiler	app
daemon	app
service	app
game	app
patch	app
software	app



MACOS SETUP

Linux




Mac



Windows



<https://sites.usc.edu/translationalgenomics/>



Setting Up Mac

Storage



Overview | Displays | Storage | Memory | Support | Service


macOS Catalina
Version 10.15.6

iMac

Processor 3.1 GHz 6-Core Intel Core i5
Memory 64 GB 2667 MHz DDR4
Graphics Radeon Pro 575X 4 GB
Serial Number C02ZP09BJV3Y

System Report... Software Update...

™ and © 1983-2020 Apple Inc. All Rights Reserved. License and Warranty



Overview | Displays | Storage | Memory | Support | Service


Macintosh HD
814.54 GB available of 3.12 TB

Managing...

3.12 TB Fusion Drive

Manage...

Memory



Overview | Displays | Storage | Memory | Support | Service


64 GB Installed

Your Mac contains 4 memory slots, each of which accepts a 2667 MHz DDR4 memory module.
All memory slots are currently in use.

16 GB 16 GB
16 GB 16 GB

Memory Upgrade Instructions

Details



iMac Pro

Hardware Overview:

Model Name:	iMac Pro
Model Identifier:	MacPro1,1
Processor Name:	10-Core Intel Xeon W
Processor Speed:	3 GHz
Number of Processors:	1
Total Number of Cores:	10
L2 Cache (per Core):	1 MB
L3 Cache:	32.8 MB
Hyper-Threading Technology:	Enabled
Memory:	128 GB
System Firmware Version:	1721.140.2.0.0 (iBridge: 19.16.16064.0.0.0)
OS Loader Version:	540.120.3-19
Serial Number (system):	C02V088HVF8F
Hardware UUID:	4714A504-3888-5FFB-A767-989F33B10CC4
Provisioning UUID:	4714A504-3888-5FFB-A767-989F33B10CC4
Activation Lock Status:	Disabled

Hardware Overview:

- ATA
 - Apple Pay
 - Audio
 - Bluetooth
 - Camera
 - Card Reader
 - Controller
 - Diagnostics
 - Disc Burning
 - Ethernet
 - Fibre Channel
 - FireWire
 - Graphics/Displays
 - Memory
 - NVMeExpress
 - PCI
 - Parallel SCSI
 - Power
 - Printers
 - SAS
 - SATA
 - SPI
 - Storage
 - Thunderbolt/USB4
 - USB
- Network
 - Firewall
 - Locations
 - Volumes
 - WWAN
 - Wi-Fi
- Software
 - Accessibility
 - Applications
 - Developer
 - Disabled Software
 - Extensions


craig_imacpro > Hardware

Directories

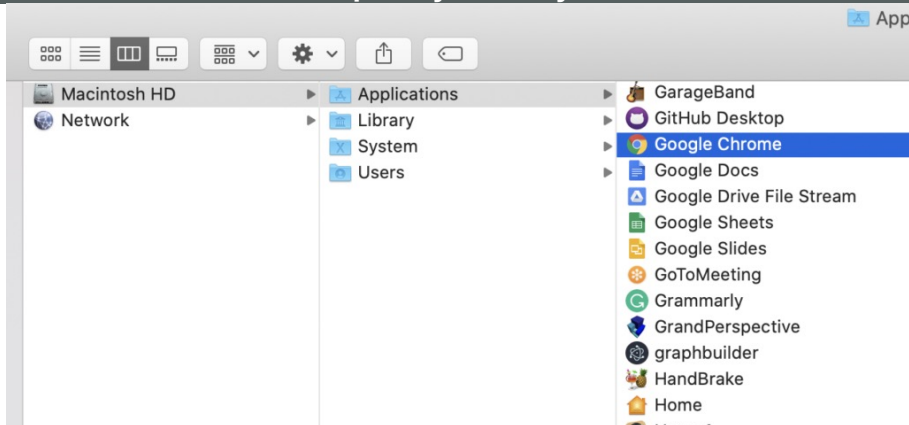
Graphical Interface



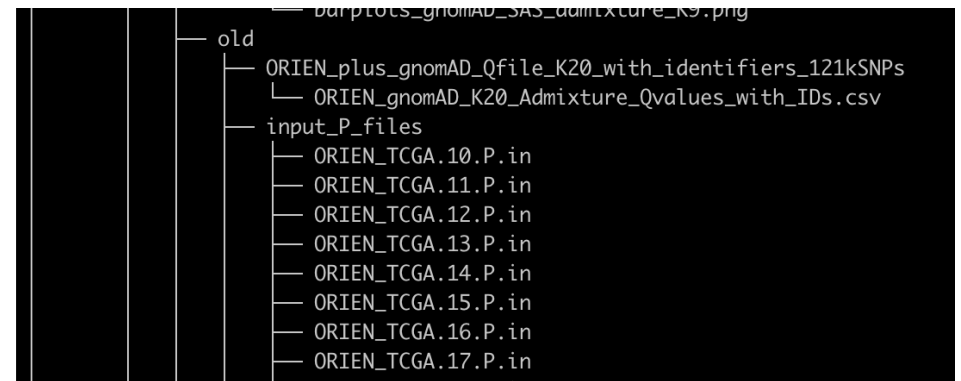
Graphical Interface: Directories of Directories



There are Locations frequently used by convention



In the end, we will focus on the command line to navigate



Setup



COMMAND-LINE SHELLS

Key Commands

- Navigating Directories

Special Characters

- ~ / * | > < . . .

Pipe Redirect

- | >

Core commands

- Find ls , among others

VIM

A right of passage. Why Mac brought back the “esc” character.



EXPECTATIONS

We will cover a lot of concepts at low depth

- You will learn & demonstrate basics of 6 scripting/database/programming languages
 - Bash
 - Python
 - R
 - Javascript
 - HTML/CSS
 - SQL
- You will learn project management and computing
- You will learn about deploying servers, high-performance computing, cloud computing.
- You will develop a multi-faceted web-application using these languages as a final project

This is a graduate level course

- You will learn to learn independently; Moving forward in the face of ambiguity is the single most important thing to learn.

The amount of weekly time Varies By Experience.

- Average 20 hours per week, with many weeks 25 hours
 - Some people report 3 to 4 hours per week; Some report more
- Graduate level courses
 - Moving forward in the face of ambiguity
- Graduate level courses: This course has extensive material provided, but you are expected to go beyond and seek external resources too!
- The material has purposeful gaps; Learning to move forward in the face of ambiguity is key
- You want to walk away with confidence at tackling most Data Science problems
- We want to make you dangerously good at biomedical informatics by the end of the course.

EXPECTATIONS II

You are expected to become independent

- This is a graduate level course.
- You'll experience frustration.
- I expect you to spend at least 4 and up to 8 hours before reaching out to help. I will not respond immediately.
- I will not respond if a deadline is less than 6 hours away.
 - *Starting an assignment with 2 hours to go, and then sending multiple emails after 1 hour demanding help does not facilitate the type of learning we are seeking.*
- Do not spend more than 8 hours solving a problem without reaching to me

Experience from Prior Years

- Those who were given the answer quickly, complained they felt they didn't learn enough in course.
- Those who received less feedback and more challenging exercises considered this to be one of most important courses they've taken.

Engagement with servers & resources is tracked

- Understanding privacy & computers and the internet is key to the course.
- The internet does not give privacy, expect that any resource is tracking your every interaction.

Seeking help

- Do use screen snapshots & videos
- Do reach out to colleagues

Giving Help

- Do help others; but don't handicap them to depend on you.
- Their goal is to find resources to learn – like your own; Is your goal to be their quick answer resource?
- Do share helpful learning resources with others
- Do share fun learning resources with others.
- Google. Can't emphasize that the secret to programming is here.
- <https://www.youtube.com/watch?v=HluANRwPyNo>

RESOURCES

Blackboard

Class Webpage

- <https://sites.usc.edu/translationalgenomics/trgn510/>