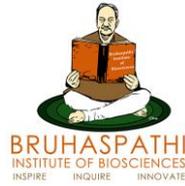


AN ABSOLUTE BEGINNER'S GUIDE TO FOUNDATIONAL BIOINFORMATICS CONCEPTS

**A BEGINNER'S
COMPANION**

DR NEELIMA CHITTURI



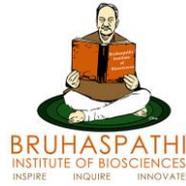


Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave,
Hyderabad, Telangana, 500081
Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

An Absolute Beginner's Guide to Foundational Bioinformatics Concepts

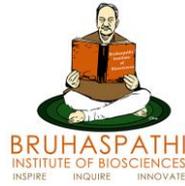
Contents

Page1: Introduction	4
About.....	4
Page 2: Understanding FASTQ Files - The Building Blocks	5
The Four-Line Structure of a FASTQ Record	5
A Complete Example of a Single Read	6
Why Are Quality Scores So Important?	6
Page 3: Quality Control (QC) - Is My Data Any Good?	7
What is FastQC?	7
Key FastQC Modules for Beginners:.....	7
What Do You Do After QC?	8
Page 4: Aligning Reads to a Reference Genome.....	9
The Goal of Alignment	9
How Does Alignment Work? (The Big Picture)	9
Understanding the SAM/BAM File Format	10
Page 5: The Idea Behind Variant Calling	11
What Kinds of Variants Are We Looking For?.....	11
How Does Variant Calling Work? The Concept of Evidence	11
The Output: The VCF File	12
Page 6: Using the Command Line Daily	13
Why is the CLI so essential?	13
Your First Steps: Navigating Your Filesystem.....	13
Daily Practice.....	14
Page 7: Automating a Simple Pipeline	15
What a Simple Pipeline Looks Like.....	15



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave,
Hyderabad, Telangana, 500081
Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

Your First Taste of Automation: The Shell Script	16
What's happening here?	17
Page 8: Review and Next Steps.....	18
The Core Bioinformatics Workflow: A Summary	18
What to Focus on After Your First 30 Days	18
Page 9: Glossary of Terms.....	20
Page 10: Final Thoughts and Resources.....	22
Page 11: About Bruhaspathi	23
Our activities	23



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave, Hyderabad, Telangana, 500081

Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

Page1: Introduction

About

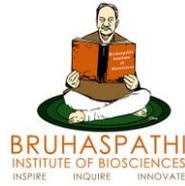
Welcome to the world of bioinformatics! It's a field that combines biology, computer science, and statistics to make sense of biological data. If you're just starting, it can feel like drinking from a firehose. There's so much to learn, and it's tough to know where to begin.

This guide is designed to accompany the 30-Day Focus Checklist and give you a solid foundation in the most critical concepts you'll encounter at the beginning of your journey. The goal here isn't to make you an expert overnight. The goal is to provide exposure, build a mental framework, and give you the confidence to take the next steps.

We will break down six fundamental topics over the next pages. We'll explore:

- 1. The FASTQ file format:** The raw data of genomics.
- 2. Quality Control (QC):** How to know if your data is any good.
- 3. Genome Alignment:** How to give your data context.
- 4. Variant Calling:** How to find what's unique in your data.
- 5. The Command Line:** Your most essential tool.
- 6. Pipelines:** How to automate your work for efficiency and reproducibility.

Remember, every expert was once a beginner. Take it one step at a time, be patient with yourself, and stay curious. Let's begin.



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave, Hyderabad, Telangana, 500081

Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

Page 2: Understanding FASTQ Files - The Building Blocks

Imagine you have a book, but instead of pages, it has been run through a shredder a billion times. Each tiny strip of paper is a short sentence fragment from the original book. This is essentially what modern sequencing machines do to an organism's DNA. They can't read the whole genome from start to finish; instead, they read millions or billions of short DNA fragments.

A FASTQ file is the standard text-based format used to store these short DNA sequence fragments (we call them reads) and their corresponding quality scores. It's the raw output from a high-throughput sequencing machine.

So, what does it look like? A single read in a FASTQ file is always represented by four lines. Let's break them down.

The Four-Line Structure of a FASTQ Record

Line 1: The Sequence Identifier (or Header)

- This line always starts with an '@' character.
- It contains information about the read, like the instrument that sequenced it, its location on the machine's flow cell, and a unique ID for the read itself.
- Example: @SEQ_ID:1:HWI-ST0001:100:H0000BCXX:1:1101:1234:2345 1:N:0:ATCACG

Line 2: The Biological Sequence (The Read)

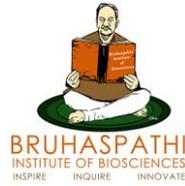
- This is the actual sequence of DNA bases (A, C, T, G).
- Sometimes, if a base cannot be determined with confidence, it will be represented by an 'N'.
- The length of this sequence is called the read length (e.g., 150 base pairs).
- Example: GATTTGGGGTTCAAAGCAGTATCGATCAAATAGTAAATCCATTTGTTCAACTCACAGTTT

Line 3: The Separator

- This line always starts with a '+' character.
- Sometimes, the sequence identifier from Line 1 is repeated here, but often it's just the '+' by itself.
- Its main purpose is to separate the sequence from its quality scores.
- Example: +

Line 4: The Quality Scores

- This line contains a string of characters that represent the quality of each base in the sequence on Line 2.
- This string MUST have the exact same number of characters as the sequence in Line 2.



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave,
Hyderabad, Telangana, 500081
Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

- Each character is an encoded representation of a Phred quality score. This score tells you the probability that the base call from the sequencing machine was incorrect.
- Example: `#"*((((***)%+%+)(%+%%).1***-+*")**55CCF>>>>>CCCCCCC65`

A Complete Example of a Single Read

```
@SEQ_ID:1:HWI-ST0001:100:H0000BCXX:1:1101:1234:2345 1:N:0:ATCACG
GATTTGGGGTTCAAAGCAGTATCGATCAAATAGTAAATCCATTTGTTCAACTCACAGTTT
+
#"*((((***)%+%+)(%+%%).1***-+*")**55CCF>>>>>CCCCCCC65
```

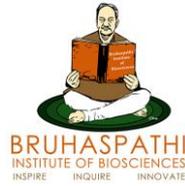
A real FASTQ file will contain millions or even billions of these 4-line blocks, one after another.

Why Are Quality Scores So Important?

The sequencing process is not perfect. The quality scores allow us to be skeptical of the data. A high-quality score means we can be very confident that the base is correct. A low-quality score is a red flag, telling us that the 'A', 'T', 'C', or 'G' at that position might be an error.

Downstream analysis tools, like aligners and variant callers, use these scores to make more informed decisions. For instance, if a read has a mismatch with the reference genome, but the quality score at that mismatch position is very low, the program might ignore it as a likely sequencing error rather than a true biological variant.

Your first goal is to simply be able to recognize a FASTQ file, understand its 4-line structure, and appreciate that it contains both the sequence read and a measure of its quality.



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave, Hyderabad, Telangana, 500081

Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

Page 3: Quality Control (QC) - Is My Data Any Good?

Now that you have your FASTQ file, filled with millions of reads, you might be tempted to jump straight into analysis. But wait! How do you know if that data is reliable?

Just like in any scientific experiment, the first step after data collection is to check its quality. This is Quality Control or QC. In sequencing, this means running checks to ensure the raw reads from the FASTQ file are good enough for downstream analysis. Bad data can lead to false conclusions, wasted time, and incorrect results.

The most popular tool for this job is called FastQC.

What is FastQC?

FastQC is a program that reads your entire FASTQ file and generates a summary report in the form of an HTML file. This report contains a series of graphs and tables that help you assess the quality of your sequencing run from many different angles.

You don't need to understand every single detail of a FastQC report at first. The goal is to learn how to spot major problems. The report gives a PASS, WARN, or FAIL for each module. As a beginner, pay close attention to the WARN and FAIL sections.

Let's look at a few of the most important modules in a FastQC report:

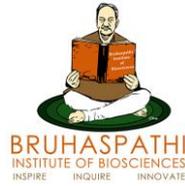
Key FastQC Modules for Beginners:

Per Base Sequence Quality

- **What it is:** This is the most important plot. It shows the distribution of quality scores at each position across all reads. The x-axis is the position in the read (e.g., from 1 to 150), and the y-axis is the quality score.
- **What you want to see:** Quality scores should be high (generally above 28, which is the green zone on the graph) across most of the read.
- **Common problems:** It's normal for quality to drop off slightly towards the end of the reads. However, if the quality starts low, or drops dramatically, it could indicate a problem with the sequencing run. If the quality is very poor, you might need to trim the low-quality ends off your reads before proceeding.

Per Sequence GC Content

- **What it is:** This plot shows the distribution of GC content (the percentage of Gs and Cs) across all reads.
- **What you want to see:** Usually, this is a nice bell-shaped curve centered around the known GC content of the organism you sequenced. For humans, this is around 40-42%.



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave, Hyderabad, Telangana, 500081

Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

- **Common problems:** A sharp spike at a different location or a very broad, strange-looking curve might suggest contamination. For example, if you were sequencing a human sample but have a sharp peak at 65% GC, you might have bacterial contamination in your sample.

Sequence Duplication Levels

- **What it is:** This plot shows the percentage of your library that is made up of duplicated sequences.
- **What you want to see:** A certain level of duplication is normal. However, very high duplication (e.g., >20-30%) can be a problem.

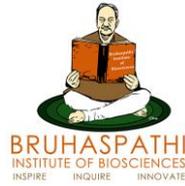
- **Common problems:** High duplication can mean you started with too little DNA, leading to the same few molecules being amplified and sequenced over and over (a technical artifact). It can reduce the complexity of your data and your power to discover new things. In some experiment types (like RNA-seq), high duplication is expected for highly expressed genes, so context is important.

What Do You Do After QC?

Running FastQC gives you a diagnosis of your data's health. The next step is often data cleaning or pre-processing. This might involve:

- **Trimming:** Using tools (like Trimmomatic or fastp) to cut off the low-quality bases from the ends of the reads.
- **Adapter Removal:** Removing sequences from the sequencing machine's adapters that can sometimes get read into your FASTQ file.

For now, your goal is to understand why QC is a critical first step and to become familiar with the basic plots in a FastQC report. You need to be able to answer the fundamental question: Is my data clean enough to trust?



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave,
Hyderabad, Telangana, 500081
Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

Page 4: Aligning Reads to a Reference Genome

Let's return to our shredded book analogy. You have millions of tiny sentence fragments (your FASTQ reads), and you've checked their quality. Now what? How do you put them back together to read the story?

You need a map. You need a complete, intact copy of the book to compare your fragments against. In genomics, this map is called a reference genome.

A reference genome is a high-quality, well-annotated version of a species' complete DNA sequence, assembled by the scientific community. For example, the human reference genome (known by codes like hg38 or GRCh38) is our master copy of the human DNA sequence.

Alignment (also called mapping) is the process of taking each of your short reads from the FASTQ file and finding its exact location on the reference genome.

The Goal of Alignment

The core idea is to figure out where each read came from. By aligning millions of reads, you can pile them up on top of the reference genome. This allows you to reconstruct the genome of the individual you sequenced and, more importantly, see where that individual's genome differs from the reference.

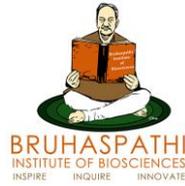
Think of it like this:

- **Reference Genome:** A master blueprint of a house.
- **Your Reads:** Thousands of photos of a specific house being built.
- **Alignment:** The process of figuring out where each photo was taken (e.g., This photo of a window matches this spot on the blueprint's second floor.).

How Does Alignment Work? (The Big Picture)

You don't need to know the complex algorithms right away, but you should understand the inputs and outputs.

- **Input:**
 1. A reference genome file (usually in FASTA format).
 2. Your quality-checked FASTQ file(s).
- **The Tool:** You use a specialized program called an aligner. The most common ones for short reads are BWA (Burrows-Wheeler Aligner) and Bowtie2. These tools are incredibly fast and efficient at finding the best match for billions of reads against a massive reference genome.
- **Output:** The result of the alignment is a SAM (Sequence Alignment/Map) file.



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave,
Hyderabad, Telangana, 500081

Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

Understanding the SAM/BAM File Format

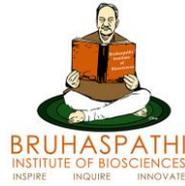
A SAM file is a large text file that, for each read, tells you:

- The read's ID.
- Whether it mapped to the reference genome or not.
- If it did map, which chromosome it mapped to.
- The exact start and end coordinates of its position on that chromosome.
- Information about how well it matched (e.g., were there any differences?).

SAM files are text-based and human-readable, but they are enormous. Because of their size, we almost always compress them into a BAM (Binary Alignment/Map) file. A BAM file contains the exact same information as a SAM file, but it's in a binary (computer-readable) format that is much smaller and faster for programs to work with.

You will also almost always see a BAI file alongside a BAM file. This is an index file that acts like a table of contents for the BAM, allowing programs to quickly jump to a specific region of the genome (e.g., show me all the reads at chromosome 3, position 5,000,000) without having to read through the entire file.

Your goal for this section is to understand the concept of alignment: taking short reads and finding their home on a reference genome. You should know that this process requires an aligner (like BWA), and it produces a SAM/BAM file as its primary output.



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave, Hyderabad, Telangana, 500081
Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

Page 5: The Idea Behind Variant Calling

You've done the hard work: you have your raw data (FASTQ), you've checked its quality (FastQC), and you've given it context by mapping it to a reference genome (BAM). Now we get to the exciting part: discovery.

The human reference genome is just that—a reference. It's a representative sequence, but every individual has a slightly different DNA sequence. These differences are called genetic variants.

Variant calling is the process of identifying these differences between your sequenced sample (represented by the reads in your BAM file) and the reference genome.

What Kinds of Variants Are We Looking For?

There are many types of genetic variation, but for a beginner, the two most important to understand are:

1. SNPs (Single Nucleotide Polymorphisms):

- This is the most common type of variation. It's a single base change at a specific position.
- **Analogy:** The reference book says thought, but in your sample's book, the word is sought. It's a single letter difference. Example: The reference genome has an 'A' at a position, but in your sample, you consistently see a 'G' in the reads that cover that position.

2. Indels (Insertions and Deletions):

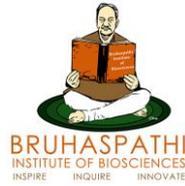
- **Insertion:** One or more bases are added into your sample's DNA that are not in the reference.
- **Deletion:** One or more bases are present in the reference genome but are missing from your sample.
- **Analogy:** The reference book says the cat sat, but your sample's book says the black cat sat (an insertion) or the sat (a deletion).

How Does Variant Calling Work? The Concept of Evidence

A variant caller is a program that systematically scans your BAM file, position by position, looking for evidence of a difference from the reference.

Imagine you are looking at a single position on chromosome 1. Let's say the reference base is a 'T'. You look at your BAM file and see 50 reads piled up over that exact spot.

- **Scenario 1 (No Variant):** All 50 reads also have a 'T' at that position. Your sample matches the reference here. No variant is called.
- **Scenario 2 (Likely Variant):** All 50 reads have a 'C' at that position. The quality scores for that 'C' in all the reads are very high. This is strong evidence for a homozygous SNP (meaning the variant is present on both copies of the chromosome, from mother and father). The variant caller would report a T -> C variant at this location.



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave, Hyderabad, Telangana, 500081

Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

- **Scenario 3 (Possible Heterozygous Variant):** Roughly half the reads (around 25) have a 'T', and the other half (around 25) have a 'C'. All quality scores are high. This is evidence for a heterozygous SNP (the individual has a 'T' on one copy of the chromosome and a 'C' on the other).
- **Scenario 4 (Likely Sequencing Error):** 49 reads have a 'T', and only one read has a 'C'. Furthermore, the quality score for that single 'C' is very low. The variant caller will likely dismiss this as a sequencing error, not a true biological variant.

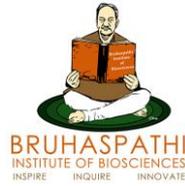
The variant caller's job is to be a detective, weighing all the evidence (read depth, base quality, mapping quality, etc.) to make a statistically informed call about whether a variant exists at each position.

The Output: The VCF File

For each variant, a VCF file will tell you:

- The chromosome and position.
- The reference base(s).
- The alternate base(s) (what was found in your sample).
- A quality score for the variant call itself (how confident the program is).
- Lots of other detailed information, including the genotype (e.g., homozygous or heterozygous).

Your goal is to understand that variant calling is about finding confident differences between your aligned reads and the reference. You should know the basic difference between a SNP and an indel, and recognize that the final output of this process is a VCF file.



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave, Hyderabad, Telangana, 500081
Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

Page 6: Using the Command Line Daily

So far, we've talked about concepts and file formats. But how do you actually do any of this? How do you run FastQC, BWA, or a variant caller?

The answer is the command-line interface (CLI), also known as the terminal or shell.

In bioinformatics, the vast majority of software is designed to be run from the command line, not with a fancy graphical user interface (GUI). While it might seem intimidating at first, it is an incredibly powerful, flexible, and efficient way to work with data. Learning to use it is not optional; it is the single most important practical skill you can develop.

Why is the CLI so essential?

- **Automation:** You can write scripts to run a sequence of commands, creating reproducible workflows.
- **Power:** You can connect tools together, sending the output of one program directly to another.
- **Remote Access:** It allows you to log in and work on powerful servers and high-performance computing (HPC) clusters, which is where most large-scale bioinformatics analysis happens.
- **Universality:** Nearly all bioinformatics tools are built for the command line first.

The key to getting comfortable is to use it every day, even for small tasks.

Your First Steps: Navigating Your Filesystem

Before you can run complex tools, you need to learn how to move around and manage your files. Here are the absolute essential commands to practice daily.

(Note: These are for Linux/macOS. Windows has equivalents like dir and move, but it's highly recommended to use the Windows Subsystem for Linux (WSL) to get a true Linux environment.)

1. pwd (Print Working Directory)

- **What it does:** Tells you which directory you are currently in. It's your You are here map marker.

```
pwd
```

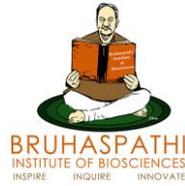
```
/home/user/project_alpha
```

2. ls (List)

- **What it does:** Lists the files and directories in your current directory. Use ls -lh to get a more detailed, human-readable list with file sizes and permissions.

```
ls -lh
```

```
-rw-r--r-- 1 user user 1.2G Jan 15 10:30 sample1.fastq
```



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave,
Hyderabad, Telangana, 500081
Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

```
-rw-r--r-- 1 user user 1.3G Jan 15 10:32 sample2.fastq
```

```
drwxr-xr-x 2 user user 4.0K Jan 15 11:00 results
```

3. cd (Change Directory)

- **What it does:** Moves you into a different directory.

```
cd results; pwd
```

```
/home/user/project_alpha/results
```

- **Pro-tips:** cd .. moves you up one level. cd ~ or just cd takes you back to your home directory.

4. mkdir (Make Directory)

- **What it does:** Creates a new directory. It's crucial for organizing your work.

```
mkdir qc_reportS; ls
```

```
qc_reports
```

5. cp (Copy)

- **What it does:** Copies a file.

```
cp sample1.fastq sample1_backup.fastq
```

6. mv (Move)

- **What it does:** Moves a file to a new directory, or renames a file. mv <source> <destination>

```
mv sample1_fastqc.html qc_reports/
```

```
mv sample1_backup.fastq old_sample1.fastq
```

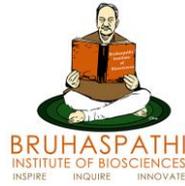
7. rm (Remove)

- **What it does:** Deletes a file. BE CAREFUL! There is no undo or recycling bin on the command line. Once it's gone, it's gone.

```
rm old_sample1.fastq
```

Daily Practice

Challenge yourself to use only the command line for your file management for a week. Create directories for a new project. Copy your data into them. Rename files. Make a backup. Your goal is to make these commands second nature. Once they are, running a bioinformatics tool will be as simple as typing its name and telling it which files to use.



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave,
Hyderabad, Telangana, 500081
Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

Page 7: Automating a Simple Pipeline

You've learned the individual steps. You know how to run a single command on the command line to perform a task like QC or alignment. But a real analysis involves many steps, executed in a specific order.

Doing this manually is:

- **Tedious:** Imagine running the same 10 commands on 100 different samples. That's 1,000 manual steps!
- **Error-prone:** You might forget a step, mistype a filename, or use the wrong parameter.
- **Not reproducible:** If you come back to your project in six months, will you remember the exact commands and order you used?

This is where automation comes in. A pipeline (or workflow) is a series of computational steps chained together to process data. Your goal is to automate these steps so they can be run with a single command.

What a Simple Pipeline Looks Like

Let's design a basic variant calling pipeline based on what we've learned.

Goal: Go from a raw FASTQ file to a VCF file of called variants.

The Steps:

1. QC on Raw Reads:

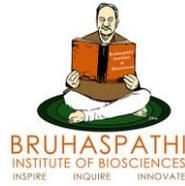
- **Input:** sample1.fastq
- **Tool:** fastqc
- **Output:** sample1_fastqc.html

2. Alignment:

- **Input:** reference.fasta, sample1.fastq
- **Tool:** bwa mem
- **Output:** sample1.sam

3. Convert SAM to BAM:

- **Input:** sample1.sam
- **Tool:** samtools view
- **Output:** sample1.bam



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave,
Hyderabad, Telangana, 500081
Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

4. Sort the BAM:

- **Input:** sample1.bam
- **Tool:** samtools sort
- **Output:** sample1.sorted.bam (Sorted BAM is required by most variant callers)

5. Index the BAM:

- **Input:** sample1.sorted.bam
- **Tool:** samtools index
- **Output:** sample1.sorted.bam.bai

6. Variant Calling:

- **Input:** reference.fasta, sample1.sorted.bam
- **Tool:** bcftools mpileup and bcftools call (a common combination)
- **Output:** sample1.vcf

This is a six-step process for a single sample. Now, let's think about how to automate it.

Your First Taste of Automation: The Shell Script

The simplest way to build a pipeline is to write a shell script. A shell script is just a plain text file that contains all the command-line commands you want to run, in the correct order.

Here's what a very basic script to automate our pipeline might look like. Don't worry about the syntax details yet, just focus on the concept.

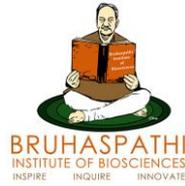
```
#!/bin/bash
# This is a simple variant calling script.

# --- Define filenames ---
SAMPLE=sample1
REFERENCE=reference.fasta

# --- Run the pipeline ---
echo Step 1: Running FastQC...
fastqc ${SAMPLE}.fastq

echo Step 2: Aligning with BWA...
bwa mem ${REFERENCE} ${SAMPLE}.fastq > ${SAMPLE}.sam

echo Step 3 & 4: Converting to sorted BAM with Samtools...
samtools view -S -b ${SAMPLE}.sam | samtools sort -o ${SAMPLE}.sorted.bam
```



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave,
Hyderabad, Telangana, 500081

Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

echo Step 5: Indexing the BAM...

```
samtools index ${SAMPLE}.sorted.bam
```

echo Step 6: Calling variants with BCFtools...

```
bcftools mpileup -f ${REFERENCE} ${SAMPLE}.sorted.bam | bcftools call -mv -Ov -o ${SAMPLE}.vcf
```

echo Pipeline finished!

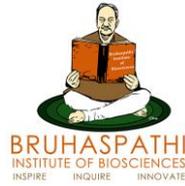
What's happening here?

- We put all our commands into one file (e.g., run_pipeline.sh).
- We used a variable (SAMPLE) to make it easy to run this on a different sample just by changing one line at the top.
- We added echo statements to print progress messages to the screen.
- We used the pipe character (|) to send the output of one command directly into the input of another, making the process more efficient (this is called piping).

Now, instead of typing six long commands, you could just run one: ./run_pipeline.sh.

This is the essence of automation. Your first goal is not to become a master programmer, but to simply try writing a small script that combines two or three commands you've learned. Maybe it's a script that runs FastQC and then moves the report into a specific folder.

Start small. The confidence you gain from automating one simple task will be immense and will set you on the path to building more complex and powerful workflows. More advanced workflow management systems like Snakemake or Nextflow build on this fundamental concept.



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave,
Hyderabad, Telangana, 500081
Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

Page 8: Review and Next Steps

Let's take a moment to pause and review the entire process we've mapped out. It's a journey that takes raw, meaningless data and transforms it into biological insight. Understanding this flow is the key to thinking like a bioinformatician.

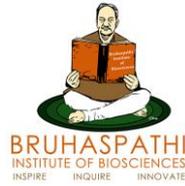
The Core Bioinformatics Workflow: A Summary

- 1. The Question:** It all starts with a biological question. (e.g., Which genetic variants does this patient have that might explain their disease?)
- 2. The Experiment:** A biological sample (like blood or tissue) is prepared, and its DNA is sequenced.
- 3. The Raw Data (FASTQ):** The sequencing machine produces FASTQ files. These files contain millions of short sequence reads and their associated quality scores. At this stage, the data is just a massive, unordered collection of fragments.
- 4. Quality Control (FastQC):** You scrutinize the raw data. You run tools like FastQC to check for sequencing errors, contamination, or other technical problems. If necessary, you clean the data by trimming bad-quality bases or removing adapter sequences. This step ensures you're not building your analysis on a faulty foundation.
- 5. Alignment (BWA/Bowtie2 -> BAM):** You take your clean reads and map them to a reference genome. This gives your reads context. You go from having a pile of random sentence fragments to having them neatly organized and stacked on top of the master copy of the book. The output is a BAM file.
- 6. Variant Calling (BCFtools/GATK -> VCF):** With your reads aligned, you can now play spot the difference. You systematically compare the stack of reads to the reference genome at every position. Where there is enough evidence of a consistent difference, you call a variant. The output is a VCF file, which is a list of all the SNPs and Indels you found.
- 7. Annotation and Interpretation:** The journey doesn't end with a VCF file! A VCF file just tells you where the variants are (e.g., a C->T change at chr7:55249071). The next step, called annotation, is to figure out what that variant means. Does it fall within a gene? Does it change an amino acid in a protein? Is it known to be associated with a disease? This final step is where you connect your findings back to the original biological question.

What to Focus on After Your First 30 Days

You have been exposed to the core workflow. You don't need to be an expert in any single step yet, but you should have a mental map of the entire process.

Here are some good directions to go from here:

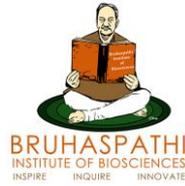


Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave,
Hyderabad, Telangana, 500081

Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

- **Get Comfortable with Linux:** Double down on your command-line skills. Learn about permissions (chmod), finding files (find), and searching inside files (grep). These are tools you will use every single day.
- **Learn a Scripting Language:** While shell scripts are great for simple pipelines, you'll quickly need something more powerful. Python is the most popular and versatile language in bioinformatics. Start with the basics: variables, loops, and functions. Then, you can learn how to use libraries like pysam (for reading BAM files) or pyvcf (for reading VCF files) to programmatically analyze your data.
- **Pick One Area and Go Deeper:** Choose one of the steps and learn more about it.
 - **Interested in QC?** Learn about other tools like MultiQC to aggregate reports.
 - **Interested in alignment?** Read about the difference between BWA-MEM and Bowtie2 and when you might choose one over the other.
 - **Interested in variant calling?** Read about the GATK framework from the Broad Institute, which is the industry standard for variant calling in humans.
- **Find a Project:** The best way to learn is by doing. Find a public dataset from a study you find interesting and try to replicate their analysis pipeline from scratch.

This guide has given you the map. Now it's time for you to start exploring the territory. Stay curious, be persistent, and don't be afraid to break things and then figure out how to fix them. That is the path to true learning.



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave,
Hyderabad, Telangana, 500081
Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

Page 9: Glossary of Terms

Adapter: A synthetic DNA sequence ligated to the ends of DNA fragments during library preparation for sequencing. Sometimes these get sequenced and need to be removed.

Aligner: A software tool (e.g., BWA, Bowtie2) that maps sequencing reads to a reference genome.

BAM (Binary Alignment Map): The compressed, binary version of a SAM file. This is the standard format for storing aligned sequencing reads.

BCFtools: A set of utilities that manipulates variant calls in the VCF and BCF format. Often used for variant calling in conjunction with SAMtools.

BWA (Burrows-Wheeler Aligner): A popular and widely used software package for mapping short sequencing reads to a large reference genome.

CLI (Command-Line Interface): A text-based interface used for running programs and managing files on a computer. The primary way bioinformaticians interact with their tools and data.

FASTQ: A text-based file format that stores both a biological sequence (read) and its corresponding quality scores. The raw output of most sequencing platforms.

FastQC: A popular tool used to perform quality control checks on raw sequence data.

GATK (Genome Analysis Toolkit): A comprehensive toolkit from the Broad Institute, considered the gold standard for variant discovery in humans.

GC Content: The percentage of Guanine (G) and Cytosine (C) bases in a DNA sequence.

Genome: An organism's complete set of DNA, including all of its genes.

Genotype: The combination of alleles an individual has at a particular genetic locus. For a diploid organism, this could be homozygous (e.g., A/A) or heterozygous (e.g., A/G).

Indel: A type of genetic variant involving the INsertion or DEletion of one or more bases.

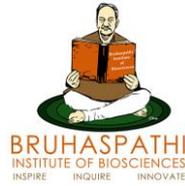
Index File: A companion file (e.g., .bai for BAM, .tbi for VCF) that provides an index of the main file, allowing for rapid access to specific genomic regions.

Phred Quality Score: A numerical score that represents the probability of an incorrect base call in sequencing. Higher scores mean higher confidence.

Pipeline/Workflow: A series of computational tools or commands chained together to perform a multi-step analysis.

Read: A short fragment of DNA sequence generated by a high-throughput sequencing machine.

Reference Genome: A digital, high-quality master copy of a species' genome assembled by the scientific community. Used as a map for alignment.



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave, Hyderabad, Telangana, 500081

Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

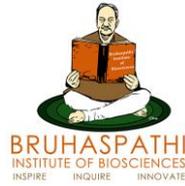
SAM (Sequence Alignment Map): A text-based format for storing aligned sequencing reads. BAM is the compressed version of SAM.

SAMtools: A suite of programs for interacting with and processing SAM and BAM files.

Shell Script: A text file containing a sequence of command-line commands that can be executed as a single program. Used for simple automation.

SNP (Single Nucleotide Polymorphism): A variation at a single position in a DNA sequence among individuals. The most common form of genetic variation.

VCF (Variant Call Format): The standard text file format for storing information about genetic variants found in a sample.



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave,
Hyderabad, Telangana, 500081

Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

Page 10: Final Thoughts and Resources

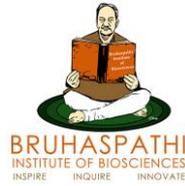
You have taken a huge first step into a very exciting and rewarding field. The journey of a thousand miles begins with a single step, and you have just walked the first ten pages.

The feeling of being overwhelmed is normal. No one masters all of this at once. The key is consistent, small efforts. Spend 20-30 minutes every day on the command line. Reread the section on FASTQ files until the 4-line structure is second nature. Download a public dataset and just try to run FastQC on it.

Success in this field is less about genius and more about persistence. It's about developing a problem-solving mindset. When you encounter an error (and you will, constantly!), don't get discouraged. See it as a puzzle. Read the error message carefully. Google it. Try to understand what the program is telling you it needs. This cycle of error -> research -> solution is how real learning happens.

Keep this guide as a reference. When you're in the middle of an analysis and forget what a BAM file is for, or can't remember the name of that QC tool, come back and review.

You have the checklist. You have the foundational knowledge. Now go and build something. Good luck!



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave, Hyderabad, Telangana, 500081

Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

Page 11: About Bruhaspathi

At Bruhaspathi, we believe in making bioinformatics accessible and actionable. We are passionate about bridging the gap between cutting-edge biological research and powerful computational tools.

Our mission is to equip the next generation of scientists, researchers, and innovators with the skills needed to tackle the biggest challenges in genomics, health, and biotechnology. We go beyond theoretical concepts, providing hands-on, project-based training that mirrors real-world bioinformatics pipelines.

Whether you're a biologist looking to enhance your computational skills, a programmer curious about biological data, or a student charting your career path, Bruhaspathi offers tailored courses and resources to help you succeed.

Ready to transform raw data into profound biological insights?

This guide is just your first step. Explore our hands-on, project-based courses at and take your bioinformatics skills to the next level.

#Bruhaspathi #Bioinformatics #NGS #LearnToCode

Our activities

Internships / Projects / Commercial collaborations

Covers all domains of bioinformatics

Short term academics

Basics of Molecular Biology

Programming: Linux, Shell Scripting, Python

Biological Databases

Sequence Alignment & Analysis

Genomics: Variant Identification, DNA Methylation Analysis, Genome Annotation

Transcriptomics: Differential Gene Expression, miRNA Profiling

Proteomics: Protein Identification & Quantification

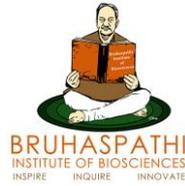
Metagenomics: 16S rRNA, Shotgun, Antimicrobial Resistance

Mid-term learning

Comprehensive Bioinformatics Training Program (6 Months)

Long-term learning

PG diploma in bioinformatics program 1 year starting June 2026



Bruhaspathi Institute of Biosciences (OPC) Private Limited, Rent A Desk, Plot No 4/2, Sector 1, Madhapur, HUDA Techno Enclave,
Hyderabad, Telangana, 500081
Phone: +91-9494113031 | Email: bruhaspathi.bio@gmail.com

WE HAVE REACHED THE END OF THIS BOOK, BUT THE
JOURNEY OF LEARNING CONTINUES—STAY TUNED FOR
MORE UPDATES

WEBSITE
[HTTPS://BRUHASPETHI.IN](https://bruhaspathi.in)

CONNECT WITH US
BRUHASPETHI.BIO@GMAIL.COM

FOLLOW US
[LINKEDIN](#) [INSTAGRAM](#) [FACEBOOK](#) [X](#)

MEET US AT
BRUHASPETHI INSTITUTE OF BIOSCIENCES, RENT A DESK,
PLOT NO 4/2, SECTOR 1, MADHAPUR, HUDA TECHNO
ENCLAVE, HYDERABAD, TELANGANA, 500081

ABOUT THE AUTHOR

DRIVEN BY HER PASSION FOR TEACHING AND RESEARCH, DR. NEELIMA CHITTURI FOUNDED THE BRUHASPATHI INSTITUTE OF BIOSCIENCES. A GOLD MEDALIST FROM THE UNIVERSITY OF MADRAS, SHE BRINGS OVER 15 YEARS OF EXPERIENCE IN MOLECULAR BIOLOGY AND RELATED FIELDS. HER CAREER INCLUDES 8 YEARS DEDICATED TO RESEARCH, INCLUDING A PH.D. IN TRANSCRIPTOMICS, AND 8 YEARS IN INDUSTRY WORK. HER EXPERTISE ENCOMPASSES DEVELOPING ANALYTICAL METHODS, UTILIZING BIOINFORMATICS TOOLS, AND TRANSFORMING COMPLEX BIOLOGICAL DATA INTO VALUABLE RESEARCH AND CLINICAL OUTCOMES. COMMITTED TO BRIDGING THEORY AND PRACTICAL TRAINING, SHE ESTABLISHED THE INSTITUTE IN 2025 TO HELP STUDENTS AND PROFESSIONALS INTEGRATE BIOLOGY WITH COMPUTER SCIENCE.



BRUHASPATHI
INSTITUTE OF BIOSCIENCES
INSPIRE INQUIRE INNOVATE