

Index

Chapter 1. Introduction to Bioinformatics

1.1 Definition, Scope, and Goals of Bioinformatics	1
1.2 Historical Perspective and Evolution	3
1.3 Interdisciplinary Nature: Biology, Computer Science & Statistics	5
1.4 Overview of Bioinformatics Tools and Techniques	5
1.5 Emerging Trends	6
1.6 Importance and Applications in Modern Biology	8
1.7 Ethical, Legal, and Social Implications (ELSI)	8

Chapter 2. Major Bioinformatic Resources

2.1 Introduction to Biological Databases	11
2.2 Types of Sequence Databases	12
2.3 Nucleic Acid Sequence Databases	14
2.4 Protein Sequence Databases	16
2.5 Gene Expression Databases	18
2.6 Three-Dimensional Structure Databases	19
2.7 Pattern and Motif Databases	21
2.8 Functional Genomics & Proteomics Databases	23
2.9 File Formats & Retrieval	24

Chapter 3. Basic Concepts of Sequence Analysis

3.1 Why Sequence Analysis Matters	27
3.2 Sequence Identity, Similarity, and Conservation	28
3.3 Homology Concepts: Homologues, Orthologues, Paralogues	32
3.4 Scoring Matrices	34
3.5 Database Search Tools	38
3.6 Pairwise Sequence Alignment	41
3.7 Repeat Finding and Motif Discovery	46
3.8 Multiple Sequence Alignment (MSA)	48
3.9 Application in Taxonomy and Phylogeny	49
3.10 Comparative Genomics	53

Chapter 4. Gene Annotation

4.1 Introduction to Gene Annotation	57
4.2 Gene Prediction Approaches (Structural Annotation)	61
4.3 Functional Annotation — Predicting Gene Function	64

4.4 Genetic Variation — Polymorphism and Deleterious Mutations	69
4.5 Phylogenetics in Gene Annotation	74
4.6 Annotation Pipelines and Practical Considerations	75
Chapter 5. Molecular Modelling and Dynamics	
5.1 Introduction to Molecular Modelling	79
5.2 Fundamentals of Molecular Modelling	81
5.3 3-D Structure Visualisation	84
5.4 Force Fields	88
5.5 Molecular Dynamics Simulation	91
5.6 Applications and Case Studies	94
Chapter 6. Classification & Comparison of Protein 3D Structures	
6.1 Introduction to Protein Structure Analysis	97
6.2 Hierarchical Organisation of Protein Structure	98
6.3 Secondary Structure Prediction	102
6.4 Tertiary Structure Prediction	104
6.5 AI-Based Protein Structure Prediction	107
6.6 Structural Classification and Comparison	110
Chapter 7. Drug Design	
7.1 Introduction to Computational Drug Design	115
7.2 Chemical and Biological Databases for Drug Discovery	117
7.3 Fundamentals of Receptor–Ligand Interactions	118
7.4 Structure-Based Drug Design (SBDD)	120
7.5 Ligand-Based Drug Design (LBDD)	122
7.6 In Silico ADMET Prediction	124
Chapter 8. Systems Biology	
8.1 Introduction to Systems Biology	127
8.2 Data Science in Biology, Health & Drug Discovery	129
8.3 Mathematical Modelling of Biological Pathways	132
8.4 Systems Modelling of Disease Mechanisms	134
8.5 Digital Health and Personalised Medicine	136
8.6 Integrative Systems Biology (Advanced)	139
Appendices	
Appendix A — Glossary of Bioinformatics Terms	143
Appendix B — Important Web Resources	157

CHAPTER

1

Introduction to Bioinformatics

Learning Objectives

1.1 Definition, Scope, and Goals

1.2 Historical Perspective and Evolution

1.3 Interdisciplinary Nature

1.4 Overview of Tools and Techniques

1.5 Emerging Trends

1.6 Importance and Applications

1.7 Ethical, Legal, and Social Implications

Bioinformatics has, in little more than three decades, evolved from a niche activity at the interface of biology and computing into a foundational discipline of the modern life sciences. The complete sequencing of the human genome in 2003, the commercialisation of next-generation sequencing from 2005, and the arrival of deep-learning systems such as AlphaFold in 2021 have, in succession, reshaped how biological questions are formulated and answered. Whereas a biologist in the 1980s might have studied a handful of genes or proteins by direct experiment, a contemporary investigation almost invariably begins with the interrogation of large, publicly available datasets. The skills needed to store, analyse, and interpret those data define the scope of this book.

1.1 Definition, Scope, and Goals of Bioinformatics

Bioinformatics is the scientific discipline that develops and applies computational methods to store, retrieve, analyse, and interpret biological data — most notably DNA, RNA, and protein sequences, together with the structural, functional, and phenotypic information associated with them. The National Institutes of Health (NIH) working definition describes bioinformatics as “research, development, or application of computational tools and approaches for expanding the use of biological, medical, behavioural, or health data, including those to acquire, store, organise, analyse, or visualise such data.” The closely related term computational biology emphasises the use of these computational methods to test hypotheses and build models of biological systems. In practice, the two terms are used interchangeably, with bioinformatics weighted slightly toward data management and computational biology toward model building.

The defining characteristic of the field is the scale and complexity of biological data. A single high-throughput sequencing run can produce tens of billions of short reads; the UniProt Knowledgebase (as of 2025) contains over 250 million protein entries; the AlphaFold Protein Structure Database holds predicted structures for nearly every sequence in UniProt — on the order of 2×10^8 models. Extracting biological insight from datasets of this magnitude requires a principled, computational approach.




Figure 1.1 Bioinformatics as the intersection of biology, computer science, and statistics. The three circles of the classical Venn view are supplemented in modern practice by contributions from the physical sciences and, increasingly, machine learning.

1.1.1 Scope of Bioinformatics

The contemporary scope of bioinformatics extends across the full hierarchy of biological organisation, from individual molecules to whole populations. Seven broad sub-areas are recognised (Figure 1.2):

- **Sequence analysis** — pairwise and multiple alignment, similarity searching, motif discovery, and statistical assessment of sequence relationships.
- **Genome annotation** — identification and functional characterisation of genes, regulatory elements, and non-coding features.
- **Phylogenetics and taxonomy** — inference of evolutionary relationships from molecular data and barcoding of biodiversity.
- **Structural bioinformatics** — prediction, comparison, and classification of biomolecular three-dimensional structures.
- **Molecular modelling and simulation** — computational study of conformational dynamics, ligand binding, and reaction energetics.
- **Systems and network biology** — mathematical description of metabolic, signalling, and regulatory networks.
- **Databases and data sharing** — design, curation, and federated access to biological resources such as GenBank, UniProt, and the PDB.




Figure 1.2 The seven principal sub-areas of bioinformatics.

1.1.2 Goals and the General Workflow

Every bioinformatics investigation — whether identifying a disease gene, predicting a protein fold, or reconstructing the phylogeny of a virus — can be abstracted into a three-stage workflow: (1) data acquisition, (2) modelling or analysis, and (3) validation and biological interpretation (Figure 1.3). In practice, findings from stage 3 frequently motivate new data collection, creating an iterative refinement loop that is characteristic of modern computational science.




Figure 1.3 The general bioinformatics workflow. Data acquisition, analysis, and interpretation are performed iteratively; results from interpretation commonly trigger new rounds of acquisition and modelling.

The primary goals of the field are threefold. First, to organise biological data so that individual observations can be retrieved and compared within a larger corpus. Second, to develop quantitative methods that reveal patterns — sequence homology, conserved motifs, regulatory circuits — which are not visible by inspection of individual records. Third, to generate testable hypotheses about biological function, evolution, and mechanism, thereby guiding subsequent experimentation.

1.2 Historical Perspective and Evolution

Bioinformatics did not emerge as a named discipline until the 1990s, but its intellectual and technical foundations were laid three decades earlier. The field's development can be usefully divided into three eras: the pre-genomic era (roughly 1962–1990), the genomic era (1990–2010), and the post-genomic or data-intensive era (2010–present). Key landmarks are summarised in Figure 1.5 and discussed below.

1.2.1 From Dayhoff's Atlas to the Human Genome Project

The origins of bioinformatics are usually traced to the work of Margaret Oakley Dayhoff (1925–1983), a physical chemist at the National Biomedical Research Foundation who is widely credited as the founder of the field. Together with her collaborators, Dayhoff compiled the first systematic collection of protein sequences — the Atlas of Protein Sequence and Structure — beginning in 1965, and published the Point Accepted Mutation (PAM) series of substitution matrices in 1978. The PAM matrices formalised, for the first time, the idea that evolutionary relationships could be quantified from sequence differences, and they remain the conceptual basis of all modern scoring schemes.

Two further methodological advances from this era are indispensable. In 1970, Saul Needleman and Christian Wunsch described the first dynamic-programming algorithm for global sequence alignment; in 1981, Temple Smith and Michael Waterman extended the approach to local alignment. The Needleman–Wunsch and Smith–Waterman algorithms are taught in every bioinformatics course and form the basis of contemporary tools such as BLAST, FASTA, and HMMER.



Figure 1.4 Margaret Oakley Dayhoff (1925–1983), the founder of bioinformatics. She built the first protein-sequence database and the PAM substitution matrices that underpin modern sequence comparison.

The 1980s saw the construction of the first large public sequence repositories: GenBank at the NIH (1982), the EMBL Data Library (1982), and the DNA Data Bank of Japan (1986). These three now form the International Nucleotide Sequence Database Collaboration (INSDC) and exchange data daily. The Protein Data Bank, although older (1971), was transformed in this period into the authoritative structural database. The decade closed with the launch of the Human Genome Project in 1990, an international consortium that set the stage for the modern field.

1.2.2 The Post-Genomic Era and Big-Data Biology

Publication of the draft human genome in 2001 and the completed reference in 2003 marked the transition to the post-genomic era. Three technological shifts characterise the period that followed. First, next-generation sequencing (NGS) — commercialised by 454 Life Sciences in 2005 and subsequently by Illumina, PacBio, and Oxford Nanopore — reduced the cost of sequencing a human genome from approximately US \$100 million in 2001 to under US \$200 by 2024, a decline far steeper than Moore's law. Second, the field shifted decisively to cloud-native infrastructure; modern analyses of tera- and petabyte-scale datasets are routinely run on Amazon Web Services, Google Cloud, or Microsoft Azure rather than on local clusters. Third, the 2020s have seen the emergence of AI systems that solve long-standing problems in biology: AlphaFold2 (2021) achieved near-experimental accuracy for protein-structure prediction; AlphaFold3 (2024) extended the paradigm to protein–ligand, protein–nucleic-acid, and multi-chain complexes; and large protein language models such as ESM-2 and ESM-3 have begun to generate functional proteins *de novo*.




Figure 1.5 Landmark events in the development of bioinformatics, from Dayhoff's Atlas (1965) to AlphaFold3 and the ESM Atlas (2024). Dates indicate the year of publication or public release.

1.3 Interdisciplinary Nature: Biology, Computer Science, and Statistics

Bioinformatics is, by construction, an interdisciplinary enterprise. The practitioner must integrate concepts and methods from biology, computer science, statistics, and — increasingly — the physical sciences (Figure 1.6). An effective approach rarely requires mastery of all four, but it does require enough fluency in each to collaborate productively with specialists.




Figure 1.6 The three foundational pillars of bioinformatics. Biology defines the questions; computer science provides the algorithmic and data-management infrastructure; statistics and mathematics supply the framework for inference and quantitative modelling. Physics and chemistry contribute physical-model components used in molecular simulation and structural prediction.

Biology supplies the questions. A working knowledge of molecular biology — genes and genomes, the central dogma, protein structure and function, gene regulation, and evolutionary mechanisms — is essential for framing meaningful analyses and avoiding biologically implausible interpretations of computational results.

Computer science provides the algorithmic toolkit: data structures for sequence indexing (suffix arrays, hash tables), algorithmic paradigms (dynamic programming, hidden Markov models, graph algorithms), and systems-level skills in scripting (Python, R), version control, relational and graph databases, and parallel or cloud computing.

Statistics and mathematics underpin the inference that turns raw computation into biological knowledge. Probability theory is central to sequence similarity (E-values, bit scores), hidden Markov models (gene prediction), and Bayesian phylogenetics. Linear algebra and optimisation theory are foundational for machine learning and structural modelling.

Physics and chemistry contribute the physical models used in molecular dynamics and energy-based structure prediction: classical force fields, statistical thermodynamics, quantum-mechanical treatments of reactive sites, and the dielectric properties of solvated biomolecules.

1.4 Overview of Bioinformatics Tools and Techniques

The modern bioinformatics toolkit spans several functional categories. Table 1.1 presents the most widely used representatives, with each tool placed in the context of the chapter where it is treated in detail. The intent here is to provide an orientation map; specific algorithms, parameters, and workflows are developed in subsequent chapters.

Table 1.1 Representative bioinformatics tools by category.

Category	Representative tools	Covered in
Pairwise / database search	BLAST, FASTA, HMMER, DIAMOND	Chapter 3
Multiple sequence alignment	Clustal Omega, MUSCLE, MAFFT, T-Coffee	Chapter 3
Phylogenetics	MEGA, PhyML, IQ-TREE, MrBayes, FigTree	Chapters 3 & 4
Gene prediction / annotation	GENSCAN, Augustus, MAKER, BRAKER, Prokka	Chapter 4
Variant calling & effect prediction	GATK, bcftools, VEP, ANNOVAR, SIFT, PolyPhen-2, CADD	Chapter 4
Structure prediction	MODELLER, SWISS-MODEL, Phyre2, AlphaFold, RoseTTAFold, ESMFold	Chapter 6
Structure visualisation	PyMOL, UCSF ChimeraX, Jmol, VMD	Chapters 5 & 6
Molecular dynamics	GROMACS, NAMD, AMBER, CHARMM	Chapter 5
Docking & virtual screening	AutoDock, AutoDock Vina, Glide, GOLD	Chapter 7
Systems & network analysis	Cytoscape, COBRA, CellDesigner, COPASI	Chapter 8
Programming environments	Python/BioPython, R/Bioconductor, Julia	Throughout

Tools evolve rapidly; URLs and version numbers are intentionally omitted. Students should consult the current documentation before running any analysis.

Major primary databases referenced throughout the book include **GenBank**, **ENA**, **DDBJ** (nucleotide sequences, INSDC members); **UniProtKB** (protein sequences; Swiss-Prot reviewed, TrEMBL unreviewed); **PDB** and **AlphaFold DB** (experimental and predicted three-dimensional structures); and **PubChem**, **ChEMBL**, **DrugBank** (small-molecule bioactivity and pharmacological data). These are catalogued in detail in Chapter 2.

1.5 Emerging Trends

Two developments define the current frontier of the discipline: the rapid integration of deep-learning methods, and the migration of analysis to cloud-native, federated infrastructure.

1.5.1 Artificial Intelligence and Deep Learning

Deep-learning models — particularly transformers, convolutional neural networks, and graph neural networks — now outperform classical methods across a wide range of bioinformatics tasks (Figure 1.7). Protein-structure prediction was transformed in 2021 by AlphaFold2, whose median GDT-TS score on CASP14 targets (92.4) approached the ~90 threshold conventionally regarded as “experimental accuracy.” Language-model approaches such as ESM have extended this capability to proteins with no known homologues and have enabled de novo design of functional enzymes. In genomics, deep learning now achieves state-of-the-art performance on variant-effect prediction (for example, AlphaMissense), splice-site prediction (SpliceAI), and regulatory-element identification (Enformer).

Despite these successes, deep-learning models are not a substitute for biological insight. They can fail silently on distributions that differ from their training data, they often lack calibrated uncertainty, and they can encode biases present in the underlying databases. Best practice combines AI-derived predictions with classical methods and experimental validation.




Figure 1.7 A generalised architecture for AI-driven bioinformatics. Heterogeneous biological inputs are encoded as numerical features, passed through a deep neural network (for example, a transformer or graph neural network), and decoded into task-specific outputs such as predicted structures, variant effect scores, or ranked drug candidates.

1.5.2 Cloud Computing and Federated Data Integration

The volume of biological data now exceeds the storage and compute capacity of typical academic clusters. A single modern whole-genome sequencing project generates 100–300 gigabytes of raw data per sample; the UK Biobank, All of Us, and similar population-scale cohorts have accumulated petabytes. The field has consequently adopted cloud-native infrastructure — Amazon Web Services, Google Cloud, Microsoft Azure — together with container technologies (Docker, Singularity/Apptainer) and workflow managers (Nextflow, Snakemake, Cromwell/WDL) that guarantee reproducibility across sites.

A complementary trend is federated data integration, in which sensitive patient data remain in secure trusted research environments and only the results of pre-authorised computations leave the hosting institution. Standards such as the GA4GH Data Repository Service and Workflow Execution Service are beginning to make federated analyses across hospitals and countries practical. Figure 1.8 illustrates the modern bioinformatics data hierarchy.




Figure 1.8 The bioinformatics data hierarchy. Biological insight emerges as raw data are successively processed, annotated, integrated across modalities, and abstracted into models and mechanistic knowledge.

1.6 Importance and Applications in Modern Biology

Bioinformatics is no longer an auxiliary discipline; it is an essential component of most biological research programmes. Its impact is felt across the life sciences and increasingly in medicine, agriculture, and public health.

- **Medical research and clinical diagnostics.** Bioinformatics has identified disease-associated genes such as BRCA1 and BRCA2 (hereditary breast and ovarian cancer), CFTR (cystic fibrosis), and HTT (Huntington's disease). Clinical exome and whole-genome sequencing pipelines now provide diagnoses for a majority of patients with previously unexplained paediatric disorders.
- **Personalised and precision medicine.** Pharmacogenomic biomarkers — for example CYP2C19 variants affecting clopidogrel metabolism and HLA-B*5701 screening prior to abacavir therapy — are now part of standard clinical practice. Tumour-sequencing pipelines such as those used in the US MSK-IMPACT and UK 100,000 Genomes projects guide selection of targeted therapies.
- **Drug discovery.** Virtual screening and structure-based design now form the first stage of most preclinical drug-discovery pipelines. AI-designed inhibitors of kinases, proteases, and ion channels are in clinical trials; repurposing pipelines identified candidate therapies within weeks of the SARS-CoV-2 genome release in January 2020.
- **Agriculture and food security.** Genome-assisted breeding, CRISPR-based trait design, and metagenomic monitoring of soil microbiomes are reshaping crop and livestock improvement. Indian institutions such as ICAR-NBPGR and the National Agri-Food Biotechnology Institute have applied these techniques to drought-tolerant rice and biofortified crops.
- **Public health and epidemiology.** Real-time genomic surveillance — including Nextstrain, GISAID, and the INSACOG consortium — tracks the spread of SARS-CoV-2, influenza, mpox, and antimicrobial-resistance determinants. Wastewater sequencing has emerged as a population-scale monitoring tool for emerging pathogens.
- **Evolutionary biology and conservation.** Molecular phylogenetics has rewritten the tree of life, most recently reorganising the eukaryotic supergroups and clarifying the origins of modern humans. Conservation genomics uses environmental DNA and population-scale sequencing to monitor endangered species.

1.7 Ethical, Legal, and Social Implications (ELSI)

Because bioinformatics routinely handles personally identifying genomic and clinical data, the field sits at the centre of a broad set of ethical, legal, and social questions. The ELSI programme was established alongside the Human Genome Project in 1990 precisely to address these issues; its remit has expanded with each subsequent technological advance.

1.7.1 Ethical Considerations

Three issues dominate contemporary ethical discussion. Privacy: genomic data are inherently identifying, cannot be truly anonymised, and carry information about biological relatives who have not themselves consented to sharing. Informed consent: participants in genomic studies must understand how their data may be used, including secondary uses not foreseen at the time of sampling. Germline editing: the 2018 births of CRISPR-edited babies in China crystallised a global consensus that heritable human genome editing is premature and, in most jurisdictions, unlawful.

1.7.2 Legal and Regulatory Frameworks

International and national legal instruments govern the handling of genomic and health data. The United States Genetic Information Non-discrimination Act (GINA, 2008) prohibits genetic discrimination in health insurance and employment. The European General Data Protection Regulation (GDPR, 2018) classifies genetic and health data as special-category personal data subject to strict consent and processing requirements. In India, the Digital Personal Data Protection (DPDP) Act, 2023, provides the first comprehensive legal framework for processing of personal data, including genomic and health data, and has direct consequences for academic and commercial bioinformatics work.

1.7.3 Social Dimensions

Equity of access, workforce diversity, and public understanding remain substantial challenges. Reference genomes, disease-association studies, and pharmacogenomic databases continue to be heavily biased toward individuals of European ancestry, with consequences for clinical utility in under-represented populations. Dual-use research — for example, the generation of pathogen sequences with enhanced transmissibility — raises biosecurity concerns that extend beyond academic biology. Responsible practice requires that bioinformaticians participate actively in the ethical, legal, and public-engagement conversations surrounding their work.

Key Concepts to Remember

- Bioinformatics is the computational discipline that organises, analyses, and interprets biological data, with primary emphasis on sequence, structure, and expression data.
- The field is defined by the scale of modern biological datasets ($>10^{12}$ bases of sequence in public archives; $>2 \times 10^8$ predicted protein structures) and cannot be effectively practised without computational tools.
- Its seven classical sub-areas — sequence analysis, annotation, phylogenetics, structural bioinformatics, molecular modelling, systems biology, and databases — correspond to Chapters 2–8
- The modern history of the field spans three eras: pre-genomic (Dayhoff, PAM matrices, Needleman–Wunsch, GenBank), genomic (BLAST, HGP, NGS), and post-genomic (cloud computing, AlphaFold, protein language models).
- Bioinformatics is inherently interdisciplinary, combining biology (the questions), computer science (the tools), statistics and mathematics (the inference), and physics and chemistry (the models).
- Current frontiers include AI-driven prediction and design (AlphaFold2/3, RoseTTAFold, ESM), cloud-native and federated analysis, and multi-omics integration.
- The field is inseparable from the ethical, legal, and social questions raised by the handling of genomic and health data, and practitioners are expected to engage with regulatory frameworks such as GINA, GDPR, and the Indian DPDP Act, 2023.

Further Reading

- Baxevanis, A. D., Bader, G. D., & Wishart, D. S. (eds.) (2020). *Bioinformatics*, 4th ed. Wiley. — Comprehensive reference work.
- NCBI Handbook (online, 3rd ed.). — Freely available at <https://www.ncbi.nlm.nih.gov/books/NBK143764/>
- EMBL-EBI Training portal. — Freely available self-paced courses at <https://www.ebi.ac.uk/training>.
- Gabriel, L., et al. (2024). BRAKER3: fully automated genome annotation using RNA-seq and protein evidence with GeneMark-ETP, AUGUSTUS, and TSEBRA. *Genome Research* 34:769–777. — Current state of the art in open-source eukaryotic annotation.
- Hoff, K. J., and Stanke, M. (2019). Predicting genes in single genomes with AUGUSTUS. *Current Protocols in Bioinformatics* 65:e57. — Practical step-by-step guide.
- Seemann, T. (2014). Prokka: rapid prokaryotic genome annotation. *Bioinformatics* 30:2068–2069. — The definitive bacterial pipeline.