



Research Techniques Made Simple: Bioinformatics for Genome-Scale Biology

Amy C. Foulkes¹, David S. Watson², Christopher E.M. Griffiths¹, Richard B. Warren¹, Wolfgang Huber³ and Michael R. Barnes²

High-throughput biology presents unique opportunities and challenges for dermatological research. Drawing on a small handful of exemplary studies, we review some of the major lessons of these new technologies. We caution against several common errors and introduce helpful statistical concepts that may be unfamiliar to researchers without experience in bioinformatics. We recommend specific software tools that can aid dermatologists at varying levels of computational literacy, including platforms with command line and graphical user interfaces. The future of dermatology lies in integrative research, in which clinicians, laboratory scientists, and data analysts come together to plan, execute, and publish their work in open forums that promote critical discussion and reproducibility. In this article, we offer guidelines that we hope will steer researchers toward best practices for this new and dynamic era of data intensive dermatology.

Journal of Investigative Dermatology (2017) 137, e163–e168; doi:10.1016/j.jid.2017.07.095

CME Activity Dates: 21 August 2017

Expiration Date: 20 August 2018

Estimated Time to Complete: 1 hour

Planning Committee/Speaker Disclosure: Amy Foulkes is a consultant/advisor for AbbVie, Almiral, Eli Lilly, Leo Pharma, Novartis, Pfizer, Janssen and UCB. Christopher Griffiths is on the speakers' bureau and is a consultant/advisor for AbbVie, GSK, Janssen, Pfizer, Lilly, Novartis, Celgene, Leo Pharma, UCB, Sun Pharmaceuticals, and Almiral; in addition, Dr. Griffiths receives research grant support from AbbVie, GSK, Janssen, Pfizer, Lilly, Novartis, Sandoz, Celgene, and Leo Pharma. All other authors, planning committee members, CME committee members and staff involved with this activity as content validation reviewers have no financial relationships with commercial interests to disclose relative to the content of this CME activity.

Commercial Support Acknowledgment: This CME activity is supported by an educational grant from Lilly USA, LLC.

Description: This article, designed for dermatologists, residents, fellows, and related healthcare providers, seeks to reduce the growing divide between dermatology clinical practice and the basic science/current research methodologies on which many diagnostic and therapeutic advances are built.

Objectives: At the conclusion of this activity, learners should be better able to:

- Recognize the newest techniques in biomedical research.

- Describe how these techniques can be utilized and their limitations.
- Describe the potential impact of these techniques.

CME Accreditation and Credit Designation: This activity has been planned and implemented in accordance with the accreditation requirements and policies of the Accreditation Council for Continuing Medical Education through the joint providership of William Beaumont Hospital and the Society for Investigative Dermatology. William Beaumont Hospital is accredited by the ACCME to provide continuing medical education for physicians.

William Beaumont Hospital designates this enduring material for a maximum of 1.0 *AMA PRA Category 1 Credit(s)*[™]. Physicians should claim only the credit commensurate with the extent of their participation in the activity.

Method of Physician Participation in Learning Process: The content can be read from the *Journal of Investigative Dermatology* website: <http://www.jidonline.org/current>. Tests for CME credits may only be submitted online at <https://beaumont.cloud-cme.com/RTMS-Sept17> – click 'CME on Demand' and locate the article to complete the test. Fax or other copies will not be accepted. To receive credits, learners must review the CME accreditation information; view the entire article, complete the post-test with a minimum performance level of 60%; and complete the online evaluation form in order to claim CME credit. The CME credit code for this activity is: 21310. For questions about CME credit email cme@beaumont.edu.

INTRODUCTION

Modern dermatology has been revolutionized by the many so-called 'omic' profiling platforms enabled by high-throughput sequencing (HTS, also referred to as next-

generation sequencing). Plunging data generation costs have enabled dermatology researchers to generate genome scale data relating to genome sequence variation (Scott et al., 2013), epigenomes (Zhou et al., 2016), and transcriptomes

¹The Dermatology Centre, Salford Royal NHS Foundation Trust, The University of Manchester, Manchester Academic Health Science Centre, Manchester, UK;

²William Harvey Research Institute, Centre for Translational Bioinformatics, Barts and The London School of Medicine and Dentistry, Charterhouse Square, London, UK; and ³European Molecular Biology Laboratory, Heidelberg, Germany

Correspondence: A.C. Foulkes, NIHR Academic Clinical Lecturer in Dermatology, The Dermatology Centre, Salford Royal NHS Foundation Trust, The University of Manchester, Manchester Academic Health Science Centre, M6 8HD. E-mail: Amy.foulkes@manchester.ac.uk

Abbreviations: HTS, high-throughput sequencing; RNA-seq, RNA sequencing

ADVANTAGES

- Bioinformatics methods allow efficient and powerful analysis of multi-omic data in a way that could not be achieved using simpler methods.
- Bioinformatics software are customizable to all ranges of computational ability; however, some informatics tasks are difficult and require experience.
- Involving bioinformatician colleagues from project conception should improve project design, maximizing the opportunity to detect relevant association.
- Sharing data, metadata, and code, and propagating the culture of bioinformaticians, will fuel best practices in dermatology research, promoting open research and reproducibility.

LIMITATIONS

- Some statistical analysis methods require an understanding of underlying assumptions—erroneous assumptions can lead to false results.
- The use of some analytical pipelines requires access to high-performance computing facilities: this may be achieved by access to omic core facilities that provide researchers with compressed datasets that are amenable to computer-based analysis.

(Li et al., 2014; Swindell et al., 2016), and these developments have increased the dermatology-relevant data openly available in repositories (Table 1).

Bioinformatics refers to the tools used to collect, classify, and analyze such datasets, collectively enabling the field of *computational biology*. Bioinformatics techniques have been developed to make sense of the output of omic platforms, including HTS, microarrays, liquid chromatography-mass spectrometry, and others (Kimball et al., 2012).

Table 1. High-throughput sequencing repositories

Repository	Website	Curator
Europe		
European Nucleotide Archive (ENA)	http://www.ebi.ac.uk/ena	European Bioinformatics Institute
ArrayExpress	http://www.ebi.ac.uk/arrayexpress	European Bioinformatics Institute
European Genome-phenome Archive (EGA)	https://www.ebi.ac.uk/ega/home	European Bioinformatics Institute
United States		
dbGAP	https://www.ncbi.nlm.nih.gov/gap	The National Center for Biotechnology Information
Gene Expression Omnibus (GEO)	https://www.ncbi.nlm.nih.gov/geo	The National Center for Biotechnology Information
Short Read Archive (SRA)	https://www.ncbi.nlm.nih.gov/sra	The National Center for Biotechnology Information

Physicians are key instigators of research data collection requiring computational biology. Structured and validated analysis pipelines for most omic data have been implemented for researchers at various levels of complexity. Software has been designed for all ranges of computational ability, from simple “point and click” graphic user interfaces to highly customizable command line interfaces, with the latter approach offering superior flexibility and analytical complexity. Although programming may seem like a daunting challenge for those without backgrounds in math, computer science, or statistics, with practice, computational methods for exploratory and inferential analytics can become a familiar part of the research toolkit. Of course, there is no substitute for expertise, and we advise all research teams working with omic data to consult a bioinformatician early and often. Here we highlight several points of special relevance to the dermatologist and dermatology researcher, based on the first-hand experience of a junior clinician.

CONSIDERATIONS BEFORE DATA COLLECTION

Experimental Design

Researchers in dermatology use a wide variety of HTS techniques, many of which have been discussed previously in the Research Techniques Made Simple series. These include transcriptome analysis with RNA sequencing (RNA-seq) (Antonini et al., 2017; Whitley et al., 2016), immunosequencing (Matos et al., 2017), genome-wide epigenetics (Capell and Berger, 2013), proteomics, metabolomics, metagenomics, and assessment of the microbiome (Jo et al., 2016). Additionally, the Molecular Revolution in Cutaneous Biology series provided an overview of HTS techniques (Anbunathan and Bowcock, 2017; Botchkareva, 2017; Johnston et al., 2017; Kong and Segre, 2017; Sarig et al., 2017), as did Grada and Weinbrecht (2013) in an earlier Research Techniques Made Simple publication. However, researchers often do not reach out to data analysts until a study is practically complete. At that point, they may look for a mathematically inclined colleague to fill in the blanks of a statistical model and provide a friendly *P*-value suitable for publication. This order of events is all wrong. As Ronald Fisher famously put it back in 1938, “To consult the statistician after an experiment is finished is often merely to ask him to conduct a post-mortem examination. He can perhaps say what the experiment died of” (Fisher, 1938).

The data analysis strategy, including the choice of statistical approaches, should be integral to planning any research study. Hypothesis testing, regression, and other statistical methods rely on rigorous collection and quality of the data, and any lapses here usually cannot be fixed retrospectively. How many samples are required to adequately power your experiment? If samples cannot be processed all at once, does it matter how they are grouped into separate batches? If the data do not corroborate your hypothesis, can a modified research question generate interesting results? Failure to consider these questions before data collection may doom a study before it even begins. Statistical expertise is required to answer these questions, which is why we urge researchers to team up with a data analyst who can help guide them through these tricky issues. This will typically either be a statistician, with a background in math and statistics, or a

Table 2. Principles of hypothesis testing from Li et al. (2014)

Step in Hypothesis Testing	Example
Ask a clinically relevant, testable question	Is there a significant difference between this set of genes expressed in subjects with psoriasis versus those without?
Choose an experimental design and statistical framework	Gene expression is modeled as a linear function of disease condition
Set up a null hypothesis, that is, a testable claim that becomes the target of statistical analysis	There is no significant difference between the average expression of gene <i>g</i> in subjects with and without psoriasis
Fix a rejection region, that is, the degree of evidence against the null hypothesis at which it may be rejected	Genes whose <i>t</i> -statistics correspond to false discovery rates $\leq 5\%$ are declared differentially expressed
Conduct the experiment: collect data, compute the test statistics	Expression levels for each gene <i>g_i</i> are regressed onto one or several clinical predictors, generating a vector of <i>t</i> -statistics
Report results: all and only those genes that fall within the rejection region are declared differentially expressed	A number of genes were significantly differentially expressed in plaques of psoriasis when compared with control samples

bioinformatician, more likely with a background in computer science and machine learning. Although there is considerable overlap in their respective areas of expertise, statisticians and bioinformaticians may offer differing (and sometimes complementary) perspectives on a given biological question.

One of the most fundamental tools in statistical analysis is hypothesis testing. The principles of hypothesis testing are illustrated in Table 2, which highlights the work of Li et al. (2014) as an exemplar study in the field (see [Supplementary Slides](#) online). In this exploratory study, RNA-seq was used to evaluate the transcriptomes of lesional psoriatic and

normal skin (from a large cohort of 174 individuals). A subset of these samples has been studied previously using microarrays, allowing for comparison of the methodologies; RNA-seq identified many more differentially expressed transcripts enriched in immune system processes.

Detailed discussion of requirements for testing a hypothesis will facilitate better downstream clinical data collection, ultimately maximizing the opportunity to detect a clinically relevant association. Several key themes tend to dominate experimental design considerations, including selection of appropriate numbers of biological replicates (Schurch et al., 2016), minimization of batch effects (Leek et al., 2010), and appropriate correction for multiple testing (Allison et al., 2006). For a general overview of issues related to HTS study design, we recommend other excellent reviews (Allison et al., 2006; Conesa et al., 2016).

The steps outlined in Table 1 apply to most forms of omic data. Methods for computing test statistics vary depending on the data and underlying statistical assumptions. Common data types and test statistics used in dermatological research are discussed elsewhere (Silverberg, 2015).

Batch Effects

Often a study's sample size exceeds the maximum number of samples that can be simultaneously processed by the available equipment. In such cases, it is common to process the samples in multiple batches. This inevitably introduces batch effects, in which technical artifacts become significant, perhaps even dominant drivers of variation in a dataset. There are several methods for batch adjustment (Oytam et al., 2016).

Each method has its merits, but none can overcome poor study design. If a batch is confounded with a clinical covariate—say, all disease samples were processed in Batch A, and all healthy samples were processed in Batch B—then there is no way to disentangle the technical from the biological variation. Ideally, each batch would represent a microcosm of

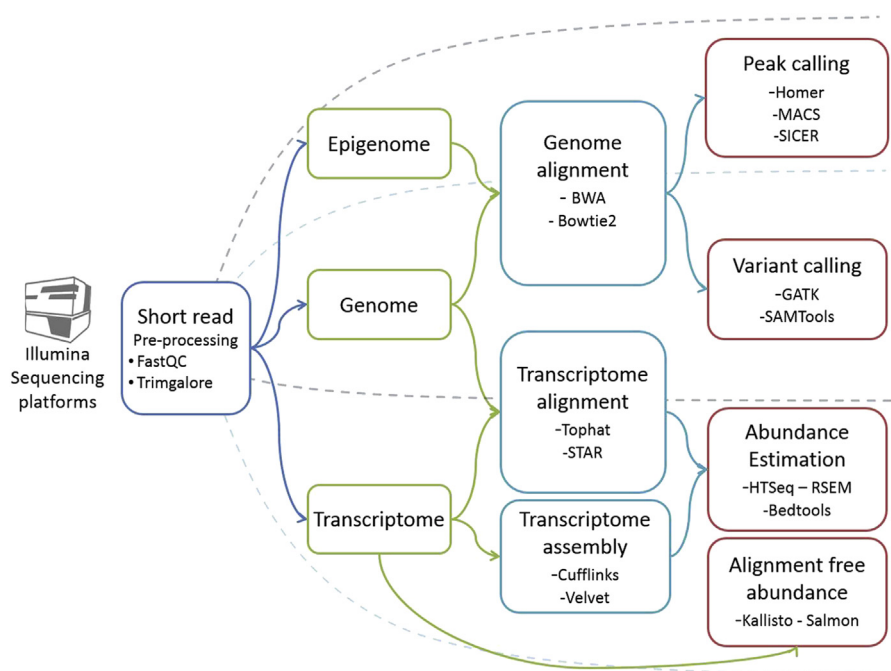


Figure 1. Common methodology for processing of short reads.

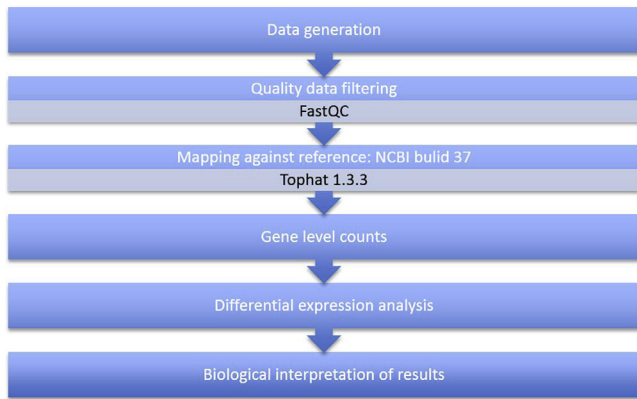


Figure 2. Example bioinformatic pipeline used by Li et al. (2014).

the experiment itself, with proportionate numbers of samples from all relevant groups. Although this cannot always be done in practice, the closer researchers come to attaining this goal, the more accurate their results will be.

CONSIDERATIONS AFTER DATA COLLECTION

Software and Workflows for Omic Analysis

As a rule of thumb, processing of raw HTS data, including genome alignment and assembly, is likely to require access to one or several devoted computers that can execute jobs in parallel. However, once the initial data processing is complete, in most cases the biological downstream analysis can be performed using a laptop. The analysis of omic data, including HTS, is supported by a range of widely used software packages that can be arranged into analysis workflows. Many packages have been made freely available by their authors with an open source license, and in this field there is very little correlation between the price of software and its usefulness. A workflow is a software pipeline that takes raw data as input, transforms and summarizes the data, conducts exploratory and/or inferential analytics, and exports results ready for biological interpretation. Command line genomic analysis tools can be scaled to use available computing resources and are highly customizable to meet the requirements of an experiment. Many standard analysis tools can also be accessed remotely using the Galaxy workflow environment (<https://usegalaxy.org>). Galaxy offers users a simple but highly customizable graphic user interface environment to perform many bioinformatics tasks. Galaxy is also well documented and serves as an excellent introduction to HTS analysis pipelines.

Processing HTS Data

The short read is the common currency of HTS methods, but the way the read is processed is highly dependent on the analysis objective (Figure 1). In most cases processing commences with alignment to a reference genome using a tool, such as Burrows-Wheeler Aligner or bowtie2, producing binary alignment map files. The alignment files can serve as the input to many other processes; in genetics they are used for variant calling, in epigenomics for peak calling, and in transcriptomics to estimate transcript abundance. A recent revolution in transcriptomics is alignment-free mapping methods, such as Kallisto (Bray et al., 2016) and Salmon (Patro et al., 2017). These tools circumvent the cumbersome alignment step and directly estimate transcript abundance; they are several orders of magnitude faster than

Table 3. Open source programming languages and resources for bioinformatics analysis of omic data

Open Source Resource	URL
Analysis code repositories	
Bioconductor	bioconductor.org
CRAN	www.cran.org
Bioperl	bioperl.org
Biopython	biopython.github.io
GitHub	github.com
BioJulia	github.com/BioJulia
Workflow tools	
Galaxy	usegalaxy.org
Visualization	
ShinyR	Shiny.rstudio.com
Plotly	plot.ly

alignment-based methods and so computationally efficient that they can be run on a laptop computer. The workflow used by Li et al. (2014) is illustrated in Figure 2.

Programming Environments

Although many programming environments are used in bioinformatics, the most popular choices tend to be R (R Core Team, 2014) and Python (Python Software Foundation, 2013). Software packages for these languages are often released under open source licenses, which means the tools are free to use and the code is publicly accessible. Large user communities have developed around these languages, and R in particular has become a lingua franca for bioinformaticians. This has been aided in no small part by the Bioconductor project (Huber et al., 2015), a major repository for biostatistical software based primarily on R. The site also hosts discussion forums, encouraging active user engagement and collaborative learning.

Several programming environments are widely used in bioinformatics, including R, Matlab (Mathworks, 2012), and Java (see Table 3). These are open source and freely available, enabling statistical and graphical data manipulation within large, active user communities.

Hypothesis Testing in the Age of Big Data

Hypothesis tests and *P*-values are the workhorses of medical research, but some additional complexities enter the scene when we do not perform only one or a few tests, but thousands or millions. Interpreting *P*-values is quite different in omic contexts than in more traditional low-throughput research. Say you test 10,000 genes in search of biomarkers to distinguish between case and control samples. You find 500 with *P*-values below 0.05, not to mention 10 with *p*-values below 0.001. Not bad, right? Wrong! Because *P*-values are uniformly distributed under the null hypothesis, we should expect 5% of all tests to reach the nominal significance level of 0.05 by chance alone. That's a manageable problem when testing one or two hypotheses, but in omic experiments we typically test something on the order of thousands to millions of hypotheses.

Some early articles attempted to mitigate the issue by controlling the family-wise error rate, defined as the probability of finding at least one false positive in a series of hypothesis tests. For example, the Bonferroni correction used by

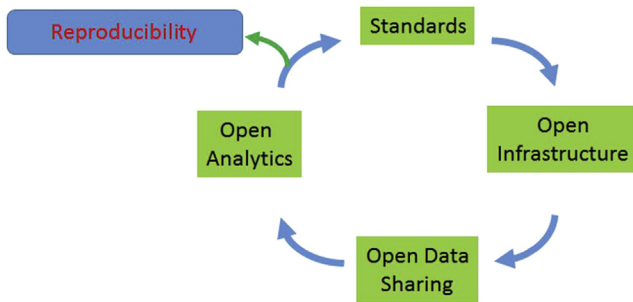


Figure 3. Reproducibility: Creating a virtuous circle.

Li et al. (2014) strongly controls the family-wise error rate by setting the significance threshold as the quotient of the type I error α and the total number of hypothesis tests m , so that all and only tests with $P \leq \alpha/m$ are declared significant. Although the Bonferroni correction is guaranteed to control the family-wise error rate, it is an overly conservative method that is likely to lead to many false negatives as m grows.

Current practice is to control not the *false positive rate* (i.e., the proportion of truly null features that are nominally significant) but the *false discovery rate* (i.e., the proportion of nominally significant features that are truly null). This latter value is typically estimated using the Benjamini-Hochberg algorithm (Hochberg and Benjamini, 1990) or some variant thereof. This method takes a list of P -values as input and returns a matched list of adjusted P -values, also known as Q -values. Applying a 5% false discovery rate threshold means that 1 in 20 genes in the hit list will be a false positive. Given 10,000 uniformly distributed P -values, as hypothesized earlier, minimum Q -values are typically greater than 0.5.

VISUALIZATION

The communication of results is key for data exploration, summarization, and ultimately publication. Readers can more readily absorb a well-made graphic than any table of numbers. Visualizing HTS results can be challenging because of the data's high dimensionality, but projection techniques like principal component analysis (Pearson, 1901), multi-dimensional scaling (Torgerson, 1952), and t -distributed stochastic neighbor embedding (van der Maaten et al., 2008) can render large matrices as easily digestible two-dimensional or three-dimensional scatterplots. Matos et al. (2017) show how these methods can give powerful insights for dermatological research. More recent interactive tools such as plotly (<https://plot.ly/>), shiny (<https://shiny.rstudio.com/>), and ggvis (<http://ggvis.rstudio.com/>) can also aid in data exploration or even create widgets for HTML publication.

CODE SHARING AND REPRODUCIBILITY

A number of studies have found an alarming lack of reproducibility in modern omic and clinical research (Open Science Collaboration, 2015). Many factors contribute to this problem, including the widespread failure to publish analysis code (Baker, 2016). Although some inroads have been made toward establishing best practices in molecular biology (Brazma et al., 2001), script sharing remains rare overall. Results may vary greatly depending on subtle, unstated analytic choices that are invisible without access to both raw data and the complete analysis script. Code sharing

MULTIPLE CHOICE QUESTIONS

- Which is an accurate description of *batch effect*?
 - Technical source of variation added to samples during handling
 - An uncommon problem in HTS experiments
 - Where proportionate samples are analyzed in each experiment
 - A problem that is not possible to adjust for using bioinformatic techniques
- The relevant significance measure in omic data is
 - the P -value.
 - the false discovery rate.
 - the false positive rate.
 - the family-wise error rate.
- Which of the following is an analysis code repository?
 - GEO
 - R
 - Galaxy
 - GitHub
- Which of the following statements is true regarding sharing of analysis code?
 - This allows reproducibility of an analysis.
 - Sharing of analysis code is technically challenging.
 - Analysis code is required alongside submission of data and metadata for submission of original articles to major journals.
 - There is no code sharing repository.
- Which of the following is a major repository for biostatistical software?
 - ShinyR
 - Plotly
 - Ggvis
 - Bioconductor

is a critical ingredient for open science; this will be apparent to researchers who have tried to reuse data in repositories, where code is absent and subject data are often incomplete, making reproduction challenging at best. Excellent platforms exist for publishing code. Taking advantage of sites like GitHub (<https://github.com>) can assist during peer review, enabling precise debate on the merits of particular methods. Set-up can be technically challenging, but user-friendly guides exist (<http://happygitwithr.com>). Researchers should ensure that they or their bioinformatician colleagues document and archive code, analogous to the use of a laboratory book as a record of research. This will ensure that bioinformatician turnover will not prevent ongoing analysis, because code will be clear, maintained, and transferable.

SUMMARY AND FUTURE DIRECTIONS

Embedding biostatisticians and computational biologists within clinical and academic research teams, as well as promoting better data and code sharing practices, will allow dermatologists to better document and communicate their research. The days of assembly line research—in which clinicians recruit patients, laboratory scientists process samples, and analysts crunch numbers—are coming to an end. The age of big data demands a rigorous, integrated approach. Appropriate statistical design and analysis methods should be discussed and decided on up front to meet most research objectives. By incorporating good experimental design and analytical work practice early, research quality and reproducibility will improve, and peer review by journals and grant awarding bodies is likely to be more favorable (Figure 3). Patients will be the ultimate beneficiaries of dermatology's drive to the forefront of life science research.

ORCID

Amy C Foulkes: <http://orcid.org/0000-0003-2680-750X>

CONFLICT OF INTEREST

ACF has received educational support to attend conferences from or acted as a consultant or speaker for Abbvie, Amgen, Almirall, Eli Lilly, Leo Pharma, Novartis, Pfizer, Janssen, and UCB. CEMG has acted as a consultant and/or speaker for Abbvie, Janssen, Novartis, Sandoz, Rock Creek Pharma, Pfizer, Eli Lilly, UCB, Leo Pharma, Galderma, and Celgene. RBW has acted as a consultant and/or speaker for Abbvie, Amgen, Almirall, Boehringer, Medac, Eli Lilly, Janssen, Leo Pharma, Pfizer, Novartis, Sun Pharma, Valeant, Schering-Plough (now MSD), and Xenoport.

ACKNOWLEDGMENTS

This forms part of the research themes contributing to the translational research portfolio of Barts and the London Cardiovascular Biomedical Research Centre, which is supported and funded by the National Institute of Health Research.

SUPPLEMENTARY MATERIAL

Supplementary material is linked to this paper. Teaching slides are available as supplementary material.

REFERENCES

Allison DB, Cui X, Page GP, Sabripour M. Microarray data analysis: from disarray to consolidation and consensus. *Nat Rev Genet* 2006;7:55–65.

Anbunathan H, Bowcock AM. The molecular revolution in cutaneous biology: the era of genome-wide association studies and statistical, big data, and computational topics. *J Invest Dermatol* 2017;137:e113–8.

Antonini D, Mollo MR, Missero C. Research techniques made simple: identification and characterization of long noncoding RNA in dermatological research. *J Invest Dermatol* 2017;137(3):e21–6.

Baker M. 1,500 scientists lift the lid on reproducibility. *Nature* 2016;533(7604):452–4.

Botchkareva NV. The molecular revolution in cutaneous biology: noncoding RNAs: new molecular players in dermatology and cutaneous biology. *J Invest Dermatol* 2017;137:e105–11.

Bray NL, Pimentel H, Melsted P, Pachter L. Near-optimal probabilistic RNA-seq quantification. *Nat Biotechnol* 2016;34:525–7.

Brazma A, Hingamp P, Quackenbush J, Sherlock G, Spellman P, Stoeckert C, et al. Minimum information about a microarray experiment (MIAME)-toward standards for microarray data. *Nat Genet* 2001;29:365–71.

Capell BC, Berger SL. Genome-wide epigenetics. *J Invest Dermatol* 2013;133:e9.

Conesa A, Madrigal P, Tarazona S, Gomez-Cabrero D, Cervera A, McPherson A, et al. A survey of best practices for RNA-seq data analysis. *Genome Biol* 2016;17:13.

Fisher RA. Presidential address to the first Indian Statistical Congress. *Sankhya* 1938;4:14–7.

Grada A, Weinbrecht K. Next-generation sequencing: methodology and application. *J Invest Dermatol* 2013;133:e11.

Hochberg Y, Benjamini Y. More powerful procedures for multiple significance testing. *Stat Med* 1990;9:811–8.

Huber W, Carey VJ, Gentleman R, Anders S, Carlson M, Carvalho BS, et al. Orchestrating high-throughput genomic analysis with Bioconductor. *Nature Methods* 2015;12:115–21.

Jo JH, Kennedy EA, Kong HH. Research techniques made simple: bacterial 16S ribosomal RNA gene sequencing in cutaneous research. *J Invest Dermatol* 2016;136:e23–7.

Johnston A, Sarkar MK, Vrana A, Tsoi LC, Gudjonsson JE. The molecular revolution in cutaneous biology: the era of global transcriptional analysis. *J Invest Dermatol* 2017;137:e87–91.

Kimball AB, Grant RA, Wang F, Osborne R, Tiesman JP. Beyond the blot: cutting edge tools for genomics, proteomics and metabolomics analyses and previous successes. *Br J Dermatol* 2012;166(Suppl. 2):1–8.

Kong HH, Segre JA. The molecular revolution in cutaneous biology: investigating the skin microbiome. *J Invest Dermatol* 2017;137(5):e119–22.

Leek JT, Scharpf RB, Bravo HC, Simcha D, Langmead B, Johnson WE, et al. Tackling the widespread and critical impact of batch effects in high-throughput data. *Nat Rev Genet* 2010;11:733–9.

Li B, Tsoi LC, Swindell WR, Gudjonsson JE, Tejasvi T, Johnston A, et al. Transcriptome analysis of psoriasis in a large case-control sample: RNA-seq provides insights into disease mechanisms. *J Invest Dermatol* 2014;134:1828–38.

MathWorks. MATLAB and Statistics Toolbox Release. Boston, MA; 2012.

Matos TR, de Rie MA, Teunissen MBM. Research techniques made simple: high-throughput sequencing of the T-cell receptor. *J Invest Dermatol* 2017;137:e131–8.

Open Science Collaboration. PSYCHOLOGY. Estimating the reproducibility of psychological science. *Science* 2015;349(6251):aac4716.

Oytam Y, Sobhanmanesh F, Duesing K, Bowden JC, Osmond-McLeod M, Ross J. Risk-conscious correction of batch effects: maximising information extraction from high-throughput genomic datasets. *BMC Bioinformatics* 2016;17:332.

Patro R, Duggal G, Love MI, Irizarry RA, Kingsford C. Salmon provides fast and bias-aware quantification of transcript expression. *Nat Methods* 2017;14:417–9.

Pearson K. On lines and planes of closest fit to systems of points in space. *Philosophical Magazine* 1901;2:559–72.

Python Software Foundation. Python Language Reference. Wilmington, DE; 2013.

R Core Team. R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing; 2014. <http://www.R-project.org/>. Accessed 2 August 2017.

Sarig O, Sprecher E. The molecular revolution in cutaneous biology: era of next-generation sequencing. *J Invest Dermatol* 2017;137:e79–82.

Schurch NJ, Schofield P, Gierlinski M, Cole C, Sherstnev A, Singh V, et al. How many biological replicates are needed in an RNA-seq experiment and which differential expression tool should you use? *RNA* 2016;22:839–51.

Scott CA, Pagnon V, Nitoiu D, Bland PJ, Blaydon DC, Chronnell CM, et al. Targeted sequence capture and high-throughput sequencing in the molecular diagnosis of ichthyosis and other skin diseases. *J Invest Dermatol* 2013;133:573–6.

Silverberg JL. Study designs in dermatology: practical applications of study designs and their statistics in dermatology. *J Am Acad Dermatol* 2015;73:733–40.

Swindell WR, Sarkar MK, Liang Y, Xing X, Gudjonsson JE. Cross-disease transcriptomics: unique IL-17A signaling in psoriasis lesions and an autoimmune PBMC signature. *J Invest Dermatol* 2016;136:1820–30.

Torgerson WS. Multidimensional scaling: I. Theory and method. *Psychometrika* 1952;17:401–19.

Whitley SK, Horne WT, Kolls JK. Research techniques made simple: methodology and clinical applications of RNA sequencing. *J Invest Dermatol* 2016;136:e77–82.

van der Maaten L, Hinton G, van der Maaten GH. Visualizing Data using t-SNE. *Journal of Machine Learning Research* 2008;9:2579–605.

Zhou F, Wang W, Shen C, Li H, Zuo X, Zheng X, et al. Epigenome-wide association analysis identified nine skin DNA methylation loci for psoriasis. *J Invest Dermatol* 2016;136:779–87.