The Biomedical Research Pyramid: A Model for the Practice of Biostatistics

Jesse D. Troy[1,*], Megan L. Neely[1], Steven C. Grambow[1] & Gregory P. Samsa[1]

[1]Department of Biostatistics and Bioinformatics, Duke University School of Medicine, Durham, NC, USA

*Correspondence: Department of Biostatistics and Bioinformatics, Duke University School of Medicine, Durham, NC, USA. E-mail: jesse.troy@duke.edu

Abstract

Biostatisticians apply statistical methods to solve problems in the biological sciences. Successful practioners of biostatistics have advanced technical knowledge, are skilled communicators, and can seamlesslessly integrate with interdisciplinary scientific teams. Despite the breadth of skills required for success in this field, most biostatistics education programs place heavier emphasis on development of technical skills than skills necessary for collaborative work, including critical thinking, writing, and public speaking. Our master's degree program in biostatistics aims for stronger integration of education in collaborative work alongside development of technical knowledge in biostatistics. Toward that end, we propose a model that provides students with a mental map for practicing biostatistics, and that can serve as a tool for faculty to create hands-on learning experiences for biostatistics students. The model helps students organize their knowledge of biostatistics, unifying the technical and collaborative aspects of the discipline in a single framework that can be applied across the broad array of activities that biostatisticians engage in. In this article we describe the model in detail and provide an initial assessment of whether the model might meet its intended purpose by applying the model to a common task for practicing biostatisticians and biostatistics students: describing the results of a medical research study.

Keywords: collaborative biostatistics, curriculum design, statistics education, statistical consulting

1. Introduction

Biostatistics is the application of statistical methods to solve problems in the biological sciences and is inherently a collaborative discipline (Samsa, 2018). Our master's program in biostatistics prepares students to work in an interdisciplinary, collaborative environment. The program emphasizes simultaneous development of skills in communication and leadership with rigorous training in biostatistical methods, with the objective of placing the work of the biostatistician in the context of the larger pursuit of science. For example, elsewhere we've described an explicit model of how an interaction between a statistician and another investigator should operate, and also described some work products (e.g., table shells, data set shells) that can assist in this communication (Samsa et al., 2018).

Our approach to biostatistical education emphasizes development of traits that are representative of the successful collaborative biostatistician, including excellent oral and written communication skills, and the ability to translate research problems into statistical language and give appropriate advice to collaborators (Zapf et al., 2019). Our program also emphasizes biology and offers experience working in a scientific team (Begg & Vaughan, 2011). Historically, in our program as well as others, much of this education has been accomplished in apprenticeship-type settings rather than collaboration-focused coursework (Belli, 2001). Unfortunately, we find few recent example implementations of such courses for biostatistics graduate students and those we have found, while resting on solid pegagogical foundation, do not emphasisze a theory of biostatistical practice itself (Davidson et al., 2019). In this article we propose that such a theoretical view of biostatistical practice can enable classroom-based instruction in the "softer skills" of biostatistics.

We present here an innovative and flexible approach for organizing educational curricula for collaborative biostatisticians. Fundamentally, our approach is to begin by modeling the practice of biostatistics itself. As we

demonstrate, this approach has the dual benefit of serving as both a pedagogical tool and a guide for practitioners. Our model, which we call The Biomedical Research Pyramid, was developed to meet our objective of training collaborative biostatisticians who can effectively integrate biostatistics into the scientific process and, following a constructivist approach (Biggs, 1996), to facilitate creation of experiential learning opportunities. Importantly, the model is intended to aid students in organizing their thinking as they become lifelong learners in biostatistics. This is consistent with the constructivist approach in that it emphasizes making connections between new ideas and existing knowledge in an organized fashion (Biggs, 1996). We refer to this in a general sense as "making a mental map." Others have proposed approaches, like concept mapping, to actually draw such maps (Novak & Cañas, 2007). Our model does not imply the use of any specific approach *per se*. Rather, our model is structured in such a way that it enables visualizing the required mental map and, by extension, creation of concrete educational activities.

In this article, we describe The Biomedical Research Pyramid and assess its potential utility for our intended purpose by evaluating the application of the model to a common task for applied biostatisticians and teachers of biostatistics, which is the interpretation of a peer-reviewed medical research article. Our intent in this article is to describe the model and communicate the concepts behind our approach. Evaluation of efficacy will be the subject of a future article.

2. Methods

2.1 Research Question and Objectives

Our research question focused on whether a conceptual model could be developed to guide the applied biostatistician in their thinking about practicing biostatistics. Our primary objective for this model was to provide students with a "big picture" to help orient them to the overall research process and to serve as an anchor point for lower-level mental maps of its individual elements. Our secondary objective was to develop a model that has dual use as a pedagogical tool and an applied tool. In other words, we sought a model that would enable us to create learning experiences and that would also serve as a tool that students could carry with them when they left school and started their careers. In this article we've tried to maintain a clear separation between the model as a tool for describing the research process in terms which are comprehensible to biostatisticians (our primary objective and the focus here) and the use of the model for curriculum development (a secondary objective, and the focus of a subsequent report). Nevertheless, some of the potential classroom exercises should be apparent to the reader and, indeed, we can report that our initial experiences in using the model as a basis for developing classroom exercises have been encouraging.

2.2 Approach

Our approach to developing the model was to leverage commonalities between three major areas of scientific pursuit that we defined as relevant for collaborative biostatisticians: planning, conducting, and reporting of research studies. We note that this is not the only way to conceptualize the scientific process, but we reasoned that it aligns well with the daily life of a practitioner of biostatistics.

As an initial assessment of whether the model might meet its intended purpose, we considered the task of reading and interpreting a peer-reviewed publication describing a medical research study. In practice, biostatisticians routinely review such publications (e.g., to extract information for planning additional studies) and also contribute to the writing of such publications (e.g., by describing the statistical methods applied to analyze the data and presenting the results of the analysis). The task of reading and interpreting a journal article therefore seemed a natural first assessment of our model. If we could successfully apply the model to complete this task then we would consider the model to have some face validity, and this would suggest further development and testing of the model is warranted. We accomplished this by evaluating the fit of our model to a report of a well-designed and well-described clinical trial published in a leading medical journal (Du Toit et al., 2015).

3. Results

3.1 The Biomedical Research Pyramid Model

The biomedical research pyramid (Figure 1) represents all of the major steps in a scientific inquiry, with the research objectives forming the base of the pyramid and the other building blocks creating a stepwise progression toward the peak, which is the output of the scientific process: dissemination of the results. Importantly, the build-up from the base of the pyramid toward the peak follows a process that is broadly applicable to the major scientific activities as seen through the lens of a biostatistician: planning research studies, conducting research, and reporting research

results. Therefore, the pyramid provides both a mental map for the practicing biostatistician and a pedagogical map for curriculum development in any one of these three important areas of science. The pyramid building blocks are described in detail below.

Figure 1. The Biomedical Research Pyramid

3.1.1 Research Objectives

We define the research objective to be one or more statements of what the researcher wants to accomplish. Following the format of a National Institutes of Health (NIH) grant proposal, we explain that the research objectives are usually stated as specific aims. There are various templates, examples, and guidelines for writing specific aims available from the NIH and elsewhere (Draft Specific Aims, 2020; Monte & Libby, 2018). In our model, we point out the following characteristics of a specific aim that make for clearly stated research objectives: a problem statement, public health importance of the problem, what is currently known about the problem, the gap in knowledge the research will fill, what specifically will be done as part of the research, and a summary of how the research will contribute to the body of scientific knowledge.

3.1.2 Biological Basis for Research

Our model links the biological basis for the research to the research objectives using an established model for the determinants of health, which describes how biology and genetics interact with policy making, social factors, health services and individual behavior to influence health (Healthy People 2020: Determinants of Health, 2020). Furthermore, we emphasize that the biological basis for research can be described in numerous ways, including (but not limited to) anatomy, genetics, cellular processes, biochemistry, and molecular biology. The overarching concept is that no matter which of these areas of biological focus are germane to the research, all are indicative of a mechanism that is potentially influential on health, sometimes alone but often in concert with other (non-biological) factors.

3.1.3 Hypotheses

Our model makes a distinction between scientific and statistical hypotheses. Sientific hypotheses are formulated directly from the research objectives and biological basis for the research. Statistical hypotheses are an articulation of the scientific hypotheses that enables the application of statistical methods to address the research objectives.

3.1.4 Study Design

Research objectives can generally be addressed using a variety of different study designs with pros and cons associated with each different design approach. In many circumstances there is no clear optimal choice and the design choice requires careful consideration and justification. Our model facilitates flexibility in thought by allowing the applied biostatistician (or student) to explore these pros and cons, typically framed in terms of sources of bias (e.g., selection bias, misclassification bias), and how they will impact the actual statistical hypotheses tested as well as the interpretation of the results (Rothman et al., 2012).

3.1.5 Statistical Analysis

The central location of statistical analysis in the pyramid conveys to the applied biostatistician (or student) how their work is not only built upon a foundation, but also receives information from and communicates information to all of the other building blocks of the research pyramid. Further, its location after study design helps reinforce the

connection between study design and statistical analysis. For example, an appropriate statistical analysis plan cannot be written until the study design is selected.

3.1.6 Critical Evaluation and Interpretation

This piece of the pyramid emphasizes that science is not an exercise in uncovering the truth but is instead a process by which knowledge is developed through observation, experimentation, and careful evaluation. There are various paradigms for critical evaluation that might fit this section of the pyramid and the choice of approach for critical evaluation can be tailored to the research at hand (Elwood, 2017).

3.1.7 Dissemination of Results

Dissemination is at the peak of the pyramid because it is the last and most important step in the scientific process. In the scientific world, dissemination takes many forms: journal articles, oral presentations, written abstracts, posters, etc. Our model makes dissemination seamless because it is the natural outgrowth of the building blocks below it in the pyramid.

3.2 Application of the Biomedical Research Pyramid Model to the Description of a Scientific Manuscript

Figure 2 shows how we propose to apply the biomedical research pyramid model to the interpretation of a scientific manuscript. We first note that the model provides a valuable contextual orientation for the biostatistician who is reading the manuscript. Specifically, the manuscript is a description of research results (top of the pyramid) and therefore the reader should expect it to rest on a strong foundation that begins at the base of the pyramid. In fact, most scientific manuscripts follow this construction, consisting of Introduction, Methods, Results and Discussion sections. The Introduction should address the research objectives, biological basis for research, and the hypotheses tested. The Methods section typically includes the study design and a description of the statistical methods used to analyze the data. The Results should give a detailed description of the results of the statistical analyses that test pre-specified hypotheses, and the Discussion summarizes the primary findings and offers a critical evaluation of the results, typically presented as a discussion of strengths and weaknesses of the research.

To evaluate the goodness of fit of our model for the task of describing a scientific manuscript, we selected a high impact publication from the New England Journal of Medicine. The publication describes a clinical trial of a dietary intervention to prevent development of peanut allergy in infants at high risk for this potentially life-threatening condition (Du Toit et al., 2015). We walk through the application of our model to each section of the manuscript as follows.

3.2.1 Introduction

The article begins by describing the overarching objective, to address a very real public health threat: the prevalence of peanut allergy among children in Western countries has doubled over a 10 year period, and is the leading cause of death from food allergies (Du Toit et al., 2015). The specific objective of this study was to test an intervention strategy for reducing the risk of developing peanut allergy in this population. The test was to be made by means of a randomized experiment in which infants were assigned to either incorporate a measured amount of peanut-containing food into their diet or to avoid peanuts entirely.

The hypothesis proposed that introduction of small amounts of peanuts into the diet early during infancy might confer protection against subsequent development of peanut allergy. This hypothesis was supported partly by observational data that showed Jewish children in the United Kingdom had 10 times the risk of peanut allergy compared with Israeli Jewish children. This observation was important because the same study noted that children in the UK typically did not eat peanuts during the first year of life whereas the Israeli children were exposed to peanuts in their diet on average by 7 months of age.

The hypothesis was also supported by evidence from an experiment in mice that showed immune tolerance to peanut antigen was induced by deliberate administration of small amounts of the same antigen. The authors admitted that the exact biologic mechanism through which this might occur was "an incompletely understood phenomenon". But they did recognize the potential that children with prior exposure to peanut antigen may have been sensitized to some degree already and might reasonably respond in a different way to the introduction of peanuts into their regular diet as compared to children who have no prior exposure to peanuts. Therefore, the trial was conducted in two groups of children: those who tested positive to the skin-prick test for peanut exposure, and those who tested negative.

3.2.2 Methods

The study was a multicenter, randomized, controlled, parallel group trial conducted during the years 2006 to 2009. Infants were eligible if they were 4-11 months of age and had severe eczema, egg allergy or both (putting them at

high risk for developing peanut allergy). At the outset of the study, the children were stratified into two cohorts: skin-prick test positive for prior peanut exposure and skin-prick test negative. Random assignment to the study groups was conducted separately within these strata.

The trial used an open-label design, meaning that the participants and the investigators both knew whether they had been assigned to eat or avoid peanuts. The primary outcome of the trial was development of peanut allergy at 60 months of age, which was assessed using a food challenge.

The statistical analysis used the intention-to-treat approach, including all participants in their assigned treatment group regardless of compliance with the study protocol. The comparison of study arms was performed separately for each of the cohorts—skin-prick test positive and negative—using a chi-square test.

3.2.3 Results

A total of 640 participants were enrolled, 542 of whom were skin-prick negative (270 randomized to peanut avoidance and 272 randomized to peanut consumption) and 98 were skin-prick positive (51 randomized to peanut avoidance and 47 randomized to peanut consumption). Results were presented separately for each of these cohorts and the fundamental finding was the same in both groups. In the skin-prick positive cohort the prevalence of peanut allergy at 60 months of age was 35.3% in the peanut avoidance group and 10.6% in the peanut consumption group with a Chi-square p-value of 0.004, suggesting the observed result (or a difference even more extreme) is exceedingly unlikely under the null hypothesis that the dietary intervention does not prevent peanut allergy. A similar result was observed in the skin-prick negative cohort with 13.7% of participants in the peanut avoidance group having peanut allergy at 60 months of age and only 1.9% of participants in the consumption group having peanut allergy ($p < 0.001$). Thus, the null hypothesis of no treatment effect was rejected in both cohorts, suggesting exposure to peanuts in early life can reduce the risk of developing peanut allergy regardless of whether there is some baseline sensitivity to peanut antigen.

We note that some participants were excluded from the final analysis and that complete details are given in the article concerning these exclusions. We point this out only to illustrate that application of our model to this article identified at least one possibility for a teachable moment in biostatistics: the validity of statistical analyses in the presence of missing primary outcome data. There are also various analyses of immune biomarkers which we don't mention here for brevity. But suffice it to say that the article includes an assessment and comparison of the biological response between the randomized groups.

3.2.4 Discussion

The Discussion restates the most important finding, which is that introduction of peanuts into the diet of infants appears to confer protection from peanut allergy regardless of whether the child has any previous exposure to peanuts or not. The primary strengths cited are the reliance on randomized comparison, which reduces the possibility that the observed result can be attributed to confounding factors, and the excellent adherence to the study protocol by the enrolled participants.

One of the major limitations the authors point out is with respect to the generalizability of the results, being that the study included only infants known to be at high risk of developing peanut allergy. Therefore, the study results do not necessarily answer the question of whether this prevention strategy is useful in the broader population that includes lower risk infants. The authors also point out that the duration of protection of peanut allergy is uncertain based on their results. In other words, the infants assigned to consume peanuts in this study eventually stopped, and it is not clear whether the protection from peanut allergy will last. The authors indicate that this question is being addressed in a future study.

4. Discussion

In this article we described a conceptual model that can serve as a guide to the practicing biostatistician. We assessed the potential utility of this model by applying it to the description of a real scientific research study and found that the model fit well. During that process we also noted how deeper instruction in the details of biostatistical practice could be anchored to the model (e.g., the validity of statistical analyses in the presence of missing primary outcome data). In this way, we feel the model has face validity for its purpose as a tool for the practicing biostatistician and as a basis for instruction in biostatistics. Specifically, the model can serve as a mental map for the research process and it offers the opportunity to incorporate instruction in the details of biostatistical methods.

Our fundamental teaching goal is to train biostatistics students to become effective practitioners of collaborative

biostatistics, which includes "hard" skills, such as the ability to perform specific statistical tests, and "soft" skills, such as the ability to successfully interact with investigators. How to effectively teach soft skills is understudied (Davidson et al., 2019) and this topic is one of the foci of our long-term research in statistical pedagogy. We hypothesize that our model will enable students to successfully integrate hard and soft skills by teaching them how to organize their thinking around the application of biostatistics to specific research questions. Future work is required to evaluate this hypothesis.

Consistent with the tenets of constructivism, our overall approach to teaching collaborative biostatistics combines explicit mental maps of various elements of statistical practice with specific tools intended to assist in applying those maps (Biggs, 1996). As an example, we have previously described an explicit model of the soft skill of conducting a meeting with an investigator, with each of the statistician's contributions falling into (a) providing information about statistics; (b) commenting on the investigator's understanding of statistics; (c) sharing their understanding of clinical content; and (d) prompting for information (e.g., about clinical content, about the design and execution of the study being discussed) (Samsa et al., 2018). Having been explicitly defined, each of those contributions can then become the focus of didactic instruction and active practice.

To implement such mental maps of statistical practice, specific tools are useful. For example, such tools include a schematic diagram of the data collection schedule, printouts of representative data records, a checklist of items to include in a typical statistical analysis plan, blank analysis tables to assist in understanding the statistical analysis plan, etc. Indeed, the use of such tools is not only helpful in actual statistical practice but can also speed the acquisition of knowledge because of their explicit and concrete nature. A theme throughout our efforts in statistical pedagogy is that of translating the implicit "tacit" knowledge of the experienced biostatistician into an explicit (and thus teachable) form.

Within this framework, the Biomedical Research Pyramid can be classified as a high-level "meta-tool" intended to provide students with a "big picture" to help orient them to the overall research process; and serve as a repository for lower-level mental maps of its individual elements. Understanding the overall research process assists in subsequent learning as it allows information to be classified in the appropriate place. As an example of lower-level mental maps that could be derived from our model, Figure 2 translates the Biomedical Research Pyramid into the components of a typical journal article, with "hypotheses" as one of its elements. The start of developing a lower-level mental map of hypotheses is to make the distinction between a scientific hypothesis (e.g., that introduction of peanuts into the diet of infants will reduce the risk of developing peanut allergy) and a statistical hypothesis (e.g., that the difference in probability of developing peanut allergy 6 months post-intervention is equivalent among infants who eat peanuts and those who avoid peanuts in their diet). As the student gains more experience, the concept of a statistical hypothesis will eventually be represented by an even lower-level and more detailed mental map, involving concerns such as the distinction between one- and two-sided tests, Type I and II error probabilities, statistical power, etc. A key component of this process of learning and construction is an effective "meta-map", which the Research Pyramid is intended to provide.

Figure 2. Application of the Biomedical Research Pyramid to Describing a Medical Research Article

5. Conclusion

In summary, we defined a flexible model for the practice of biostatistics that also has potential to serve as a pedagogical tool. The uniqueness of the model lies in the fact that its focus is the discipline, not the learning framework. While the model is loosely rooted in constructivist theory, it does not demand this as a pedagogical approach. Instead, the model offers flexibility in designing curricula and specific educational activities with various underlying models of pedagogy and using different learning methodologies and environments, e.g., role-playing, flipped classroom, online asynchronous, in-person didactic, and others. In this article we only briefly alluded to the use of the model as the organizing focus for course work in applied biostatistics. A subsequent report will describe an innovative course in biostatistical practice organized around the Biomedical Research Pyramid, and this report will include additional information about how the model can be translated into specific classroom activities.

References

Begg, M. D., & Vaughan, R. D. (2011). Are biostatistics students prepared to succeed in the era of interdisciplinary science? (and how will we know?). *American Statistician*, *65*(2), 71-79. https://doi.org/10.1198/tast.2011.10222

Belli, G., & Batanero, C. (2001). The Teaching / Learning Process in University Statistical Consulting Labs in the United States (Ed.). *International Association for Statistical Education*, 325-338. http://www.ugr.es/~batanero/documentos/training.pdf

Biggs, J. (1996). Enhancing teaching through constructive alignment. *Higher Education*, *32*(3), 347-364. https://doi.org/10.1007/BF00138871

Davidson, M. A., Dewey, C. M., & Fleming, A. E. (2019). Teaching Communication in a Statistical Collaboration Course: A Feasible, Project-Based, Multimodal Curriculum. *The American Statistician*, *73*(1), 61-69. https://doi.org/10.1080/00031305.2018.1448890

Draft Specific Aims. (2020). National Institute for Allergy and Infectious Disease. Retrieved from https://www.niaid.nih.gov/grants-contracts/draft-specific-aims

Du Toit, G., Roberts, G., Sayre, P. H., Bahnson, H. T., Radulovic, S., Santos, A. F., Brough, H. A., Phippard, D., Basting, M., Feeney, M., Turcanu, V., Sever, M. L., Gomez Lorenzo, M., Plaut, M., & Lack, G. (2015). Randomized Trial of Peanut Consumption in Infants at Risk for Peanut Allergy. *New England Journal of Medicine*, *372*(9), 803-813. https://doi.org/10.1056/NEJMoa1414850

Elwood, M. (2017). *Critical Appraisal of Epidemiological Studies and Clinical Trials* (Vol. 1). Oxford University Press. https://doi.org/10.1093/med/9780199682898.001.0001

Healthy People 2020: Determinants of Health. (2020). HealthyPeople.Gov. Retrieved from https://www.healthypeople.gov/2020/about/foundation-health-measures/determinants-of-health

Monte, A. A., & Libby, A. M. (2018). Introduction to the Specific Aims Page of a Grant Proposal. *Academic Emergency Medicine*, *25*(9), 1042-1047. https://doi.org/10.1111/acem.13419

Novak, J. D., & Cañas, A. J. (2007). Theoretical Origins of Concept Maps, How to Construct Them, and Uses in Education. *Reflecting Education*, *3*(1), 29-42.

Rothman, K., Greenland, S., & Lash, T. (2012). *Modern Epidemiology* (3rd ed.). Philadelphia, PA: Wolters Kluwer/Lippincott, Williams & Wilkins.

Samsa, G. P. (2018). A Day in the Professional Life of a Collaborative Biostatistician Deconstructed: Implications for Curriculum Design. *Journal of Curriculum and Teaching*, *7*(1), 20-31. https://doi.org/10.5430/jct.v7n1p20

Samsa, G. P., LeBlanc, T. W., Locke, S. C., Troy, J. D., & Pomann, G.-M. (2018). A Model of Cross-Disciplinary Communication for Collaborative Statisticians: Implications for Curriculum Design. *Journal of Curriculum and Teaching*, *7*(1), 1-11. https://doi.org/10.5430/jct.v7n2p1

Zapf, A., Huebner, M., Rauch, G., & Kieser, M. (2019). What makes a biostatistician? *Statistics in Medicine*, *38*(4), 695-701. https://doi.org/10.1002/sim.7998

Copyrights