Ideal Observers and Efficiency:
Commemorating 50 Years of Tanner and Birdsall: Introduction

Just over 50 years ago (1958), Wilson P. Tanner and Theodore G. Birdsall published an article in the *Journal of the Acoustical Society of America* [1] that was to have a profound effect on psychophysics in general and visual perception in particular. From the title of the paper, "Definitions of d' and η as psychophysical measures," it was clear that the focus of the work would be on methodology in perceptual experiments. This methodology has persisted essentially unchanged in principle, despite enormous changes in the methods and equipment available for making perceptual measurements and ever more sophisticated models of perceptual phenomena. It would be quite an understatement to say that they have achieved their stated goal of clarifying "…the reasons for employing these variables in psychoacoustic experiments."

Both measures have been used extensively in visual perception experiments and psychophysics in general. What makes the efficiency measure, η, so particularly powerful as a performance metric is that it builds in the constraints of the stimuli on the sensory system. Efficiency is based on the construct of the ideal observer. An ideal observer is a theoretical machine that makes optimal use of available information in a given task. The performance of the ideal observer in any given task is constrained only by the limitations imposed by the physical availability of information and therefore represents a strict upper bound on performance. Real observers (e.g., humans) are not ideal. Unlike an ideal observer, information for a real observer must be encoded by sensory organs and is subjected to further processing before a decision is made. Information can be lost at any of these stages, yielding performance that is less than the ideal.

If we know that real observers are not ideal, then what use is it to measure the performance of an ideal observer? This might be a fair question if our reason for measuring ideal performance was simply to predict human performance. The ideal observer is often not a good model of human observers. However, the ideal observer is rarely used to model human performance in this fashion. Rather, it is the *departures* from ideal performance that can tell us a great deal about how information is coded by a real observer. This is precisely what is measured in the efficiency variable, η, which is determined by the ratio of the squared sensitivity for the observer under investigation to the ideal. An efficiency of 100% would imply no loss of task-relevant information between encoding and response. Efficiencies less than 100% imply that information has been lost at some point and therefore the processes employed by the observer under investigation depart somehow from that of the ideal.

At the most basic level, a departure from the ideal allows us to rule out one simple model of the observer under investigation (the ideal model). But the real utility of the ideal observer lies in the various ways that the comparison to the ideal can give us important insights into the kinds of processes that are taking place within an observer that is less than ideal. For example, constraints can be systematically built into an otherwise ideal observer while maintaining the optimality of the model with respect to the added constraints. The performance of this constrained ideal observer can be compared to that of a real observer. Because the constraints have been built into an otherwise ideal model, any increases in efficiency that are observed after including the constraints indicate the contribution of those processes to the departure of the real observer from the ideal. A strict version of this approach, known as *sequential ideal observer analysis*, involves measuring the physical effect of each stage of processing on the stimulus (e.g., measuring the filtering characteristics of the cornea and lens) and building these stages into the ideal model. By sequentially introducing each processing stage into an otherwise ideal model, it is possible to measure the contribution of each stage to the overall departure from the ideal and localize where information loss is taking place.

An approach related to sequential ideal observer analysis is to systematically manipulate various aspects of the task performed by real and ideal observers to see how these manipulations affect efficiency. This approach rests on the premise that real observers have evolved to perform certain tasks with high efficiency and that low efficiency implies that the real observer has been asked to perform a task for which their perceptual system is not well tuned. By systematically manipulating the task in ways that we suspect will make it approach something that is more similar to what the real observer has evolved to perform (e.g., introducing various uncertainties about the stimulus, such as spatial uncertainty), we can identify likely processes that are taking place within the real observer by tracing the effect these manipulations have on efficiency. Closely related to this approach is the comparison of real observer performance to the performance of an ideal observer that has limited access to some aspect of the stimulus information (e.g., has access only to 2-D information in a 3-D task). In this case, the performance of the real observer may exceed that of the information-limited ideal observer, implying the human observer must be using the information to which the information-limited observer does not have access.

An additional aspect of ideal observer analysis that often plays an important role in of all of the approaches described above is the ability to compare an observer's performance across different tasks or conditions. Unless carefully contrived in some fashion, tasks and conditions typically differ in terms of their intrinsic difficulty. Therefore, simply comparing an observer's performance across two or more tasks or conditions is very difficult to interpret: Is the pattern of performance across tasks or conditions due to variations in the physical availability of information? Or is it due to variations in the observer's

ability to make use of information? Or perhaps a combination of the two? Comparison to ideal performance (i.e., efficiency) decouples these two factors by taking variations in the physical constraints on performance out of the equation. Any variations in efficiency across conditions must be due to processes taking place within the observer under investigation. The importance of this application of ideal observer is clear in the simple case where one wishes to make veridical comparisons of a nonideal observer's performance across a set of tasks or conditions. But the real utility of this application of ideal observer analysis is most elegantly illustrated when it is used in conjunction with one of the approaches described above. For example, efficiency may vary across a set of tasks or conditions, but this variation might be eliminated by the inclusion of a single processing stage that is common to all of the tasks or conditions.

Much of what has been described above represents an evolution well beyond the original domain of ideal observer analysis. In their seminal article, Tanner and Birdsall developed the underpinnings of ideal observer analysis as well as the concept of efficiency within the framework of the "theory of signal detectability" (which is now generally referred to as signal detection theory, even though Tanner and Birdsall objected to this term for describing the limitations of a "receiver"). The purpose of this special feature in JOSA A is to commemorate this important advance in the development of our field and to recognize the extensive influence of the efficiency concept put forth by Tanner and Birdsall.

The topics in the collection of papers contained in this special feature cover a range of issues that all fall within the ideal observer framework. Several of the papers address theoretical and methodological issues related to past and current models of vision, including the predic-tions of template models and the influence of external noise (Klein and Levi; Abbey and Eckstein; Jeon, Lu, and Dosher), the predictions of single and multiple channel models (Taylor and Bennett), and the development of more effective ways of estimating ideal performance in the presence of statistical complexity (Park and Clarkson). Other papers involve the application of ideal observer analysis to more specific problems in vision, including the recognition of patterns in correlated noise (Conrey and Gold), the programming and execution of saccadic eye movements (Strizke, Trommershauser, and Gegenfurtner), the perception of color and natural scenes (Foster, Marin-Franch, Amano, and Nascimento), and the neural coding of motion (Lalor, Ahmadian, and Paniski).

We would like to acknowledge the hard work of the authors, reviewers, and editorial staff who have devoted a great deal of their time and energy to creating and refining the papers contained in this special feature. We believe the result is a fitting tribute to the enduring relevance of Tanner and Birdsall's work.

REFERENCE

1. W. P. Tanner and T. G. Birdsall, "Definitions of d' and η as psychophysical measures," J. Acoust. Soc. Am. **30**, 922–928 (1958).

Jason M. Gold
Craig Abbey
Bosco S. Tjan
Daniel Kersten
Feature Editors
Ideal Observers and Efficiency:
Commemorating 50 Years of Tanner and Birdsall