Cognitive Psychology 92 (2017) 101–126

Contents lists available at ScienceDirect

Cognitive Psychology

journal homepage: www.elsevier.com/locate/cogpsych

Likelihood ratio sequential sampling models of recognition memory

Adam F. Osth [a,*], Simon Dennis [b], Andrew Heathcote [c]

[a] University of Melbourne, Australia
[b] University of Newcastle, Australia
[c] University of Tasmania, Australia

A R T I C L E   I N F O

Article history:
Accepted 15 November 2016
Available online 3 December 2016

A B S T R A C T

The mirror effect – a phenomenon whereby a manipulation produces opposite effects on hit and false alarm rates – is benchmark regularity of recognition memory. A likelihood ratio decision process, basing recognition on the relative likelihood that a stimulus is a target or a lure, naturally predicts the mirror effect, and so has been widely adopted in quantitative models of recognition memory. Glanzer, Hilford, and Maloney (2009) demonstrated that likelihood ratio models, assuming Gaussian memory strength, are also capable of explaining regularities observed in receiver-operating characteristics (ROCs), such as greater target than lure variance. Despite its central place in theorising about recognition memory, however, this class of models has not been tested using response time (RT) distributions. In this article, we develop a linear approximation to the likelihood ratio transformation, which we show predicts the same regularities as the exact transformation. This development enabled us to develop a tractable model of recognition-memory RT based on the diffusion decision model (DDM), with inputs (drift rates) provided by an approximate likelihood ratio transformation. We compared this "LR-DDM" to a standard DDM where all targets and lures receive their own drift rate parameters. Both were implemented as hierarchical Bayesian models and applied to four datasets. Model selection taking into account parsimony favored the LR-DDM, which requires fewer parameters than the standard DDM but still fits the data well. These results support log-likelihood based models as providing an elegant explanation of the regularities of recognition memory, not only in terms of choices made but also in terms of the times it takes to make them.

1. Introduction

In recognition memory, participants study a list of items and during a test phase and are asked to discriminate between studied items (targets) and unstudied items (lures). Two of the most successful modeling frameworks for decision making in recognition memory are signal detection theory (SDT) and sequential sampling. In signal detection theory, different stimulus conditions are represented as continuous evidence distributions (usually Gaussian in shape), with the observer placing a criterion on the evidence axis. Models in the SDT framework are successful for accounting for the shape of the receiver operating characteristic (ROC). To construct an ROC, participants undergo recognition memory testing across a range of different bias conditions; hit rates (HR) are plotted against false alarm rates (FAR) for each bias condition. SDT models were successful in predicting the curvilinear shape of the ROC, a nearly universal finding in recognition memory (Egan, 1958; Wixted, 2007).

* Corresponding author.
E-mail address: adamosth@gmail.com (A.F. Osth).

http://dx.doi.org/10.1016/j.cogpsych.2016.11.007
0010-0285/

Nonetheless, a major weakness of SDT models is their inability to predict the shape of response time (RT) distributions: if RT is determined by the distance from the response criterion, SDT models are not able to correctly predict right skewed RT distributions for both correct and error responses (Ratcliff & McKoon, 2008).

Sequential sampling models are the most successful framework for predicting the shapes of choice RT distributions. In sequential sampling models, evidence is sampled from a stimulus until it reaches one of two response boundaries corresponding to the decision alternatives; the response boundary is the choice and the time taken to reach the boundary is the RT. We focus on the *diffusion model*, in which evidence begins to accumulate at the starting point z that is placed between two response boundaries, the upper boundary a, and the lower boundary at 0. The rate of evidence accumulation is called the *drift rate* (denoted by v): positive drift rates tend toward the upper boundary and negative drift rates tend toward the lower boundary. As the absolute value of the drift rate increases, the rate of correct responses increases and RT decreases. Errors are made because evidence accumulation is noisy; on each timestep Gaussian noise is added to the accrued evidence. To account for perceptual encoding and response output processes that are outside of the scope of evidence accumulation, a fixed constant t_{er} is added to the RT distribution to reflect nondecision time.

There were two major successes of classical diffusion models. First, in contrast to distance-from-criterion SDT models, they naturally predicted right skewed RT distributions for both correct and error responses. Second, they are able to account for the speed-accuracy tradeoff through changes in the response boundary. A decreased boundary produce faster RTs as the diffusion process has less distance to travel, but more errors result because closer response boundaries make it more likely that the diffusion process will reach the incorrect boundary by accident. Nonetheless, the classical diffusion model also has a number of weaknesses. As response boundaries increase errors disappear entirely, whereas experiments examining speed-accuracy tradeoffs find that asymptotic accuracy is usually far from perfect, especially in recognition memory (e.g., Reed, 1976). Additionally, classical diffusion models have difficulty with the relative speeds of correct and error responses. Under speeded conditions, errors are often faster than correct responses, whereas with more cautious responding errors are often slower than correct responses. Classical diffusion models, in contrast, predict equivalent RT distributions for correct and error responses under unbiased responding (Laming, 1968).

The problems with both of these modeling frameworks were solved by marrying them into a single framework, originally by Laming (1968) in discrete time, and later in continuous time in Ratcliff's diffusion decision model (DDM: Ratcliff, 1978; Ratcliff & McKoon, 2008). The DDM uses an SDT front-end for a diffusion process: drift rates for each trial are sampled relative to a drift criterion d_c from Gaussian evidence distributions with standard deviation η. Trial-to-trial variability in the drift rate ensures that there is an *asymptotic* d'; as response boundaries are increased, performance can never exceed the limit imposed by the overlap of target and lure drift rate distributions. Additionally, drift rate variability allows the model to predict error responses that are slower than correct responses (Ratcliff, 1978; Ratcliff & McKoon, 2008). A diagram of the DDM can be seen in Fig. 1.

Ironically, in the years that followed the publication of the seminal Ratcliff (1978) article, both SDT models and sequential sampling models developed largely independently of each other. A major development in SDT models of recognition memory was the rejection of the equal variance signal detection model based on investigations of the z-transformed ROC (zROC). Equal variance signal detection models predict linear zROCs with a slope of 1. However, many investigations have revealed zROC slopes of around 0.8 (Egan, 1958; Glanzer & Adams, 1990; Heathcote, 2003; Ratcliff, McKoon, & Tindall, 1994; Ratcliff, Sheu, & Gronlund, 1992), or even less when random item variability is taken into account (Averell, Prince, & Heathcote, 2016; Pratte & Rouder, 2012; Pratte, Rouder, & Morey, 2010). As a consequence, theorists have adopted the *unequal variance signal detection* (UVSD) model, which allows greater variability for targets than for lures, potentially due to the contribution of encoding variability (Wixted, 2007). Another development in SDT models, which will be described in more detail below, is the usage of *log-likelihood ratio signal detection theory* models to capture the mirror effect (Glanzer & Adams, 1985; Glanzer, Hilford, & Maloney, 2009).

The DDM was also updated with the adoption from Laming (1968) of cross-trial variability in the starting point of evidence accumulation (Ratcliff & Rouder, 1998; Ratcliff, Van Zandt, & McKoon, 1999). Rather than have the starting point fixed across trials, the starting point was sampled from a uniform distribution with range s_z. The inclusion of this variability parameter allowed for a complete account of the speed-accuracy tradeoff: the DDM with cross-trial variability in both starting point and drift rates can predict faster errors than correct responses in speeded conditions while predicting slower correct than error responses in conditions that emphasise accuracy (Ratcliff & Smith, 2004). The model was further updated with the inclusion of cross-trial variability in nondecision time (t_{er}), where nondecision time is sampled from a uniform distribution with range s_t, to allow for better predictions of the leading edge of the RT distributions across different levels of performance (Ratcliff, Gomez, & McKoon, 2004). However, the majority of DDM applications to recognition memory continued to use equal variance for targets and lures (Arnold, Broder, & Bayen, 2015; Bowen, Spaniol, Patel, & Voss, 2016; Criss, 2010; Ratcliff & Smith, 2004; Ratcliff, Thapar, & McKoon, 2004; Ratcliff, Thapar, & McKoon, 2010; Ratcliff, Thapar, & McKoon, 2011; White & Poldrack, 2014).

A re-introduction of contemporary SDT influences to the DDM came from Starns, Ratcliff, and McKoon (2012), who tested whether or estimates from the DDM were consistent with unequal variance signal detection models. They applied the DDM to binary ROC data, where an ROC was created by giving participants yes/no decisions with bias manipulated via changes in response proportions. This procedure allows for application of the DDM by manipulating the starting point along with the drift criterion across the bias conditions. Starns et al. manipulated bias using five different levels of target proportions, yield-