

Perceptual abilities undergo major development during infancy and childhood – for example, for detecting low-contrast stimuli (Adams & Courage 2002) and noisy patterns of motion (Hadad et al. 2011) or recognising complex stimuli such as faces (Mondloch et al. 2002). Classically, the focus of perceptual development research has been on improvements in sensitivity (likelihoods). As reviewed in the target article, decades of adult research show how sensitivity changes can result from changes within a decision-model framework that incorporates likelihoods, priors, cost functions, and decision rules. Applying this framework to development, we argue that *perceptual improvements must be explained in terms of changes to these components*. This will lead to a new understanding of how perceptual systems attain their more highly optimised mature state.

Specifically, we need to know the following:

- (1) *Which elements of the observer model are changing (developing), leading to improvements in perceptual function?* Recent evidence suggests that multiple components of the decision model are developing significantly during childhood. Until late into childhood, observers are still using *decision rules* less efficiently: misweighting informative cues (Gori et al. 2008; Manning et al. 2014; Sweeny et al. 2015) or using qualitatively different decision rules altogether (Jones & Dekker 2017; Nardini et al. 2008; 2010). Other studies show abilities to learn and use *priors* and *costs* also to be developing late into childhood (e.g., Dekker & Nardini 2016; Stone 2011; Thomas et al. 2010). The new, model-based approach to development pioneered in these studies paves the way for understanding how likelihoods, priors, cost functions, and decision rules are shaped as children learn, and for testing which common processes can explain perceptual development across a range of different tasks. Studies to date have successfully captured developmental changes in performance by fitting how parameters of specific components of the decision model change with age on single tasks. This usefully sets quantitative bounds on potential changes in these processes, but the data are often compatible with more than one account. For example, in a rewarded reaching task (Dekker & Nardini 2016), children up to the age of 11 years aim too close to a penalty region to maximise their score, reflecting overconfidence in likelihood of hitting the target, underestimation of cost, or a central pointing prior. An important way forward is therefore to evaluate the fit of developmental models to multiple tasks and to test their predictions on new tasks.
- (2) *How are more efficient and adult-like decision rules, priors, and cost functions acquired during development?* Beyond characterising the changes in decision-model components underlying perceptual development, the ultimate aim is to understand the mechanisms driving these changes. A major contributing factor is likely to be experience, which shapes the sensitivity of neuronal detectors, determining likelihoods (Blakemore & Van Sluyters 1975), changes priors (Adams et al. 2004), and is needed to learn the potential consequences of actions (cost factors). It is not clear in which circumstances such experience is generalizable (e.g., priors or costs learned during one task applied to another), how experience drives learning of decision rules, or whether there are sensitive periods like those for sensitivities (likelihoods) in other parts of the decision model (e.g., for learning priors). A useful approach is investigating the neural changes supporting improvements in decision-model components as perception becomes more

optimised, such as more precise representation of likelihoods (Van Bergen et al. 2015) and values (Wu et al. 2011), or more precise computing of weighted averages, perhaps implemented via divisive normalisation (Ohshiro et al. 2011). The power of this approach is demonstrated by recent studies of developmental disorders, in which there are exciting developments in linking components of observer models to specific neural mechanisms (Rosenberg et al. 2015). For example, in autism, tasks that involve combining new evidence with prior knowledge are disproportionately affected, and this has recently been linked to the overweighting of sensory likelihoods versus priors, possibly because of altered neural operations mediated by noradrenaline and acetylcholine (Lawson et al. 2017). In addition, a new, model-based approach to developmental neuroimaging lets us disentangle components of the developing decision model across different neural processing stages. We recently showed that development of cue integration during depth perception was linked to a shift from using depth cues independently to combining them, by neural detectors in sensory cortex (adopting a “fusion” rule; Dekker et al. 2015). This suggests that the late development of cue integration is driven by a change in how sensory information is combined (sensory decision rule), rather than improved readout of the fused estimate during task performance (higher-order decision rule or cost function). These studies demonstrate how a developmental approach can provide computational-level understanding of the crucial ingredients for building a mature optimised observer.

The end goal of this approach is an observer model incorporating processes of learning and development: a developing standard observer model. This will provide a more complete understanding of perceptual systems and a basis for developing intelligent machines that can learn to perceive in novel environments. For example, understanding the structure of experience that scaffolds our ability to transfer previous likelihoods, cost functions, and decision rules from one task to another can inform the development of more flexible artificial intelligence (AI) agents (Wang et al. 2017). Similarly, significant improvements in robotic grasp performance have been gained from incorporating developmental stages such as motor babbling and gradual improvements in visual acuity into the training regime (Cangelosi et al. 2015). In addition, understanding which developmental changes in the decision model (e.g., sensitivity vs. decision rule) drive perceptual improvements at different ages will provide a crucial basis for better training of perception and action in patients with sensory loss.

Supra-optimality may emanate from suboptimality, and hence optimality is no benchmark in multisensory integration

Jean-Paul Noel 

Vanderbilt Brain Institute, Vanderbilt University, Nashville, TN 37240.
jean-paul.noel@vanderbilt.edu <http://jeanpaulnoel.com/>

doi:10.1017/S0140525X18001280, e239

Abstract

Within a multisensory context, “optimality” has been used as a benchmark evidencing interdependent sensory channels. However, “optimality” does not truly bifurcate a spectrum from suboptimal to supra-optimal – where optimal and supra-optimal, but not suboptimal, indicate integration – as supra-optimality may result from the suboptimal integration of a present unisensory stimuli and an absent one (audio = audio + absence of vision).

Arguably the study of multisensory integration was born from the recording of spikes in the feline superior colliculus (Stein & Meredith 1993). These early studies presented animals with simple visual (V) flashes and auditory (A) beeps and held the occurrence of supra-additive responses (i.e., audiovisual [AV] responses greater than the sum of auditory and visual responses) as the hallmark for multisensory integration. However, this phenomenon is not common in the neocortical mantle (vs. subcortex; Frens & Van Opstal 1998), nor when multisensory integration is indexed via behavior or by measuring ensembles of neurons (e.g., local field potentials, electroencephalography [EEG], functional magnetic resonance imaging [fMRI]; Beauchamp 2005). Hence, over the last two decades there has been a greater appreciation for sub-additive responses as equally demonstrating an interesting transformation from input (i.e., A + V) to output (i.e., AV), and thus highlighting the synthesis of information across senses. That is, arguably the classic study of multisensory integration has grown to conceive of sub- and supra-additivity as being on extremes of a spectrum where both ends are interesting and informative.

In parallel, the originally described “principles of multisensory integration” (e.g., information that is in close spatial and temporal proximity will be integrated) have been translated to a computational language that is seemingly applicable throughout the cortex and widely observed in behavior. As Rahnev & Denison (R&D) underline in their review, this computational framework dictating much of the current work within the multisensory field is that of Bayesian decision theory. Indeed, among others, audiovisual (Alais & Burr 2004), visuo-tactile (Ernst & Banks 2002), visuo-vestibular (Fetsch et al. 2009), and visuo-proprioceptive (van Beers et al. 1999) pairings have been demonstrated to abide by maximum likelihood estimation (MLE) – the weighting of likelihoods by relative reliabilities and concurrent reduction in integrated (vs. unisensory) variance. Given this extensive body of literature, I believe the gut reaction of many multisensory researchers – mine included – to this review and the thesis that assessing optimality is not useful was that we must acknowledge the limitations of solely considering “optimality” without examining the underlying components (e.g., prior, cost function), but that this construct is nevertheless valuable. If subjects behave optimally (i.e., reduction of uncertainty), then at minimum, there is evidence for interdependent channels. Namely, the reduction of variance in multisensory cases (vs. unisensory) is evidence for the fact that at some point, unisensory components are fused; the next step is to understand exactly how these channels are fused. Furthering this argument, it could be conceived that supra- and suboptimality exist on a continuum where evidence for supra-optimality or optimality is evidence for multisensory integration (admittedly without providing much mechanistic insight given the points raised by R&D), while suboptimality does not bear evidence of a synthesis across the senses. In other words, indexing optimality as a benchmark for integration is

useful because Bayesian computations are ubiquitous in the brain and behavior, and in that it reduces the state space of integration from “anything apart from linear summation” (i.e., from sub-additive to supra-additive excluding additive) to “anything greater than or equal to optimal” (i.e., from optimal to supra-optimal but not suboptimal).

However, upon further consideration, I believe this reasoning to be erroneous (and therefore I agree with the thesis put forward by R&D). In short, contrarily to the case of additivity, optimality does not lie on a spectrum from sub- to supra-optimal, and hence optimality per se is no benchmark.

Traditionally, supra-optimality (an apparent impossibility) within multisensory systems has been hypothesized to emerge from a process of “active sensing” (Schroeder et al. 2010). That is, the presence of a second sensory stimulus (e.g., A) may sharpen the representation of a first unisensory stimulus (e.g., V) so that when these are combined (e.g., AV), sharper unisensory estimates than originally considered are combined, resulting in apparently supra-optimality. Nonetheless, as Shalom and Zaidel (2018) have recently highlighted, somewhat paradoxically, it could additionally be the case that supra-optimality results from suboptimal integration. Namely, researchers typically take unisensory likelihoods at face value. However, within a multisensory (e.g., AV) context, the presentation of auditory stimuli alone is in fact not auditory alone (e.g., A), but instead the presence of auditory information and the absence of visual information (e.g., A + no V). Therefore, in this example, researchers are underestimating the reliability of the auditory channel (which is truly A-likelihood + a flat visual likelihood), which will ultimately result in claims of supra-optimal multisensory integration. This second observation (by Shalom & Zaidel 2018) is similar to the case of active sensing, in that the sharpness of unisensory likelihoods is underestimated. However, the perspective is quite different in that supra-optimality is not the result of cross-modal feedback enhancing unisensory representation solely when presented in a multisensory context, but in fact, in this latter case, supra-optimality is merely an experimental construct that results from the erroneous underestimation of a unisensory likelihoods; the world is by nature multisensory, and hence unisensory estimates are impoverished estimates wherein a cue has been artificially removed. That is, supra-optimality can result from the non-optimal integration of a signal (e.g., A) and noise (e.g., a non-present V signal). In turn, there is no true gradient between supra- and suboptimality, and hence positioning optimality as a benchmark bifurcating between multisensory fusion and fission is ill advised. Instead, as highlighted by R&D, we ought to conceive of (multisensory) perception as a dynamic system where likelihoods, priors, cost functions, and decision criteria all interact interdependently in both feedforward and feedback manners.

When the simplest voluntary decisions appear patently suboptimal

Emilio Salinas , Joshua A. Seideman and Terrence R. Stanford

Department of Neurobiology & Anatomy, Wake Forest School of Medicine, Winston-Salem, NC 27157-1010.