Rational inattention: A new theory of neurodivergent information seeking

Jones, Samuel ; Jones, Manon; Koldewyn, Kami; Westermann, Gert

Developmental Science

DOI:
10.1111/desc.13492

Published: 29/03/2024

Publisher's PDF, also known as Version of record

Received: 2 May 2023 | Revised: 19 February 2024 | Accepted: 21 February 2024

DOI: 10.1111/desc.13492

Developmental Science WILEY

RESEARCH ARTICLE

Rational inattention: A new theory of neurodivergent information seeking

Samuel David Jones[1,2] | Manon Wyn Jones[1] | Kami Koldewyn[1] | Gert Westermann[2]

[1]Department of Psychology, Bangor University, Bangor, UK

[2]Department of Psychology, Lancaster University, Lancaster, UK

Correspondence
Samuel Jones, Room 309, Brigantia, Bangor University, Bangor, LL57 2AS, UK.
Email: samuel.jones@bangor.ac.uk

Funding information
Economic and Social Research Council, Grant/Award Numbers: ES/S007113/1, ES/L008955/1

Abstract

This paper presents *rational inattention* as a new, transdiagnostic theory of information seeking in neurodevelopmental conditions that have uneven cognitive and socio-emotional profiles, including developmental language disorder (DLD), dyslexia, dyscalculia and autism. Rational inattention holds that the optimal solution to minimizing epistemic uncertainty is to avoid imprecise information sources. The key theoretical contribution of this report is to endogenize imprecision, making it a function of the primary neurocognitive difficulties that have been invoked to explain neurodivergent phenotypes, including deficits in auditory perception, working memory, procedural learning and the social brain network. We argue that disengagement with information sources with low endogenous precision (e.g. speech in DLD, orthography-phonology mappings in dyslexia, numeric stimuli in dyscalculia and social signals in autism) constitutes resource-rational behaviour. We demonstrate the strength of this account in a series of computational simulations. In experiment 1, we simulate information seeking in artificial agents mimicking an array of neurodivergent phenotypes, which optimally explore a complex learning environment containing speech, text, numeric stimuli and social cues. In experiment 2, we simulate optimal information seeking in a cross-modal dual-task paradigm and qualitatively replicate empirical data from children with and without DLD. Across experiments, simulated agents' only aim was to maximally reduce epistemic uncertainty, with no difference in reward across information sources. We show that rational inattention emerges naturally in specific neurodivergent phenotypes as a function of low endogenous precision. For instance, an agent mimicking the DLD phenotype disengages with speech (and preferentially engages with alternative precise information sources) because endogenous imprecision renders speech not conducive to information gain. Because engagement is necessary for learning, simulation demonstrates how optimal information seeking may paradoxically contribute negatively to an already delayed learning trajectory in neurodivergent children.

KEYWORDS
attention, computational rationality, information seeking, neural networks, neurodevelopmental disorder, neurodivergence

Research Highlights

- We present the first comprehensive theory of information seeking in neurodivergent children to date, centred on the notion of rational inattention.
- We demonstrate the strength of this account in a series of computational simulations involving artificial agents mimicking specific neurodivergent phenotypes that optimally explore a complex learning environment containing speech, text, numeric stimuli, and social cues.
- We show how optimal information seeking may, paradoxically, contribute negatively to an already delayed learning trajectory in neurodivergent children.
- This report advances our understanding of the factors shaping short-term decision making and long-term learning in neurodivergent children.

1 | INTRODUCTION

The idea that children play an active role in their learning is well recognised. Whether through gaze direction, auditory attention or physical engagement, there is growing evidence that children preferentially allocate their cognitive and motor resources to environmental stimuli about which they are uncertain (Chen et al., 2022; Fandakova & Gruber, 2021; Kidd & Hayden, 2015; Poli et al., 2020; Twomey & Westermann, 2018). Yet, to date, active informationseeking has only been studied in neurotypical children. We are not aware of any account of information seeking in children affected by neurodevelopmental conditions such as developmental language disorder (DLD), developmental dyslexia, developmental dyscalculia and autism spectrum condition. The purpose of the current report is to present an account of information seeking that is consistent with both neurotypical and neurodivergent behaviour, and to demonstrate the strength of this account in a series of computational simulations that qualitatively replicate behavioural data from neurodivergent children.

Our account is centred on the notion of *rational inattention*, which we argue may be a fundamental, transdiagnostic characteristic of neurodivergent learning and development. The central message of this paper is that perhaps paradoxically, optimal information seeking may contribute negatively to an already delayed learning trajectory in neurodivergent children. Through the provision of a computational framework of neurodivergent information seeking, this report lends necessary traction to established ideas of compensatory strategies and the Matthew effect. Moreover, this report motivates a further reassessment of the factors contributing to apparently unconventional attentional behaviour in neurodivergent children. It is often claimed that specific neurodivergent phenotypes (e.g. DLD or dyslexia) can be attributed to an atypically constrained attentional capacity, a factor which may be improved through attention training (e.g. Ebert & Kohnert, 2009; Gathercole & Baddeley, 1990; Holmes et al., 2015). In contrast, the computational simulations presented in the current report show how apparently unconventional attentional behaviour in neurodivergent children may reflect optimal information seeking given low endogenous precision in the absence of any functionally discrete attentional bottleneck. This position is consistent with empirical evidence suggesting that attention training programmes have limited impact, and raises important questions about the deployment of such programmes in clinical settings (Shipstead et al., 2012; Jones & Westermann, 2022). The current report is, therefore, of both theoretical and clinical significance.

1.1 | Information seeking in neurotypical and neurodivergent children

The drive to learn about our environments has been explained in a number of complementary ways. Curiosity-based accounts centre around intrinsic reward, pragmatic accounts centre around learning to support good decision making (Gershman & Burke, 2022), and affective accounts maintain that uncertainty is stress provoking and should be avoided (Peters et al., 2017). What these accounts share is the belief that implicit and explicit decision making tends towards acting and engaging with the environment in a way that resolves the individual's epistemic uncertainty.

On the surface, neurodevelopmental disorders present a problem for accounts arguing that children preferentially engage with information sources about which they are uncertain (e.g. Chen et al., 2022; Poli et al., 2020). Information gaps are the defining feature of neurodevelopmental disorders. To simplify greatly, focussing on what might reasonably be considered canonical areas of difficulty: Children with DLD demonstrate greater uncertainty about speech than their peers;

children with dyslexia demonstrate greater uncertainty about spelling and phonology; children with dyscalculia demonstrate greater uncertainty about numeracy; and many children with autism demonstrate greater uncertainty about social cues[1]. Yet, these information gaps do not encourage preferential engagement. On the contrary, neurodivergent children are characteristically seen to engage less with those information sources with which they have primary difficulty (Ashkenazi et al., 2009; Jones et al., 2018; Moriuchi et al., 2017; Smolak et al., 2020; Soriano-Ferrer & Morte-Soriano, 2017; Stanovich, 2009; West et al., 2021; Annaz et al., 2009). There is, then, an apparent incongruity between neurotypical and neurodivergent information seeking: Neurotypical children characteristically engage with information sources about which they are uncertain, while neurodivergent children characteristically disengage with those specific information sources about which they are uncertain.

A range of neurocognitive mechanisms have been invoked to explain neurodivergent developmental profiles. These mechanisms range from low-level neural noise to deficits in auditory perception or working memory, to deficits in procedural or statistical learning or the development of the social brain network (Bishop & McArthur, 2005a, 2005b; Gray et al., 2019; Hancock et al., 2017; Lindsay et al., 1999). Yet, while these neurocognitive deficits might explain difficulties in *learning* a specific information type, for instance, speech, text, numeric stimuli or social cues, they cannot explain a child *disengaging* with these information types. This is because it is in principle possible to remain engaged with a stimulus about which you are uncertain. Indeed, as described above, engaging with stimuli about which we are uncertain in an effort to reduce epistemic uncertainty and act accordingly appears to be a relatively stable feature of human decision making (Gottlieb et al., 2013; Parr & Friston, 2017). This considered, primary neurocognitive deficits might in fact be expected to promote preferential engagement, just as blurring images encourages preferential engagement in infants (Chen et al., 2022). Yet this is not what is seen.

The notion of *expected information gain* may be key to reconciling this apparent incongruity between neurotypical and neurodivergent information seeking. The point here is that information seeking is guided not only by information gaps, which suggest an implicit or explicit awareness of a difference between a knowledge state and inferred states of the world, but also by expectations about which information sources are likely to reliably reduce epistemic uncertainty (Addyman & Mareschal, 2013; Baer et al., 2018; Bazhydai et al., 2020; Cittern et al., 2018; Gershman et al., 2015; Gottlieb et al., 2013; Poli et al., 2020, Twomey & Westermann, 2018). For instance, engagement preferences for informative over redundant visual sequences are observed empirically among 5-month olds (Addyman & Mareschal, 2013; see also Poli et al., 2020). Similarly, when solving an object labelling task, 12-month-olds preferentially engage with adult informants who they expect to be reliable rather than those who they expect not to have the required information (Bazhydai et al., 2020). Related work with infants indicates that stimuli that are neither redundant nor complex to the point of resembling noise may be maximally engaging (Kidd et al., 2012).

Collectively this work illustrates how engagement preferences are shaped by the nature of the different information sources that populate a child's environment. Information sources with a distinct and meaningful signal may reliably support learning, while noisy information sources may not (Parr & Friston, 2017; Addyman & Mareschal, 2013). A 'resource-rational' information seeking child should, therefore, preferentially allocate their finite cognitive and motor resources to precise information sources that reliably reduce their epistemic uncertainty and allocate few cognitive and motor resources to imprecise information sources that do not reliably reduce epistemic uncertainty (Gershman et al., 2015; Gershman & Bhui, 2020; Gershman & Burke, 2022). This latter disengagement profile has been termed *rational inattention*; a phrase borrowed from economic theory to describe why rational agents ignore certain information sources that are available to them (Gershman & Burke, 2022). Applied to cognitive psychology, the rational inattention framework proposes that the active avoidance of unreliable information sources with low precision is essential to optimising learning under constraints on time and cognitive and motor resources, enabling the individual to navigate their environment successfully (Gershman et al., 2015). To be clear, it is in our view not necessary for a child to possess formal metacognitive awareness regarding the precision of the many different information sources that populate their environment. Instead, all that is required is for the child to recognise that engagement with a certain information source did or did not effectively resolve epistemic uncertainty in the past and that without evidence to the contrary, this pattern may well hold in similar situations in the future.

The key theoretical contribution of the current report is to endogenize low information source reliability to make it a function of the primary neurocognitive differences that have been proposed to explain neurodivergent profiles, including differences in auditory perception, working memory, procedural learning and the development of the social brain network. Rational inattention is not a theory of the fundamental nature of neurocognitive differences, with respect to which there is a long tradition including much computational modelling work (e.g. Joanisse & Seidenberg, 2003; Johnson et al., 2021; Jones et al., 2023; Thomas et al., 2011, 2019). Instead, rational inattention is a general theory of how such fundamental differences, the nature of which is 'bracketed out' using low endogenous signal precision as proxy, shape children's information seeking and learning. Indeed, given assumed differences in underlying mechanism and complex patterns of comorbidity, it is plausible that a truly transdiagnostic theory of neurodivergent information seeking can only operate at this level of abstraction. Our view is that primary neurocognitive differences render different information sources endogenously imprecise, warranting disengagement with these information sources because they are not reliably conducive to information gain. A child with DLD may disengage from speech because endogenous imprecision renders expected information gain low with respect to speech. Likewise, a child with dyslexia may disengage from orthography and phonology, a child with dyscalculia may disengage from numeric stimuli, and a child with autism may disengage from social cues. Crucially, disengagement of this sort reflects entirely resource-rational behaviour. Imprecise information sources are unreliable and should be avoided by an ideal learner that optimally explores its environment whether the source of imprecision is extrinsic or intrinsic to the learner (Parr & Friston, 2017). This line of reasoning resolves the apparent incongruity between neurotypical and neurodivergent information seeking outlined above, in which neurotypical children preferentially engage with information sources about which they are uncertain and neurodivergent children preferentially disengage with those specific information sources about which they are thought to be uncertain. The point of reconciliation is that *all* children preferentially engage with stimuli that reliably support information gain, and this profile may be moderated by both exogenous and endogenous noise. In this sense, rational inattention is not only unconstrained by primary neurocognitive differences and diagnostic labels (e.g. DLD or dyslexia), but also by the broader *neurotypical* and *neurodivergent* distinction, which may be conceptualised in terms of continuous rather than discrete differences in degrees of endogenous signal noise. Our focus on neurodivergence in this paper therefore reflects our position that the attentional and exploratory behaviour of children with pronounced areas of learning difficulty may be explained under a general principle that traverses neurotypical and neurodivergent development. Namely, the principle of minimising epistemic uncertainty. Once again, we want to emphasise that the neurodivergent child need not possess formal metacognitive awareness that their experience of a given information source is imprecise relative to some abstract baseline or normative standard. Indeed, we consider explicit knowledge of this kind to be implausible. Rather, the child *takes the world as it comes* and engages or disengages with an information source to the degree that it resolves their epistemic uncertainty.

1.2 | Empirical evidence of rational inattention in neurodivergent children

Evidence that rational inattention may be at play in neurodivergent child behaviour comes from a number of sources. For instance, some autism research indicates preferential engagement with non-social over social cues and also enhanced knowledge of non-social cues such as a mechanical systems, maps or timetables (Baron-Cohen, 2009; Moriuchi et al., 2017; Annaz et al., 2009). Findings like this have motivated explanatory frameworks of autism centred around a form of *empathizing and systematizing* distinction (e.g. Baron-Cohen, 2009). Under this view, while socio-cognitive understanding is seen as a relative weakness, the understanding of non-social phenomena, including mechanical systems, is in some cases seen as a relative strength (Baron-Cohen, 2009, p. 71). Such gains may plausibly be attributed to an attentional boost effect (Gershman & Burke, 2022). That is, actively discounting an information source with low endogenous precision (here social signals) frees up cognitive and motor resources for allocation to endogenously precise information sources (here non-social signals) that have relatively high subjective epistemic value, resulting in enhanced learning for those information sources. Attentional boost effects like this are an important prediction of rational inattention theory (Gershman & Burke, 2022) and may be an important signal to probe for in future empirical inquiries into neurodivergent information seeking. The picture is, however, unlikely to be clear cut, with numerous factors including the child's specific constellation of difficulties and their learning environment determining whether attentional boost effects are in fact observed.

Similar effects are reported by Leclercq et al. (2015), who provide perhaps some of the best examples of adaptive engagement preferences in neurodivergent children to date. Leclercq et al. (2015) used what is known as a cross-modal dual-task paradigm, in which a linguistic and a non-linguistic task are administered concurrently. In this study, children with and without DLD aged 9–11 years were tested in a non-word span task and a non-verbal visual search task that required participants to identify the complete circle in an array of broken circles. Importantly, the non-word encoding phase happened at the same time as the visual search task, with a non-word recall test conducted later after the visual search task was complete. A striking feature of Leclercq et al.'s (2015) data was that children with DLD outperformed same-age peers without DLD in the visual search task (note that Kaldy et al., 2011 report a similar advantage in children with autism). Despite the small sample size of Leclercq et al.'s (2015) study, which makes replication essential, this performance advantage is intriguing because there is little empirical basis for assuming that children with DLD are in general better at visual search than their peers. Indeed, a much larger study of 242 children with DLD reported slightly below-population-mean performance in tasks measuring non-verbal cognition (Conti-Ramsden et al., 2012). Furthermore, as Leclercq et al. (2015) note, the apparent visual search performance advantage is not easily reconciled with the attentional capacity bottleneck hypothesis of DLD (e.g. Anobile et al., 2013; Askenazi & Henik, 2010; Ebert & Kohnert, 2011; Rabiner & Coie, 2000). Were a *domain-general* attentional bottleneck causally implicated in these children's language learning difficulties, pronounced deficits may be expected in both the speech and shape tasks. Alternatively, were a *domain-specific* attentional capacity deficit implicated in these children's language learning difficulties (i.e. a deficit in language processing alone), we would predict performance deficits in the speech task, but crucially we would not expect any performance gain in the shape processing task. Only through the adaptive discounting of the subjectively imprecise speech stimuli and the resulting attentional boost for the simultaneously presented shape stimuli, which are expected to reliably support information gain and overall task performance, may we explain this performance advantage. Importantly, Leclercq et al. (2015) report that language-matched control children, too, show a performance advantage relative to age-matched control children in the visual search task. This speaks to the idea that rational inattention is unconstrained by the neurotypical and neurodivergent distinction. Rather, *all* children preferentially engage with stimuli that reliably support information gain. In this instance, both children with DLD and younger children without DLD, who themselves have relatively immature language skills, preferentially engage with visual stimuli plausibly because they find this component of the dual task easier than the verbal component.

Each of the empirical examples described so far in this section illustrates gross differences in engagement preferences across very different types of signals, that is, preferential engagement with non-social over social signals in children with autism, and preferential engagement with visual shapes over auditory speech in children with DLD. These case studies provide good empirical examples of rational inattention in neurodivergent children, and it is this sort of gross difference in engagement profile that we recreate in the computational simulations that follow. However, rational inattention may similarly explain fine-grained differences in attention to specific features of a single information source (Bates et al., 2019; Gershman & Burke, 2022, 2022). For instance, preferential engagement with the semantic rather than phonological features of written text is sometimes seen among children with dyslexia, who characteristically struggle to decode letter sounds (a phenomenon termed *semantic bootstrapping*; Muter & Snowling, 2009). Rational inattention is, therefore, assumed to operate at various levels of granularity both within and across information sources, reflecting the child's efforts to maximise their learning progress in order to achieve their subjective aims. This means that even when a neurodivergent child appears to be engaged with a given information source, they may be engaged with that information source in a very different way than their neurotypical peers, with very different consequences for their learning over time. Sims (2016) describes this sort of variation in engagement preference in terms of *cost function mismatch*, where the term 'cost function' describes the uncertainty (or error) that the individual is trying to reduce through engagement and learning. Here, the claim is that individuals faced with matched stimuli in a matched environment will differ in how they engage with that material as a function of expected information gain given their subjective aims. Crucially, Sims (2016) demonstrates empirically that how individuals are *expected to engage* (e.g. by an experimenter) and how subjects *actually engage* can differ quite dramatically. Such differences in engagement are of course not unique to neurodivergent children – rational inattention is a general principle and there will be variance in how *all* children attend to a shared array of environmental stimuli. However, the empirical research touched on in this section indicates that specific neurobiological obstacles to learning may force specific, relatively unconventional modes of engagement that may be as pronounced as an aversion to the eyes or as subtle as semantic bootstrapping.

1.3 | Rational inattention and long-term learning

The paradox at the heart of the theoretical framework developed in this report is that even though disengaging with information sources associated with low endogenous precision is 'optimal' or 'rational' in the short-term, it is expected to have a detrimental effect on the child's overall learning trajectory. Thus, while rational inattention is resource-optimal, it comes with a cost. This is because rational inattention precludes practice with difficult information sources, which intervention studies suggest are often in principle learnable (e.g. Rinaldi et al., 2021), and so may, therefore, deepen long-term learning delays that negatively affect educational outcomes and wellbeing (Conti-Ramsden et al., 2018). For instance, a child with DLD may reasonably attend to peer behaviour (an endogenously reliable cue) over verbal instructions (an endogenously unreliable cue) in their classroom to rapidly reduce their epistemic uncertainty and determine what they have to do at the moment. However, this engagement profile (sometimes termed a *compensatory strategy*[2]; Muter & Snowling, 2009) may come with the long-term cost of fewer well-developed memories of natural speech.

Negative cycles of learning like this are often referred to as a (negative) Matthew effect, in reference to the ideas that 'the rich get richer' and 'the poor get poorer' (Stanovich, 2009). The Matthew effect is perhaps most commonly cited in dyslexia research, where it is argued that children who struggle to read engage less with text and so learn fewer grapheme–phoneme mappings and words than their peers, which in turn makes it more difficult to read (Soriano-Ferrer & Morte-Soriano, 2017; Stanovich, 2009). To the extent that we consider the active avoidance of an information source to be an important contributing factor in the delayed learning trajectory of a neurodivergent child, there is an overlap between the rational inattention framework developed in this report and the notion of the Matthew effect. In this sense, the current report may be understood as making the case for greater transdiagnostic awareness of the Matthew effect; an important idea with respect to which there is already some work (e.g. Foster, 2023).

Yet, there is also a critical difference between the rational inattention framework developed in this paper and the Matthew effect as it is commonly described. Namely, much of the literature citing the Matthew effect argues that a failure to acquire a given skill reduces the child's implicit motivation to engage with a relevant information source, further curtailing learning (Pfost et al., 2014; Soriano-Ferrer & Morte-Soriano, 2017). Put plainly, the assumption is often that the child *does not like*, for instance, reading and so does everything in their power to avoid it. In contrast, the rational inattention framework promoted in this paper makes no assumptions regarding a child's affective motivation to disengage with a specific information source. We do not doubt that this dynamic is also at play. However, our claim is instead that actively discounting information sources associated with low precision is resource-rational behaviour for any agent that is trying to learn maximally about its environment, whether the source of imprecision is extrinsic or part of the individual's neurocognitive makeup as it is in neurodivergent children. This is not a question of reduced enthusiasm stemming from failures to acquire a specific skill, but of acting rationally to reduce epistemic uncertainty in order to achieve subjective aims given finite cognitive and motor resources and a neurobiological obstacle that restricts the endogenous precision of certain information sources. Rational inattention theory entails a description of neurodivergent child behaviour as optimal in the face of primary neurobiological constraints, rather than as *deficient* or *demotivated*.

This trade-off between short-term effectiveness and long-term costs is not unique to the novel application of rational inattention theory pursued in this report, though it is perhaps thrown into sharp relief by the application of rational inattention to neurodivergent child development. On the contrary, the tendency to allocate cognitive and motor resources in a way that secures immediate rather than deferred gain is well recognised (Gershman & Bhui, 2020; Gershman & Burke, 2022). People act to reduce epistemic uncertainty because they want and need to get things done in the here and now, and the so-called temporal discounting of expected long-term information gain likely reflects the drive to respond to immediate curiosities and challenges coupled with a difficulty envisioning how current decisions will predict outcomes way down the line (Gabaix & Laibson, 2017; Gershman & Bhui, 2020; Gershman & Burke, 2022). Rational inattention remains 'rational' because it implies an optimal way to gather information so that we can act accordingly within a reasonable time horizon – whether out of implicit or explicit motivation – whatever the consequences for long-term learning.

2 | SIMULATIONS

This project is associated with a fully documented Open Science Framework repository, which enables readers to replicate our simulations and experiment further with our model: https://osf.io/jey7q/.

2.1 | Experiment 1: Simulating rational inattention in a complex learning environment

In this section, we use a computational model to demonstrate that disengagement with an information source emerges naturally as the optimal solution to the reduction of epistemic uncertainty given constraints on endogenous precision. Several mathematical and computational frameworks have been used to study human information seeking, and the current manuscript shares with some of these frameworks a description of human behaviour as resource rational (Bates et al., 2019; Feldman & Friston, 2010; Gershman & Burke, 2022; Gottlieb et al., 2013; Oudeyer et al., 2007; Parr & Friston, 2017; Twomey & Westermann, 2018). In the current study, we take a connectionist approach, using artificial neural networks comprising simple nodes and connecting weights that are a broad computational analogy to biological neurons and synapses. The weights linking nodes adapt dynamically to the input that the network receives, mimicking biological synaptic plasticity in response to environmental stimulation. By controlling features of the input and the network architecture, we are able to develop and test theories of emergent cognitive phenomena and child behaviour (Mareschal & Thomas, 2007; Twomey & Westermann, 2018; Westermann et al., 2009; Westermann & Ruh, 2012).

2.2 | Model architecture

We modelled neurotypical and neurodivergent information seeking using autoencoder neural networks (Figure 1). Autoencoders have been used to simulate a wide range of child behaviours, from categorisation and visual object processing to curiosity-driven learning and language acquisition (Jones & Brandt, 2020; Mareschal et al., 2000; Twomey & Westermann, 2018; Westermann et al., 2009; Westermann & Mareschal, 2004, 2012). Autoencoders are a class of self-supervised neural network, which form representations of their learning environment by adaptively updating internal weights to minimise the difference between the input they receive and a re-construction of that input that they produce. The difference between network input and output is quantified as a mean squared error (MSE). For instance, a 28×28 black and white bitmap image may be passed to the network, transformed across the networks hidden layers and then output as a 28×28 re-constructed version of the input image[3]. The hidden layers of an autoencoder are smaller than the input and output layers and this enforces a processing bottleneck which means that reconstruction is normally imperfect. The MSE is the sum of the squared differences between each input pixel value and each corresponding output pixel value. Over time (or training *epochs*), an error is reduced by updating (raising or lowering) connection weights across the autoencoder using the backpropagation algorithm. This results in a distributed internal representation of the input stimuli that is optimal given the network's capacity and the task at hand, a notion that aligns well with the principle of efficient neural encoding (Barlow, 2012). Autoencoders therefore constitute an ideal choice for modelling optimal learning under constraints on cognitive and motor resources (Sims, 2016; Twomey & Westermann, 2018; Westermann et al., 2009; Westermann & Ruh, 2012).

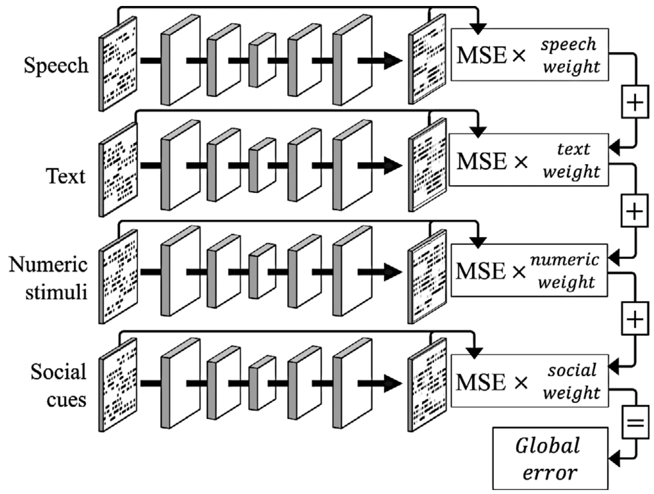

FIGURE 1 Four-path autoencoder neural network.

Figure 1 illustrates the structure of the autoencoder used in this first series of simulations. This is a four-path autoencoder, which takes in and reconstructs abstract schematic inputs from four different information sources: speech, text (i.e. orthography-phonology mappings), numeric stimuli and social cues. The network reconstructs each input from each information source and calculates the corresponding MSE. Each MSE is then added together to produce a global error, which is the quantity that the system minimises through backpropagation-driven parameter updating. Our intention here is to capture the notion that the agent seeks to reduce their overall uncertainty (i.e. global error) by engaging selectively with their environment. To achieve this, a key innovation of our computational model is that the MSE from each information source is weighted (i.e. by the speech weight, text weight, numeric weight and social weight). These weights are normalised to sum to one across the network and so are initially each set to 0.25 (i.e. $\frac{1}{n}$, where n is the number of channels). The global error is, therefore, the

Developmental Science

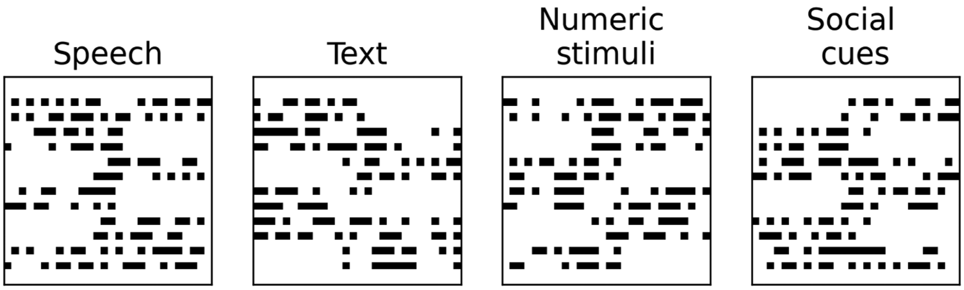

FIGURE 2 Abstract input patterns, 'schematic spike trains'.

sum of each MSE times the weight for each information source:

$$\begin{aligned} \textit{Global error} \ = \ & \textit{speech MSE} \times \textit{speech weight} + \textit{text MSE} \times \textit{text weight} \\ & + \textit{numeric MSE} \times \textit{numeric weight} + \textit{social MSE} \\ & \times \textit{social weight} \end{aligned}$$

Over 100 training epochs, the weights connecting hidden layers are optimised using the backpropagation algorithm to reduce the global error. Crucially, the speech MSE weight, text MSE weight, numeric MSE weight and social MSE weight – what we might collectively call the *engagement weights* – are also optimised during training via the backpropagation of global error. The network therefore dynamically optimises the weight for each information source as a function of how much engaging with that information source reduces the global error, that is, how much engagement enhances the network's overall ability to represent information in its environment. If each information source is equally epistemically valuable in terms of its contribution to the reduction of global error, each information source engagement weight will remain at its initialisation value of 0.25. However, if one information source is less epistemically valuable, then that information source will be de-weighted, meaning that its MSE feeds less into the pooled global error that is backpropagated across the network in order to shape overall network structure. The fact that the engagement weights are normalised to one captures the idea of finite cognitive and motor resources. Preferentially engaging with one information source necessarily entails dis-engaging with another information source to some degree. In essence, the domains of speech, text, numeric stimuli and social cues are competing for attention to enable the learner to maximise their learning progress. The neural network shown in Figure 1 is, therefore, able to simulate fully emergent engagement preferences that are driven by optimising overall learning in a complex environment given resource constraints.

2.3 | Stimuli

In this first series of simulations, we used abstract 28 × 28 bitmap images as model input. These inputs may be thought of as schematic spike trains that are associated with stimuli from the different information sources that we are modelling, namely speech, text, numeric stimuli and social cues (Figure 2). The input schematic spike trains elicited by stimuli from each information source had a unique general pattern. For instance, the speech spike trains followed the pattern (from top to bottom) long bar, short left-aligned bar, short right-aligned bar, short left-aligned bar, short right-aligned bar and long bar (see the 'speech' panel in Figure 2). However, each spike train presented to the network was unique in terms of the specific location of its scattered black pixels, the number of which was identical across input exemplars. During training, each network was exposed to 100 randomly generated unique input exemplars from each of the four information sources.

In order to simulate neurotypical and neurodivergent information seeking, we controlled the degree of input noise that networks experienced to align with canonical areas of difficulty commonly associated with various neurodivergent phenotypes. Noise was generated by randomly redistributing a percentage (0%–25%) of the input's black pixels into whitespace. The result of this process is shown in Figure 3. A baseline, neurotypical model received input with zero noise in each information source. Neurodivergent phenotypes were then modelled by adding noise to one or more information sources to align with the areas of difficulty that a child with a certain neurodevelopmental disorder may be reasonably expected to have. To simulate DLD, for instance, we added noise to speech, to simulate dyslexia we added noise to text (orthography-phonology mappings), to simulate dyscalculia we added noise to the numeric information source, and to simulate autism we added noise to the social information source. Comorbidity was simulated by adding noise to multiple information sources concurrently. Once again, we note that rational inattention is unconstrained by diagnostic labels and indeed by the broader *neurotypical* and *neurodivergent* distinction. The labels used throughout these simulations (e.g. DLD) are of course arbitrary and simply reflect different degrees of endogenous noise in a canonical information channel (e.g. speech). We consider this necessary to clearly build intuition regarding how a generalised learning principle can help shape very different developmental trajectories.

The purpose of designing stimuli at this level of abstraction is to remain agnostic about the underlying nature of the various neurocognitive differences that can affect child development, which remains an area of much controversy. For instance, as we touched on above, whether DLD is the result of auditory perceptual deficits or deficits in working memory or procedural learning remains keenly debated. Each of these cognitive mechanisms is aligned with a different neural

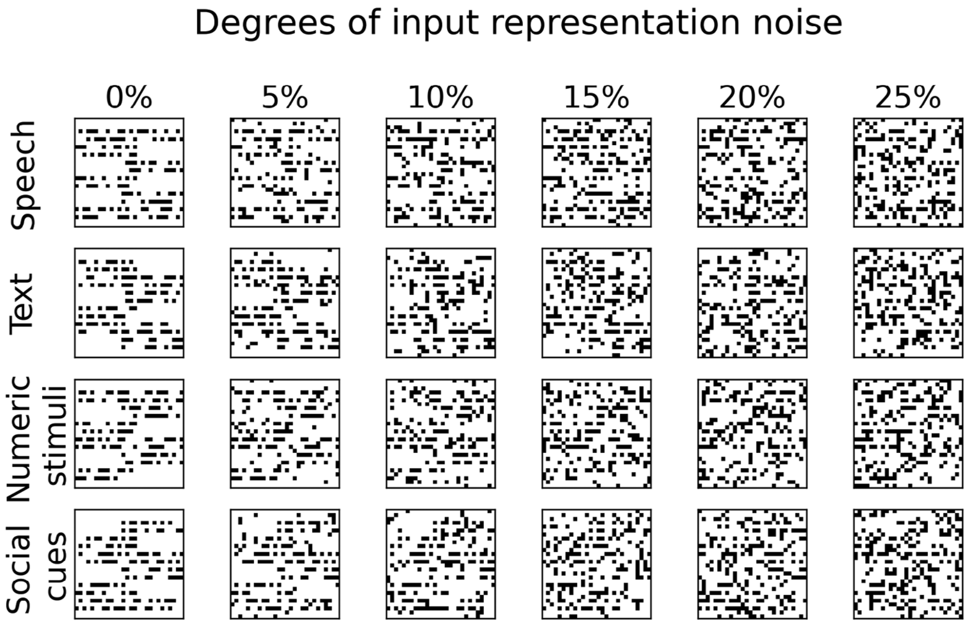

FIGURE 3 Input schematic spike trains with different degrees of noise.

substrate (e.g. the basilar membrane or frontal-basil ganglia circuits). Moreover, comorbidity is the norm and not the exception in neurodivergent development, a factor which further complicates theorizing about the nature of the complex constellation of mechanisms that underly neurodiverse phenotypes (Astle et al., 2022). Accordingly, in the simulations that follow, we take the very general view that a neurodivergent child has relatively noisy representations of the information source – or multiple information sources – that sits at the heart of their condition *at some level*. Focussing again on canonical areas of difficulty: Children with DLD commonly demonstrate uncertainty about speech; children with dyslexia demonstrate uncertainty about orthographic-phonemic mapping (which we label *text* for brevity), children with dyscalculia demonstrate uncertainty about numeracy, and some children with autism demonstrate uncertainty about social cues. Throughout these simulations, we are asking *what does optimal information seeking look like given generalised noisy representations affecting disparate information sources?* without making any claims about the primary nature or neurobiological locus of the underlying neural noise.

2.4 | Results and discussion

To show that networks are indeed forming the types of internal representations of their input that we expect, we begin by presenting the raw inputs and output reconstructions from a single network modelling each phenotype. In the results shown, networks simulating neurotypical information seeking received clean input stimuli from each information source. Networks simulating dyscalculia and autism were affected by 25% noise in the numeric and social information streams respectively. Comorbid DLD and dyslexia were simulated by adding 10% noise to both the speech and text (orthographic-phonemic) information sources. Figure 4 shows that the trained autoencoders in each group are indeed generally picking up on the underlying patterns which are common to each of the four classes of randomised schematic spike trains that they received as input (see Figure 2). Looking carefully, you may also be able to identify a greater number of dispersed grey pixels in the re-constructed spike trains from the relevant noisy information source (or sources) in each neurodivergent simulation (e.g. with respect to numeric stimuli in dyscalculia). This is evidence that the associated input noise is constraining learning.

Figure 5 shows dynamic changes in the weighting of each information source over time, which as a reminder was initialised at 0.25 across information sources at the outset of learning. This point is important because it means that there is no prior engagement preference for any information source and that engagement preferences are an emergent property of maximising overall learning. The results shown represent the average engagement profile over time for 10 networks in each condition and may be read as exploratory trajectories through time in a complex learning environment containing speech, text, numeric stimuli and social cues.

Figure 5 confirms that engagement with each information source is equivalent at the outset of learning regardless of the simulated phenotype (i.e. weights are uniformly initialised at 0.25; see the initial value at the *y*-axis in each panel). Yet, engagement becomes increasingly asymmetrical in the neurodivergent simulations as a function of learning optimisation (i.e. minimising global error) given endogenous noise. For instance, the networks simulating dyscalculia and autism disengage with numeric and social signals entirely by around 60 epochs, as neural noise renders these information sources not conducive to information gain. Similarly, the networks modelling comorbid DLD and dyslexia slowly disengage with speech and text in order to maximise their overall learning. In the absence of any constraints on disengagement,

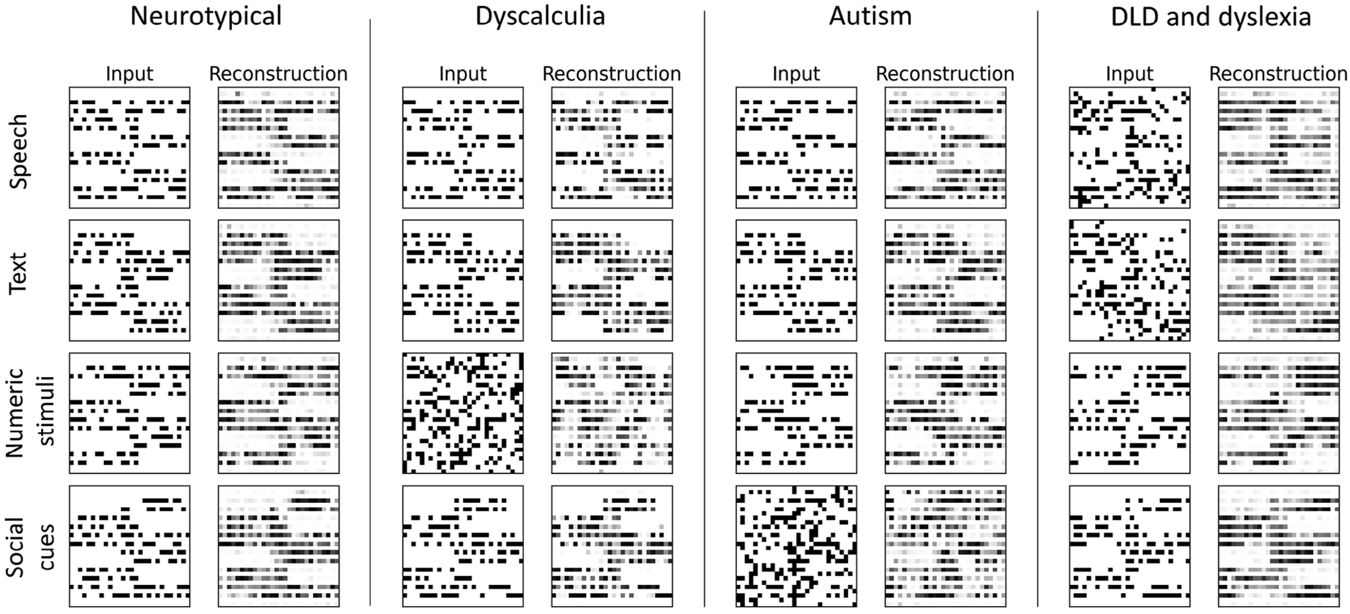

FIGURE 4 Example inputs and reconstructions from a randomly sampled trained network in each group.

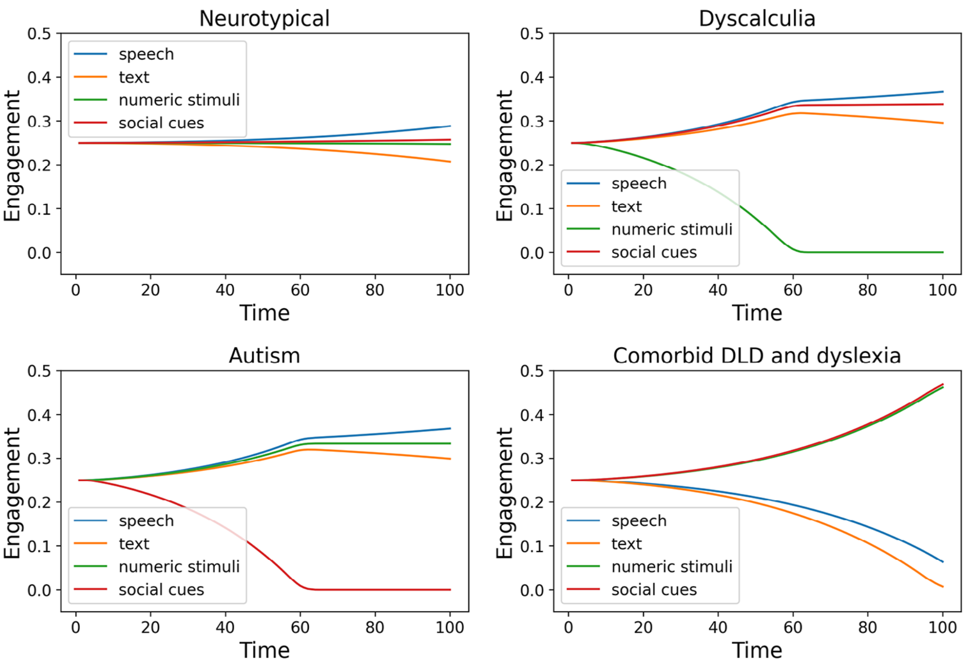

FIGURE 5 Engagement by phenotype over time.

ideal preferential engagement weights will reach asymptote at approximately $\frac{1}{n}$, where n is the number of information sources unaffected by neural noise.

Importantly, the engagement profiles shown in Figure 5 are in a cyclical relationship with network learning, which is shown by the simulated phenotype in Figure 6. Figure 6 shows the mean squared re-construction error rates associated with stimuli from each information source over time, given a certain canonical input noise profile that aligns broadly with each phenotype (see above). Error is understandably higher for the relevant noisy information source, and this effect deepens learning delays over time (i.e., error rates may steadily increase) in line with the disengagement profiles shown in Figure 5.

In conjunction, Figures 5 and 6 illustrate the cyclical relationship that exists between engagement and learning. Reducing epistemic uncertainty fosters continued preferential engagement and, when there is still more to learn, preferential engagement further reduces epistemic uncertainty. On the other hand, networks disengage with information sources that do not reduce epistemic uncertainty, deepening long-term learning deficits with respect to specific information sources.

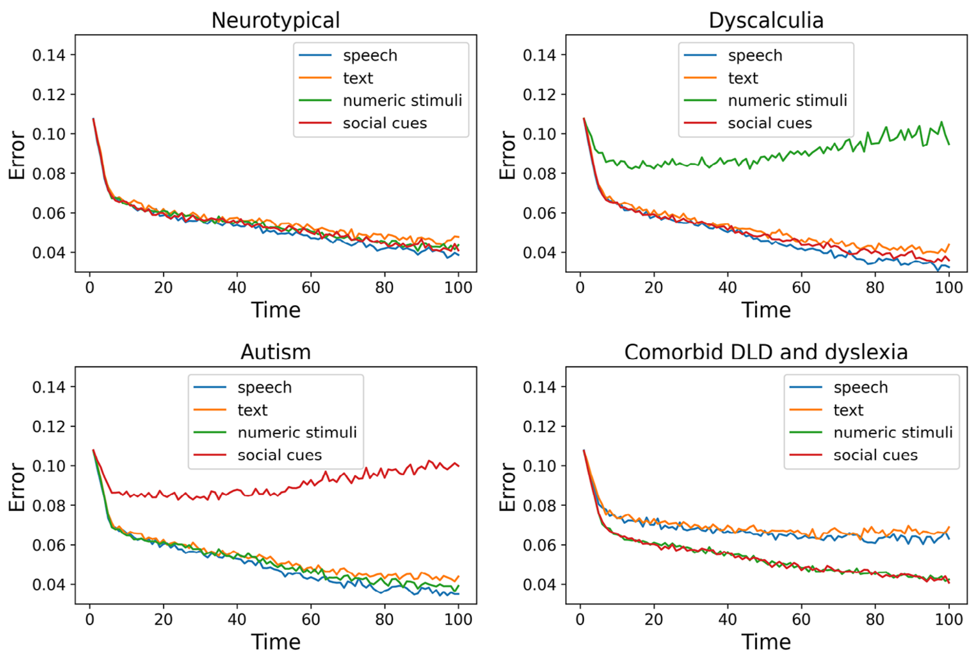

FIGURE 6 Learning (error reduction) by phenotype over time.

Using broadly the same setup, we conducted further simulations to demonstrate that rational inattention is essential to maximising overall learning given finite attentional resources and constraints on endogenous precision. These simulations involved two conditions: active learning and clamped learning. In the active learning condition, the procedure was identical to that described so far in this section. That is, the error from each information source is weighted by the relevant *engagement weight* before being fed into the pooled error and then backpropagated. This is a dynamic learning condition in which individual information sources can be engaged with more or less in order to maximise overall learning. In contrast, the new clamped learning condition involved fixing each engagement weight (i.e. speech weight, text weight, numeric weight and social weight) to 0.25 throughout training. This, then, is a non-active learning simulation, in which the artificial agent must allocate 25% of its resources to each information source and attempt to learn from it regardless of the degree to which that information source is immediately uncertainty reducing. Rationally inattentive information seeking is not permitted in this clamped condition. The results of this comparison between active learning and clamped learning are shown for a simulated neurotypical agent and a simulated DLD agent below. Results again reflect the mean learning profile of 10 networks in each group, and in the DLD group input noise on the speech channel was set to 25%.

Figure 7 (Neurotypical) shows that whether the neurotypical agent is engaged in clamped or active learning has little effect on reducing global error. This is because each information source in the simulated learning environment is precise and comparable and therefore equally epistemically valuable. On the other hand, Figure 7 (DLD) shows that active information seeking – in this case, the active avoidance of endogenously noisy speech input – is essential to reducing overall uncertainty about the complex learning environment in which the neurodivergent agent is situated. Figure 7 shows rational inattention to be a hallmark of optimal, resource-rational neurodivergent information seeking and learning. If the DLD agent is forced via clamping to engage symmetrically with all four information sources, overall error will be higher than if it is free to engage and disengage on the basis of expected information gain. This effect can be unpacked further by looking at the mean individual information source error rates for the DLD simulation networks in the clamped and active learning conditions (Figure 8).

Figure 8 shows that active learning promotes relatively low error rates for precise information sources (text, numeric stimuli and social cues in Figure 8; the circle markers have lower error than the cross markers) but that active learning promotes relatively high error rates for imprecise information source (speech). This is the result of the active learning networks being able to freely disengage with imprecise information sources, highlighting our central claim that optimal in-the-moment information seeking may, paradoxically, contribute negatively to an already delayed learning trajectory in neurodivergent children (see the upward trend in the circle-marked speech line relative to the cross-marked speech line in Figure 8). Conversely, clamping engagement weights – promoting sustained engagement – results in better learning for imprecise information sources (in this instance speech).

The critical point, then, is that rational inattention characterises optimal information seeking given finite attentional resources and constraints on endogenous signal precision, providing proof of concept for the rational inattention framework developed in the current report. Rational inattention in this context cannot be attributed to any lack of affective motivation to engage with an information source or indeed to any form of processing capacity bottleneck, neither of which features in our model. Instead, each of the exploratory trajectories shown

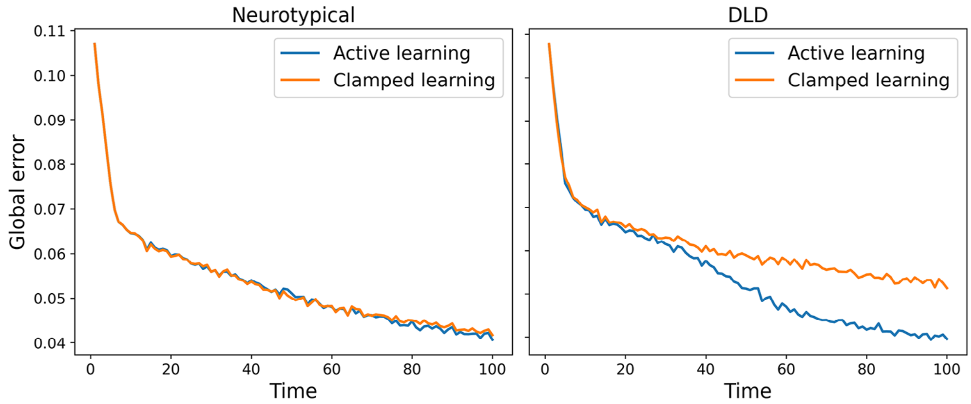

FIGURE 7 Clamped and active neurotypical and neurodivergent learning.

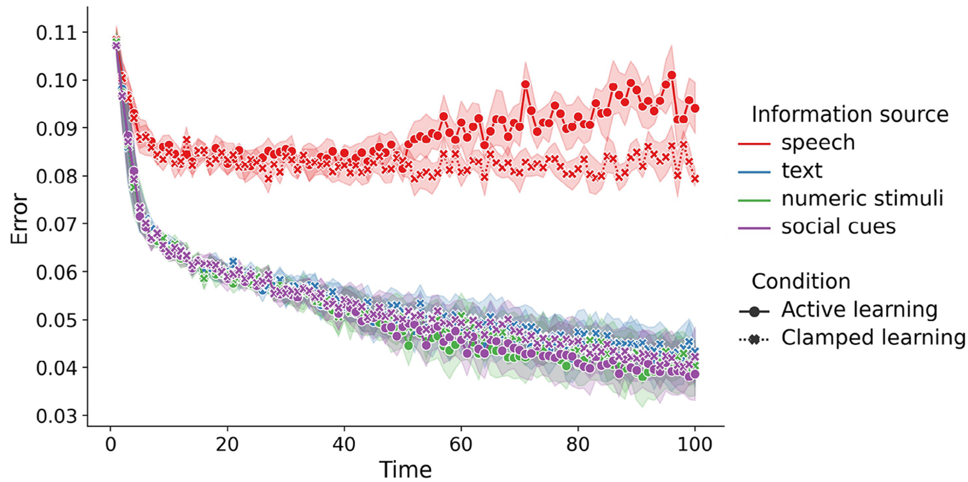

FIGURE 8 DLD simulation information source error rates in clamped and active conditions.

in Figure 5 (and indeed throughout this section), both neurotypical and neurodivergent, captures entirely optimal, resource-rational behaviour. Nevertheless, optimal in-the-moment information seeking can be shown to deepen learning delays over time, particularly when direct engagement is not encouraged (Figure 8).

2.5 | Experiment 2: Simulating rational inattention and attentional boost effects in the cross-modal dual-task paradigm

In this section, we 'zoom in' to simulate information seeking in a relatively simple learning environment containing just two information sources: speech and visual stimuli. The simulations presented in this section aim to re-create the cross-modal dual-task paradigm data collected by Leclercq et al. (2015), who tested children with and without DLD in a non-word span task and a non-verbal visual search task that required participants to identify the complete circle in an array of broken circles. In this study, the non-word encoding phase happened at the same time as the visual search task, with the non-word recall task then conducted afterwards. The striking feature of Leclercq et al.'s (2015) results is that children with DLD outperformed same-age children without DLD in the visual search task. The relevant data from Leclercq et al. (2015, p. 727, tab. 2) are reprinted in Table 1.

TABLE 1 Non-word repetition task performance (proportion of syllables accurately repeated) and visual search task performance (proportion of visual targets identified) in a cross-modal dual-task paradigm. The values shown are means, with standard deviations in brackets. Controls = age-matched controls aged 9–11.

Task	DLD ($n = 21$)	Controls ($n = 21$)
Non-word repetition	0.47 (0.16)	0.52 (0.23)
Visual search	0.83 (0.13)	0.73 (0.13)

The empirical results shown in Table 1 – namely the so-called *attentional boost effect* – may be explained by expected information gain. That is, age-matched neurotypical (control) children expect to learn from both visual and verbal information sources, and so engage symmetrically with each modality. In contrast, children with DLD find the visual task easier than the verbal task, and so preferentially engage with the visual stimuli, resulting in an attentional boost effect for visual search. Given that the neurocognitive makeup of children with DLD means

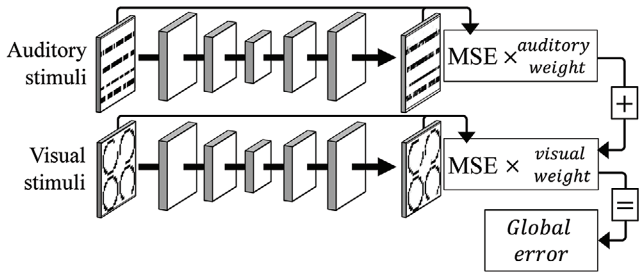

FIGURE 9 Two-path autoencoder.

that they tend to have difficulty with speech processing, this constitutes entirely resource rational behaviour – *rational inattention* – and enables the child to maximise their overall learning given the task at hand.

2.6 | Model architecture

The autoencoder neural network used in this section was identical to that shown in Figure 1, with the exception that it was capable of processing just two information sources: auditory verbal stimuli and visual shape stimuli (Figure 9). This meant that the network's two engagement weights – that is, the auditory weight and the visual weight – were initialised to 0.5 (i.e. $\frac{1}{n}$, where n is the number of channels).

Leclercq et al.'s (2015) paradigm required children to produce two responses. In the visual shape task, children were required to point to the complete circle in a visual array, and in the auditory verbal task, children were required to verbally reproduce non-words that they heard during the encoding phase. In the simulations that follow, these responses (i.e. pointing and non-word articulation) are inferred from the mean squared reconstruction errors for visual and auditory verbal inputs respectively. A low visual reconstruction error means that the system has formed a robust internal representation of the detail of the visual stimuli that it has encountered, including the location of the accurate target and neighbouring distractors. Similarly, a low auditory verbal reconstruction error means that the system has formed a robust internal representation of the detail in thespeech stimuli that it has encountered, one that would be necessary to accurately articulate the non-word at test.

2.7 | Stimuli

Stimuli were once again 28 × 28 bitmap images. However, the images used in this section were engineered to broadly replicate those used by Leclercq et al. (2015). Over 100 training epochs, the two-path autoencoder was presented concurrently with randomly generated schematic spectrograms (auditory input) and randomly generated arrays of four circles (visual input; see Figure 10). The auditory inputs represented non-words from four classes (Figure 10, schematic auditory spectrogram patterns zero to three). Within each class, the location and width of the synthetic formants (i.e. the high-energy density regions of the spectrogram) were pre-defined. However, each spectrogram exemplar generated had randomly distributed breaks of white space across these formants that made it unique. The visual inputs comprised 28 × 28 bitmap images containing four circles, one of which was complete and three of which had a break in a randomly generated location for each exemplar. The location of the complete circle also differed across training exemplars, whether in the top left, top right, bottom left or bottom right. Networks were exposed to 100 training exemplars in each input modality: 25 speech exemplars from each of the four classes of non-word and 25 visual arrays with the complete circle in each of the four quadrants.

Following the procedure described in the first series of simulations, we then degraded the auditory input representations shown in Figure 10 at different rates. This meant re-distributing a percentage of black pixels into whitespace (Figure 11). Once again, this reflects the general notion that children with DLD have canonical difficulties with speech processing and should not be interpreted as a commitment to any claim about the neurobiological nature or locus of these children's difficulties. Rather, we are asking *what does optimal information seeking look like given generalised noisy representations affecting speech processing?* Specifically, our simulations aimed to determine whether optimal information seeking can in principle explain the behavioural profile reported by Leclercq et al. (2015), namely the attentional boost effect (Table 1).

3 | RESULTS AND DISCUSSION

Figure 12 shows sample inputs and reconstructions from neural networks with different degrees of auditory neural noise. Like Figure 4, the purpose of this figure is to demonstrate what exactly it is that networks are learning in each condition. Remember that each input passed to the network is generated on the fly and is entirely unique. Yet, after 100 epochs of dynamic adaptation, networks have learned a general distribution that enables them to reproduce each new visual or auditory input encountered more or less accurately (given the consistency of underlying patterns; Figure 10). Trained networks are able to accurately represent the location of the complete circle and, in the absence of excessive neural noise, the key energy bands in the synthetic auditory spectrograms.

Increased re-construction error for auditory stimuli is visible in Figure 12 as auditory noise levels rise from 0 to 25%. Once again, this is evidence of endogenous noise necessarily inhibiting learning. The networks are unable to learn the underlying input distribution, which is simply too distorted, and so produce 'fuzzy' re-constructions. As emphasised in the previous simulations, learning interacts with engagement profile, with learning successes fostering preferential engagement and vice versa. Figure 13 shows engagement preferences across auditory speech stimuli and visual stimuli in the simulated cross-modal dual-task paradigm as a function of auditory noise level. Lines represent means, with n = 10 networks in each group. The take-home message is that, as in the preceding simulations, noise fosters

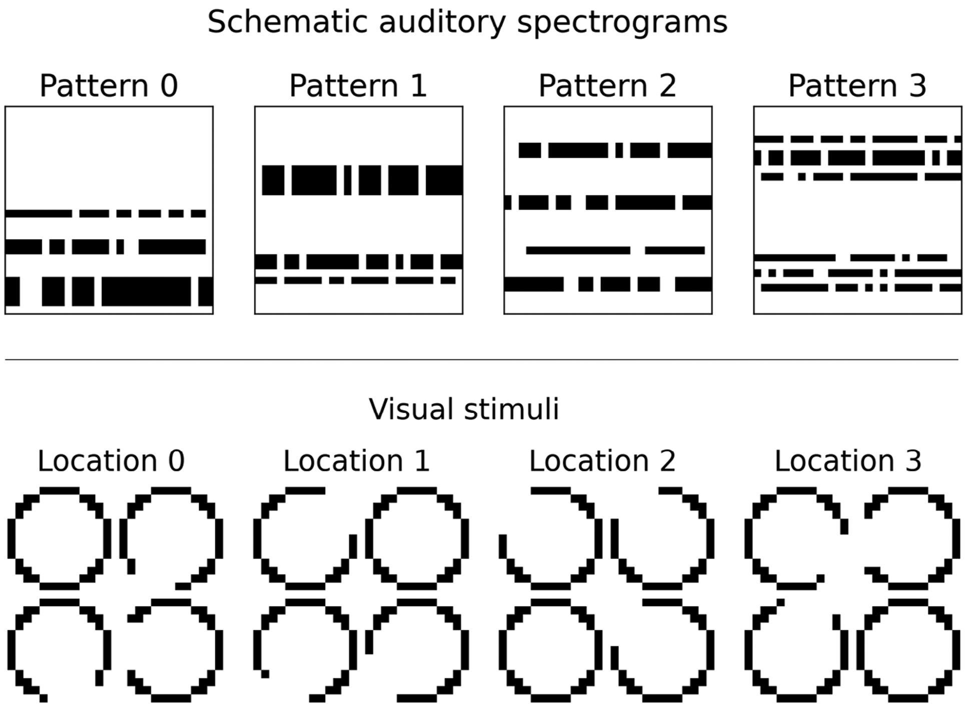

FIGURE 10 Auditory and visual input used in the cross-modal dual-task paradigm simulations.

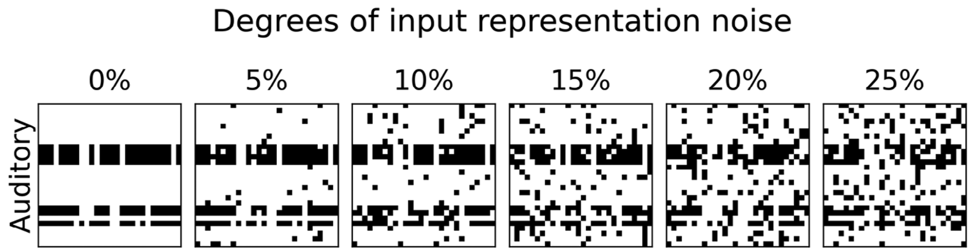

FIGURE 11 Input auditory representations with different degrees of noise.

resource-rational preferential disengagement (here, with the auditory information stream).

Interestingly, networks with no auditory noise ($n = 10$) first preferentially engage with auditory stimuli (as do those with 5 and 10% noise), suggesting that engagement with the auditory spectrograms is initially the most expedient way to reduce global error. Later, at around 80 epochs when learning is more advanced, networks appear to shift towards preferential engagement with the visual stimuli. This particular simulation therefore appears to capture synthetic exploitation–exploration behaviour.

As soon as significant noise enters the auditory stream networks become rationally inattentive to that stream and instead preferentially engage with the visual information source in order to optimise their overall learning. In the extreme case of a noise level of 25%, this occurs quite dramatically within 40 epochs, by which time the ($n = 10$) neural networks in this group engage exclusively with the visual stimuli. Figure 13 therefore neatly captures how rational inattention is unconstrained by any hard-and-fast *neurotypical* and *neurodivergent* distinction and may instead be conceptualised in terms of a response to continuous degrees of endogenous signal noise.

Of course, in the child's environment, there are factors which will preclude such absolute disengagement with an information source like speech, such as a need to communicate with caregivers and later to participate in educational tasks at school. Such encounters may be broadly aligned with the notion of clamped engagement that we explored in Experiment 1 because they limit the extent to which disengagement is possible and in doing so may mitigate long-term learning deficits (Figure 8). Formal programmes of clinical intervention may confer similar effects by facilitating deep, structured engagement with difficult information sources. Engagement preferences are expected to become increasingly symmetrical as the quality of the child's long-term speech representations improves. This may be expected to take the child from a profile resembling the higher noise engagement trends shown in Figure 13 to a profile resembling the lower noise engagement trends shown in Figure 13 (see also Figures 5 and 8).

The rationally inattentive behaviour shown in Figure 13 is in a cyclical relationship with learning. Figure 14a shows auditory task error by auditory (i.e. speech) noise level. As in the human data from Leclercq et al. (2015), generalised noisy representations affecting

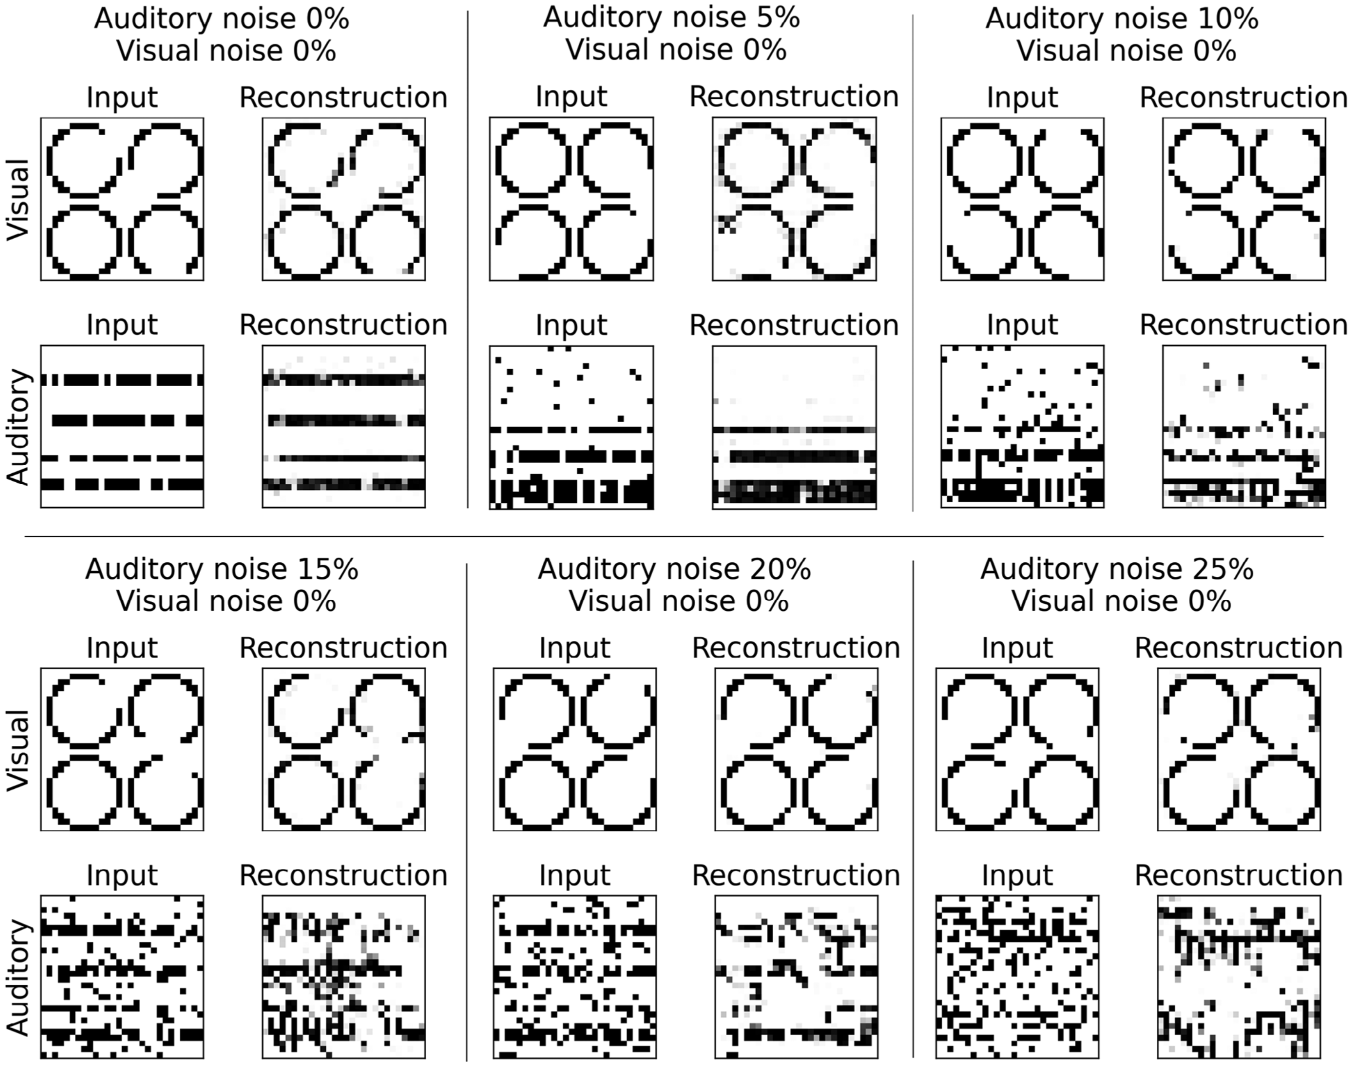

FIGURE 12 Example inputs and re-constructions from a randomly sampled trained network at each noise level.

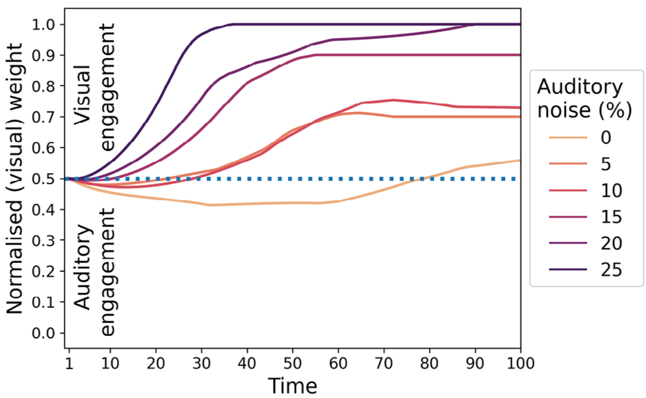

FIGURE 13 Engagement preferences in the simulated cross-modal dual-task paradigm as a function of synthetic speech processing deficits.

speech processing unsurprisingly correspond to worse performance on the auditory task. Importantly, though, Figure 14b shows that simulation using the dual-path autoencoder does recover the apparent attentional boost effects seen in Leclercq et al. (2015; see Table 1 of the current study). That is, higher auditory stream noise was associated with higher rates of success on the visual search task (i.e. visual task error is lower at higher auditory verbal noise levels). Each pattern of findings shown in Figure 14 of course relates to the preferential disengagement with the auditory signal and the preferential engagement with the visual signal that are illustrated in Figure 13; a pattern that can itself be driven only by expected information gain given endogenous noise.

The results of the simulations in this and the previous section align neatly with an expansive behavioural and computational literature on resource-rational information seeking (Bates et al., 2019; Bazhydai et al., 2020; Chen et al., 2022; Feldman & Friston, 2010; Gershman et al., 2015; Poli et al., 2020; Sims, 2016; Twomey & Westermann, 2018). In summary, agents preferentially engage with information sources associated with high expected information gain and are rationally inattentive to information sources associated with low expected information gain. Importantly, this principle holds whether

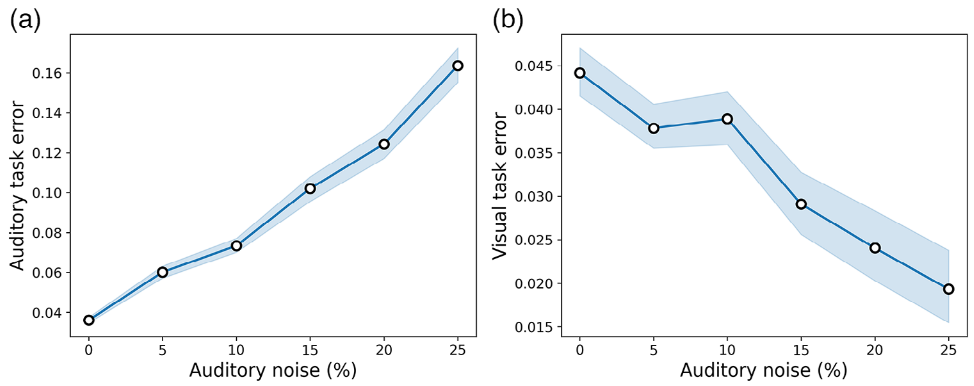

FIGURE 14 Auditory and visual task error by the degree of auditory verbal noise.

the noise source that renders an information source not conducive to information gain is exogeneous or endogenous to the agent (Parr & Friston, 2017). Rational inattentive information seeking may, where the learning and testing environment is appropriately constrained and appropriately configured, be evident in attentional boost effects (Figure 14b).

4 | GENERAL DISCUSSION

Many of the ideas considered in this paper have a long history, from the idea that neurodivergent children allocate attention unconventionally to notions of compensatory strategies and the Matthew effect (Muter & Snowling, 2009; Stanovich, 2009). This paper has presented rational inattention as a coherent scheme for integrating and developing these ideas, and for putting the under-researched notion that neurodivergent children play an active role in shaping their learning front and centre. Our focus in this report was on information seeking, which, as we emphasised at the outset, may be driven by numerous extrinsic or intrinsic factors. We presented a series of computational simulations in which rational inattention emerged naturally in artificial agents mimicking specific neurodivergent phenotypes and qualitatively replicated child empirical data from Leclercq et al. (2015).

This paper began by describing an apparent incongruity between a well-reported drive to resolve epistemic uncertainty in neurotypical children and the active avoidance of difficult information sources that is commonly seen in neurodivergent children. We argued that this incongruity can be resolved through the notion of *expected information gain*. Actively disengaging with imprecise information sources is resource-rational behaviour for any agent with finite cognitive and motor resources that is trying to learn maximally about its environment in order to achieve its aims, whether the source of imprecision is extrinsic or part of the individual's neurocognitive makeup, resulting, for instance, from a fundamental neurobiological obstacle that perturbs endogenous precision. Rational inattention therefore provides a unified framework in which both neurotypical and neurodivergent attention and engagement preferences are shaped by expectations about the degree to which the different information sources that populate the child's environment support information gain. Importantly, neurodivergent children do not need to meta-reason that their sensory experience is imprecise relative to an abstract and indeed inaccessible baseline. All children need to know is that in general engagement with a certain information source does or does not resolve their epistemic uncertainty.

Strong and lasting engagement preferences can be induced by apparently subtle perturbations in exogenous stimulus precision (Bates et al., 2019; Bourgeois et al., 2016; Chen et al., 2022; Della Libera & Chelazzi, 2009; Fandakova & Gruber, 2021; Sims, 2016). The rational inattention framework therefore suggests that a neurocognitive deficit with only a relatively subtle direct influence on the endogenous precision of sensory experience may nevertheless have a substantial effect on behavioural engagement preferences over time, with subsequent cascading effects that dramatically shape long-term learning outcomes with respect to specific information types. The neurodivergent child's environment will contain a wide array of information sources associated with different degrees of endogenous precision, and an expectation that a certain information source within that environment is not immediately conducive to information gain and goal fulfilment relative to others may quickly drive disengagement with that information source and a re-distribution of engagement preferences. A great deal of a neurodivergent child's day-to-day behaviour may be shaped by rational inattention, as the child adopts perhaps unconventional ways of engaging with their environment and their caregivers and peers become naturally responsive to those behaviours. Yet, rational inattention precludes direct practice with endogenously imprecise information sources, which are in principle learnable to some degree, and is, therefore, expected to foster deeper delays in key skill development. The central message of this paper, then, is that, perhaps paradoxically, optimal information seeking may contribute negatively to an already delayed learning trajectory in neurodivergent children.

The extent to which rational inattention may explain neurodivergent behaviour over and above an assumed primary neurocognitive deficit, such as a deficit in auditory processing or working memory, remains unclear. This is the essential question raised by the current report: *How much of the delayed learning trajectory and behaviour of a neurodivergent child is the direct result of a fundamental neurocognitive deficit, and how much may be attributed to optimal information seeking – that is, to rational inattention – given fundamental, and in certain cases feasibly very* *subtle, neurocognitive constraints?* For instance, we wonder how much of the behavioural data readily attributed to a functionally discrete attentional capacity bottleneck (e.g. difficulties attending to speech in DLD) may instead be attributed to the child's engagement being regulated by the fact that past experience has shown that a particular information source is not immediately conducive to effective information gain.

Future research should aim to determine how well the evidence of attentional boost effects reported in this manuscript (e.g. data from Leclercq et al., 2015) replicates across different information sources (including speech, text, numeric stimuli and social cues) and among children with different types of learning difficulty. Future research should also aim to devise paradigms that help to separate out disengagement profiles that are (i) affective-motivational (i.e. in line with the Matthew effect literature) and (ii) epistemic (i.e. in line with the rational inattention framework developed in the current manuscript). Studying how engagement preferences are shaped by continuous degrees of underlying difficulty rather than recruiting participants based on the presence or absence of a clinical diagnosis might be particularly insightful.

Determining the factors that contribute to online disengagement profiles associated with long-term learning delays may inform the development of effective programs of clinical intervention that improve outcomes for neurodivergent children. The simulations presented in the current report suggest that existing teaching and intervention methods that promote deep engagement and aim explicitly to boost the quality of the child's long-term awareness of specific features of a given information source confer gains in part by tempering the child's disengagement preference for that information source and promoting self-driven information seeking and learning. In contrast, the current report suggests that widely used attention training programmes (e.g. *Cogmed*; see Shipstead et al., 2012) may in certain cases be ineffective because they wrongly target the secondary effects of primary neurocognitive deficits. That is, attention training programmes may in certain cases be targeting emergent (rather than causally implicated) disengagement behaviours that are optimal in the face of primary neurobiological constraints. The rational inattention framework therefore raises further questions about the use of attention training in a clinical context (see also Shipstead et al., 2012; Jones & Westermann, 2022).

5 | CONCLUSION

This paper has presented rational inattention as a new, transdiagnostic theory of information seeking in children with neurodevelopmental conditions including DLD, dyslexia, dyscalculia and autism. Our argument is not that every unconventional engagement profile can be explained exclusively in terms of rational inattention. The picture is complex, and attentional capacity limitations, deficits in long-term learning, and a child's affective motivation to engage with an information source may all play an important role in explaining why neurodivergent children characteristically disengage (or engage unconventionally) with information sources about which they are uncertain. Rational inattention is just one piece of this puzzle. But it may well be integral to understanding neurodivergent behaviour and learning.

ACKNOWLEDGEMENTS

The authors would like to thank Professor Karla McGregor for providing feedback on an earlier draft of this manuscript. Gert Westermann was supported by the Economic and Social Research Council (ESRC) International Centre for Language and Communicative Development (LuCiD) [ES/S007113/1 and ES/L008955/1].

CONFLICT OF INTEREST STATEMENT

The authors declare no conflicts of interest.

DATA AVAILABILITY STATEMENT

All materials required to re-create the simulations and analyses presented in this manuscript are available from the following public repository: https://osf.io/jey7q/.

ORCID

Samuel David Jones https://orcid.org/0000-0002-8870-3223
Manon Wyn Jones https://orcid.org/0000-0003-3266-5472
Kami Koldewyn https://orcid.org/0000-0003-3588-6449
Gert Westermann https://orcid.org/0000-0003-2803-1872

ENDNOTES

[1] We fully appreciate that this overlooks huge complexity in the aetiology and phenotypes of different neurodevelopmental disorders. However, we believe this 'bird's eye view' approach with a focus on broad canonical areas of difficulty to be justified in the context of developing a generalised, transdiagnostic account of neurodivergent information seeking. This is a point we develop further below. Ultimately, the argument we work towards steps away not only from the underlying nature of neurocognitive differences and diagnostic labels (e.g. DLD or dyslexia) but also from the broader *neurotypical* and *neurodivergent* distinction with respect to the principles of optimal information seeking.

[2] Our own view is that the term 'strategy' may be too *rich* in the sense of Haith (1998). Elaborate strategizing is not required. Once again, the child simply *takes the world as it comes* and interacts with an information source to the degree that they expect they learn from it. This is a position reflected in the computational simulations that follow, which illustrate engagement with an information source as a function of expected learning but do not evidence strategizing or meta-reasoning.

[3] Note that our autoencoder incorporated convolutional hidden layers, rather than vanilla feed-forward hidden layers. This enables the network to learn and later combine simple features like the edges of shapes. There is a strong precedent for incorporating convolution of this sort into computational models of auditory and visual processing (e.g. Chung & Abbott, 2021; Cohen et al., 2020; DiCarlo & Cox, 2007; Francl & McDermott, 2022).

REFERENCES

Addyman, C., & Mareschal, D. (2013). Local redundancy governs infants' spontaneous orienting to visual-temporal sequences. *Child Development*, *84*(4), 1137–1144. https://doi.org/10.1111/cdev.12060

Annaz, D., Karmiloff-Smith, A., Johnson, M. H., & Thomas, M. S. C. (2009). A cross-syndrome study of the development of holistic face recognition in children with autism, Down syndrome, and Williams syndrome

. *Journal of Experimental Child Psychology*, *102*(4), 456–486. https://doi.org/10.1016/j.jecp.2008.11.005

Anobile, G., Stievano, P., & Burr, D. C. (2013). Visual sustained attention and numerosity sensitivity correlate with math achievement in children. *Journal of Experimental Child Psychology*, *116*(2), 380–391. https://doi.org/10.1016/j.jecp.2013.06.006

Ashkenazi, S., Rubinsten, O., & Henik, A. (2009). Attention, automaticity, and developmental dyscalculia. *Neuropsychology*, *23*(4), 535–540. https://doi.org/10.1037/a0015347

Askenazi, S., & Henik, A. (2010). Attentional networks in developmental dyscalculia. *Behavioral and Brain Functions*, *6*(1), 2. https://doi.org/10.1186/1744-9081-6-2

Astle, D. E., Holmes, J., Kievit, R., & Gathercole, S. E. (2022). Annual Research Review: The transdiagnostic revolution in neurodevelopmental disorders. *Journal of Child Psychology and Psychiatry*, *63*(4), 397–417. Article 4. https://doi.org/10.1111/jcpp.13481

Baer, C., Gill, I. K., & Odic, D. (2018). A domain-general sense of confidence in children. *Open Mind*, *2*(2), 86–96. https://doi.org/10.1162/opmi_a_00020

Barlow, H. B. (2012). Possible principles underlying the transformations of sensory messages. In W. A. Rosenblith (Ed.), *Sensory communication* (pp. 216–234). The MIT Press. https://doi.org/10.7551/mitpress/9780262518420.003.0013

Baron-Cohen, S. (2009). Autism: The empathizing-systemizing (E-S) theory. *Annals of the New York Academy of Sciences*, *1156*(1), 68–80. https://doi.org/10.1111/j.1749-6632.2009.04467.x

Bates, C. J., Lerch, R. A., Sims, C. R., & Jacobs, R. A. (2019). Adaptive allocation of human visual working memory capacity during statistical and categorical learning. *Journal of Vision*, *19*(2), 11. https://doi.org/10.1167/19.2.11

Bazhydai, M., Westermann, G., & Parise, E. (2020). "I don't know but I know who to ask": 12-month-olds actively seek information from knowledgeable adults. *Developmental Science*, *23*(5). https://doi.org/10.1111/desc.12938

Bishop, D. V. M., & McArthur, G. M. (2005a). Individual differences in auditory processing in specific language impairment: A follow-up study using event-related potentials and behavioural thresholds. *Cortex*, *41*(3), 327–341. Article 3. https://doi.org/10.1016/S0010-9452(08)70270-3

Bishop, D. V. M., & McArthur, G. M. (2005b). Individual differences in auditory processing in specific language impairment: A follow-up study using event-related potentials and behavioural thresholds. *Cortex*, *41*(3), 327–341. https://doi.org/10.1016/S0010-9452(08)70270-3

Bourgeois, A., Chelazzi, L., & Vuilleumier, P. (2016). How motivation and reward learning modulate selective attention. In *Progress in brain research* (Vol. 229, pp. 325–342). Elsevier. https://doi.org/10.1016/bs.pbr.2016.06.004

Chen, X., Twomey, K. E., & Westermann, G. (2022). Curiosity enhances incidental object encoding in 8-month-old infants. *Journal of Experimental Child Psychology*, *223*, 105508. https://doi.org/10.1016/j.jecp.2022.105508

Chung, S., & Abbott, L. F. (2021). Neural population geometry: An approach for understanding biological and artificial neural networks. *Current Opinion in Neurobiology*, *70*, 137–144. https://doi.org/10.1016/j.conb.2021.10.010

Cittern, D., Nolte, T., Friston, K., & Edalat, A. (2018). Intrinsic and extrinsic motivators of attachment under active inference. *PLOS ONE*, *13*(4), e0193955. https://doi.org/10.1371/journal.pone.0193955

Cohen, U., Chung, S. Y., Lee, D. D., & Sompolinsky, H. (2020). Separability and geometry of object manifolds in deep neural networks. *Nature Communications*, *11*(1), 746. Article 1. https://doi.org/10.1038/s41467-020-14578-5

Conti-Ramsden, G., Durkin, K., Toseeb, U., Botting, N., & Pickles, A. (2018). Education and employment outcomes of young adults with a history of developmental language disorder. *International Journal of Language & Communication Disorders*, *53*(2), 237–255. Article 2. https://doi.org/10.1111/1460-6984.12338

Conti-Ramsden, G., St Clair, M. C., Pickles, A., & Durkin, K. (2012). Developmental trajectories of verbal and nonverbal skills in individuals with a history of specific language impairment: From childhood to adolescence. *Journal of Speech, Language, and Hearing Research: JSLHR*, *55*(6), 1716–1735. https://doi.org/10.1044/1092-4388(2012/10-0182)

Della Libera, C., & Chelazzi, L. (2009). Learning to attend and to ignore is a matter of gains and losses. *Psychological Science*, *20*(6), 778–784. https://doi.org/10.1111/j.1467-9280.2009.02360.x

DiCarlo, J. J., & Cox, D. D. (2007). Untangling invariant object recognition. *Trends in Cognitive Sciences*, *11*(8), 333–341. Article 8. https://doi.org/10.1016/j.tics.2007.06.010

Ebert, K. D., & Kohnert, K. (2009). Non-linguistic cognitive treatment for primary language impairment. *Clinical Linguistics & Phonetics*, *23*(9), 647–664. https://doi.org/10.1080/02699200902998770

Ebert, K. D., & Kohnert, K. (2011). Sustained attention in children with primary language impairment: A meta-analysis. *Journal of Speech Language and Hearing Research*, *54*(5), 1372–1384. Article 5. https://doi.org/10.1044/1092-4388(2011/10-0231)

Fandakova, Y., & Gruber, M. J. (2021). States of curiosity and interest enhance memory differently in adolescents and in children. *Developmental Science*, *24*(1), e13005. https://doi.org/10.1111/desc.13005

Feldman, H., & Friston, K. J. (2010). Attention, uncertainty, and free-energy. *Frontiers in Human Neuroscience*, *4*, 215. https://doi.org/10.3389/fnhum.2010.00215

Foster, M. E. (2023). Evaluating the impact of supplemental computer-assisted math instruction in elementary school: A conceptual replication. *Journal of Research on Educational Effectiveness*, 1–25. https://doi.org/10.1080/19345747.2023.2174919

Francl, A., & McDermott, J. H. (2022). Deep neural network models of sound localization reveal how perception is adapted to real-world environments. *Nature Human Behaviour*, *6*(1), 111–133. https://doi.org/10.1038/s41562-021-01244-zbib>

Gabaix, X., & Laibson, D. (2017). *Myopia and discounting* [w23254]. National Bureau of Economic Research. https://doi.org/10.3386/w23254

Gathercole, S. E., & Baddeley, A. D. (1990). Phonological memory deficits in language disordered children: Is there a causal connection? *Journal of Memory and Language*, *29*(3), 336–360. Article 3. https://doi.org/10.1016/0749-596X(90)90004-J

Gershman, S. J., & Bhui, R. (2020). Rationally inattentive intertemporal choice. *Nature Communications*, *11*(1), 3365. https://doi.org/10.1038/s41467-020-16852-y

Gershman, S. J., & Burke, T. (2022). Mental control of uncertainty. *Cognitive, Affective, & Behavioral Neuroscience*, *23*(3), 465–475. https://doi.org/10.3758/s13415-022-01034-8

Gershman, S. J., Horvitz, E. J., & Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, *349*(6245), 273–278. https://doi.org/10.1126/science.aac6076

Gottlieb, J., Oudeyer, P.-Y., Lopes, M., & Baranes, A. (2013). Information-seeking, curiosity, and attention: Computational and neural mechanisms. *Trends in Cognitive Sciences*, *17*(11), 585–593. https://doi.org/10.1016/j.tics.2013.09.001

Gray, S., Fox, A. B., Green, S., Alt, M., Hogan, T. P., Petscher, Y., & Cowan, N. (2019). Working memory profiles of children with dyslexia, developmental language disorder, or both. *Journal of Speech, Language, and Hearing Research*, *62*(6), 1839–1858. Article 6. https://doi.org/10.1044/2019_JSLHR-L-18-0148

Haith, M. M. (1998). Who put the cog in infant cognition? Is rich interpretation too costly? *Infant Behavior and Development*, *21*(2), 167–179. https://doi.org/10.1016/S0163-6383(98)90001-7

Hancock, R., Pugh, K. R., & Hoeft, F. (2017). Neural noise hypothesis of developmental dyslexia. *Trends in Cognitive Sciences*, *21*(6), 434–448. https://doi.org/10.1016/j.tics.2017.03.008

Holmes, J., Butterfield, S., Cormack, F., Loenhoud, A. V., Ruggero, L., Kashikar, L., & Gathercole, S. (2015). Improving working memory in children with low language abilities. *Frontiers in Psychology*, *6*, 519. https://doi.org/10.3389/fpsyg.2015.00519

Joanisse, M. F., & Seidenberg, M. S. (2003). Phonology and syntax in specific language impairment: Evidence from a connectionist model. *Brain and Language*, *86*(1), 40–56. Article 1. https://doi.org/10.1016/S0093-934X(02)00533-3

Johnson, M. H., Charman, T., Pickles, A., & Jones, E. J. H. (2021). Annual research review: Anterior modifiers in the emergence of neurodevelopmental disorders (AMEND)—a systems neuroscience approach to common developmental disorders. *Journal of Child Psychology and Psychiatry*, *62*(5), 610–630. https://doi.org/10.1111/jcpp.13372

Jones, M. W., Kuipers, J.-R., Nugent, S., Miley, A., & Oppenheim, G. (2018). Episodic traces and statistical regularities: Paired associate learning in typical and dyslexic readers. *Cognition*, *177*, 214–225. https://doi.org/10.1016/j.cognition.2018.04.010

Jones, S. D., & Brandt, S. (2020). Density and distinctiveness in early word learning: Evidence from neural network simulations. *Cognitive Science*, *44*(1). https://doi.org/10.1111/cogs.12812

Jones, S. D., Stewart, H. J., & Westermann, G. (2023). A maturational frequency discrimination deficit may explain developmental language disorder. *Psychological Review*. https://doi.org/10.1037/rev0000436

Jones, S. D., & Westermann, G. (2022). Under-resourced or overloaded? Rethinking working memory and sentence comprehension deficits in developmental language disorder. *Psychological Review*, *129*(6), 1358–1372. Advance online publication. https://doi.org/10.1037/rev0000338

Kaldy, Z., Kraper, C., Carter, A. S., & Blaser, E. (2011). Toddlers with autism spectrum disorder are more successful at visual search than typically developing toddlers: Toddlers with ASD are more successful at visual search. *Developmental Science*, *14*(5), 980–988. https://doi.org/10.1111/j.1467-7687.2011.01053.x

Kidd, C., & Hayden, B. Y. (2015). The psychology and neuroscience of curiosity. *Neuron*, *88*(3), 449–460. https://doi.org/10.1016/j.neuron.2015.09.010

Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2012). The goldilocks effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PLoS ONE*, *7*(5), e36399. https://doi.org/10.1371/journal.pone.0036399

Leclercq, A.-L., Maillart, C., Lange, M., & Majerus, S. (2015). The impact of attentional allocation capacities on nonword repetition in children with specific language impairment. *Clinical Linguistics & Phonetics*, *29*(8–10), 719–735. https://doi.org/10.3109/02699206.2015.1022664

Lindsay, R. L., Tomazic, T., Levine, M. D., & Accardo, P. J. (1999). Impact of attentional dysfunction in dyscalculia. *Developmental Medicine & Child Neurology*, *41*(9), 639–642. https://doi.org/10.1017/S0012162299001309

Mareschal, D., French, R. M., & Quinn, P. C. (2000). A connectionist account of asymmetric category learning in early infancy. *Developmental Psychology*, *36*(5), 635–645. https://doi.org/10.1037/0012-1649.36.5.635

Mareschal, D., & Thomas, M. S. C. (2007). Computational modeling in developmental psychology. *IEEE Transactions on Evolutionary Computation*, *11*(2), 137–150. https://doi.org/10.1109/TEVC.2006.890232

Moriuchi, J. M., Klin, A., & Jones, W. (2017). Mechanisms of diminished attention to eyes in autism. *American Journal of Psychiatry*, *174*(1), 26–35. https://doi.org/10.1176/appi.ajp.2016.15091222

Muter, V., & Snowling, M. J. (2009). Children at familial risk of dyslexia: Practical implications from an at-risk study. *Child and Adolescent Mental Health*, *14*(1), 37–41. https://doi.org/10.1111/j.1475-3588.2007.00480.x

Oudeyer, P.-Y., Kaplan, F., & Hafner, V. V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation*, *11*(2), 265–286. https://doi.org/10.1109/TEVC.2006.890271

Parr, T., & Friston, K. J. (2017). Uncertainty, epistemics and active inference. *Journal of The Royal Society Interface*, *14*(136), 20170376. https://doi.org/10.1098/rsif.2017.0376

Peters, A., McEwen, B. S., & Friston, K. (2017). Uncertainty and stress: Why it causes diseases and how it is mastered by the brain. *Progress in Neurobiology*, *156*, 164–188. https://doi.org/10.1016/j.pneurobio.2017.05.004

Pfost, M., Hattie, J., Dörfler, T., & Artelt, C. (2014). Individual differences in reading development: A review of 25 years of empirical research on matthew effects in reading. *Review of Educational Research*, *84*(2), 203–244. https://doi.org/10.3102/0034654313509492

Poli, F., Serino, G., Mars, R. B., & Hunnius, S. (2020). Infants tailor their attention to maximize learning. *Science Advances*, *6*(39), eabb5053. https://doi.org/10.1126/sciadv.abb5053

Rabiner, D., & Coie, J. D. (2000). early attention problems and children's reading achievement: A longitudinal investigation. *Journal of the American Academy of Child & Adolescent Psychiatry*, *39*(7), 859–867. https://doi.org/10.1097/00004583-200007000-00014

Rinaldi, S., Caselli, M. C., Cofelice, V., D'Amico, S., De Cagno, A. G., Della Corte, G., Di Martino, M. V., Di Costanzo, B., Levorato, M. C., Penge, R., Rossetto, T., Sansavini, A., Vecchi, S., & Zoccolotti, P. (2021). Efficacy of the treatment of developmental language disorder: A systematic review. *Brain Sciences*, *11*(3), 407. Article 3. https://doi.org/10.3390/brainsci11030407

Shipstead, Z., Hicks, K. L., & Engle, R. W. (2012). Cogmed working memory training: Does the evidence support the claims? *Journal of Applied Research in Memory and Cognition*, *1*(3), 185–193. https://doi.org/10.1016/j.jarmac.2012.06.003

Sims, C. R. (2016). Rate–distortion theory and human perception. *Cognition*, *152*, 181–198. https://doi.org/10.1016/j.cognition.2016.03.020

Smolak, E., McGregor, K. K., Arbisi-Kelm, T., & Eden, N. (2020). Sustained attention in developmental language disorder and its relation to working memory and language. *Journal of Speech, Language, and Hearing Research*, *63*(12), 4096–4108. https://doi.org/10.1044/2020_JSLHR-20-00265

Soriano-Ferrer, M., & Morte-Soriano, M. (2017). Teacher perceptions of reading motivation in children with developmental dyslexia and average readers. *Procedia—Social and Behavioral Sciences*, *237*, 50–56. https://doi.org/10.1016/j.sbspro.2017.02.012

Stanovich, K. E. (2009). Matthew effects in reading: Some consequences of individual differences in the acquisition of literacy. *Journal of Education*, *189*(1–2), 23–55. https://doi.org/10.1177/0022057409189001-204

Thomas, M. S. C., Fedor, A., Davis, R., Yang, J., Alireza, H., Charman, T., Masterson, J., & Best, W. (2019). Computational modeling of interventions for developmental disorders. *Psychological Review*, *126*(5), 693–726. https://doi.org/10.1037/rev0000151

Thomas, M. S. C., Knowland, V. C. P., & Karmiloff-Smith, A. (2011). Mechanisms of developmental regression in autism and the broader phenotype: A neural network modeling approach. *Psychological Review*, *118*(4), 637–654. https://doi.org/10.1037/a0025234

Twomey, K. E., & Westermann, G. (2018). Curiosity-based learning in infants: A neurocomputational approach. *Developmental Science*, e12629. Article September. https://doi.org/10.1111/desc.12629

West, G., Shanks, D. R., & Hulme, C. (2021). Sustained attention, not procedural learning, is a predictor of reading, language and arithmetic skills in children. *Scientific Studies of Reading*, *25*(1), 47–63. https://doi.org/10.1080/10888438.2020.1750618

Westermann, G., & Mareschal, D. (2004). From parts to wholes: Mechanisms of development in infant visual object processing. *Infancy*, *5*(2), 131–151. https://doi.org/10.1207/s15327078in0502_2

Westermann, G., & Mareschal, D. (2012). Mechanisms of developmental change in infant categorization. *Cognitive Development*, *27*(4), 367–382. https://doi.org/10.1016/j.cogdev.2012.08.004

Westermann, G., & Ruh, N. (2012). A neuroconstructivist model of past tense development and processing. *Psychological Review*, *119*(3). https://doi.org/10.1037/a0028258

Westermann, G., Ruh, N., & Plunkett, K. (2009). Connectionist approaches to language learning. *Linguistics*, *47*(2), 413–452. https://doi.org/10.1515/LING.2009.015

How to cite this article: Jones, S. D., Jones, M., Koldewyn, K., & Westermann, G. (2024). Rational inattention: A new theory of neurodivergent information seeking. *Developmental Science*, e13492. https://doi.org/10.1111/desc.13492