Computational rationality and developmental neurodivergence

Jones, Samuel ; Rauwolf, Paul; Westermann, Gert

Trends in Cognitive Sciences

DOI:
10.1016/j.tics.2025.01.006

Published: 08/02/2025

Dyfyniad o'r fersiwn a gyhoeddwyd / Citation for published version (APA):
Jones, S., Rauwolf, P., & Westermann, G. (2025). Computational rationality and developmental neurodivergence. *Trends in Cognitive Sciences*. https://doi.org/10.1016/j.tics.2025.01.006

Trends in
Cognitive Sciences

CellPress
OPEN ACCESS

Forum

Computational rationality and developmental neurodivergence

Samuel David Jones[1,*],
Paul Rauwolf[1], and
Gert Westermann[2]

The role of behaviour – choices, actions, and habits – in shaping neurodivergent development remains unclear. In this forum article we introduce computational rationality as a framework for understanding dynamic feedback between brain and behavioural development, and neurodevelopmental variation.

The search for the neurocognitive bases of conditions like **dyslexia** (see Glossary), **dyscalculia**, and **developmental language disorder (DLD)** is a central focus in developmental science. Despite the lessons of the transdiagnostic revolution, which highlights the complexity inherent in neurodevelopmental conditions and the limitations of **core-deficit hypotheses**, this literature remains fundamentally divided between causal accounts centred, for instance, on either auditory or visual perception, working memory, or statistical learning, each associated with a candidate neural substrate [1]. In contrast, behavioural contributions to neurodevelopmental differences remain understudied. There is, of course, acknowledgement that phenotypic variation is the product of **probabilistic epigenesis**, that is, the dynamic interaction between genetics, neural activity, behaviour, and the environment [2]. However, how a child's behaviour – their choices, actions, and habits – shapes neurodivergent development remains hard to define.

One way to understand how behaviour is not only influenced by but also influences neurodivergent child development is through **computational rationality**, which assumes that behaviours are optimised for achieving the highest expected utility subject to neurocognitive resource constraints [3]. Computational rationality inherits from a long tradition in **decision theory** that incorporates constraints to explain deviations from axiomatic rational behaviour (e.g., **bounded rationality**). It is this core theoretical focus on what best to do when faced with constraints, combined with a novel focus on neurocognitive information processing, that makes the computational rationality paradigm well suited to determining behavioural contributions to neurodevelopmental variation.

The rational analysis of neurodivergent child behaviour

The description of neurodivergent child behaviour as 'rational' might appear counterintuitive. While neurotypical children tend to engage with stimuli about which they are uncertain, seemingly to maximise learning and reward, neurodivergent children often disengage from stimuli about which they are characteristically uncertain, or engage with them unconventionally. For a child with dyslexia, this might mean relying on whole-word recognition rather than letter-by-letter phonological decoding when reading [4]. For a child with dyscalculia, it might mean relying on counting rather than 'subitising', including using visual aids like their fingers, to determine the number of items (e.g., dots) in a set [5]. And for a child with DLD, it might mean relying on situational cues such as peer behaviour in order to decode spoken instructions, for instance those from a teacher [6]. These heuristics, which are sometimes termed **compensatory strategies**, may appear suboptimal because they do not always generalise well, perhaps leading to worse outcomes. Sight reading, for instance, may not support the accurate pronunciation of a novel word, and strategies used in a familiar environment (e.g., in parent–child interactions at home) might not be as effective elsewhere.

Glossary

Bounded rationality: a decision-making framework emphasising that agents rely on heuristics and satisficing to effectively navigate cognitive, temporal, and informational constraints.

Compensatory strategies: adaptive techniques and heuristics used by individuals with neurodevelopmental difficulties to work around specific challenges to achieve a goal.

Computational rationality: a framework in which behaviour is understood as the outcome of decision-making optimised to maximise expected utility under constraints in a given environment. The word 'computational' highlights a novel focus on biological and artificial neural processing.

Core-deficit hypothesis: the assumption that symptoms of a developmental condition arise from a single, discrete cognitive or neurological cause.

Decision theory: classically models rational decision-making under uncertainty using expected utility and probability.

Developmental language disorder (DLD): a neurodevelopmental condition affecting spoken language acquisition and use.

Developmental niche construction: a framework proposing that organisms actively modify their environments in ways that shape their development.

Dyscalculia: a neurodevelopmental condition affecting the ability to understand and use numbers and arithmetic.

Dyslexia: a neurodevelopmental condition characterised by reading difficulties, typically involving phonological processing.

Learning by thinking: the use of mental simulation, synthesis, and reasoning to solve problems or develop knowledge in the absence of direct input.

Perceptual narrowing: developmental process in which the ability to perceive stimuli becomes more specialised, reducing sensitivity to less frequently encountered information.

Probabilistic epigenesis: the idea that development results from the dynamic interaction of genetic, neural, behavioural, and environmental factors.

The computational rationality paradigm nevertheless interprets such behaviours not – as is common – in terms of 'deficiency' or 'demotivation', but instead as adaptive efforts to maximise utility given the individual's neurocognitive makeup and the environment in which they find themselves. The claim here is that optimal decision-making about which information sources to attend to and which action policies to pursue occurs in the context of a limited-capacity attentional system and perceptual experience that is imprecise due to both exogenous noise and endogenous neurocognitive noise on a continuum from typical to severe [7,8] (Figure 1A,B). When the expected cost of information processing is high, an implicit cost–benefit analysis may bias the child towards inferences and the selection of action policies with high prior probability, and likewise towards heuristics that the child associates with relatively low information processing cost given their experience (Box 1 and Figure 1C–F). Disengagement or unconventional engagement with text in dyslexia, numeric stimuli in dyscalculia, and speech in DLD may be understood as the outcomes of an implicit resource-rational trade-off of this kind: a trade-off that becomes increasingly habitual during early development.

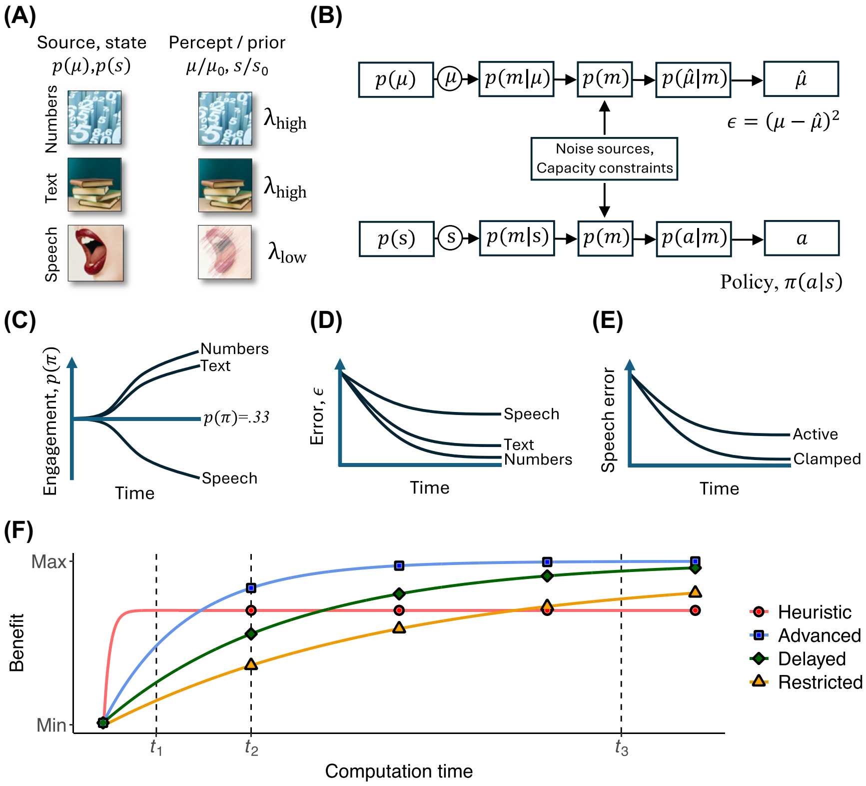

Trends in Cognitive Sciences

Figure 1. Computational rationality and neurodevelopmental variation. (A) Numeric, text, and speech information sources and states (μ, s), with low precision (λ) indicating neurocognitive constraints. (B) Abstract encoder–decoder communication channels for perception (top) and action (bottom). $p(m)$ indicates encoding, ϵ is error, a is action, and π is policy [7,8]. (C) Engagement with numbers and text (high precision) and speech (low precision) over time [11]. Engagement is initially symmetrical across information sources, but engagement with speech declines over time due to low precision limiting learning and reward. (D) Engagement-related error rates over time. Low engagement with speech is associated with a high error rate for this information source. (E) Error rates for two agents with the precision deficit illustrated in (A): the 'active' agent engages adaptively with numeric, text, and speech information sources as per computational rationality; the 'clamped' agent is programmed to engage symmetrically with all three information sources (i.e., this agent cannot disengage from speech). Clamping results in better learning for speech stimuli, illustrating that resource rational behaviour (or 'rational inattention') can deepen learning delays over time [11]. (F) Resource-rational trade-off between heuristic and direct computation strategies in advanced, delayed, and restricted agents. Direct computation is most effective in advanced agents, mimicking a neurotypical profile. Direct computation progresses more slowly in delayed agents, and asymptotes early in restricted agents, mimicking plausible neurodivergent profiles. Direct computation by each agent may be compared to the fast-and-frugal heuristic strategy. At t_1, the heuristic strategy is universally optimal due to insufficient time for direct computation (i.e., inference refinement and complex action policy planning). At t_2, the heuristic remains optimal for the delayed and restricted agents, but direct computation is optimal for the advanced agent. By t_3, all agents benefit more from direct computation than from the heuristic strategy, though this gain is relatively small for the restricted agent.

Computational rationality may explain hallmark neurodivergent behaviours, including disengagement and defaulting to common visual or situational cues or frequent structures (e.g., spellings, words, or syntax), and similarly to high-probability action policies and heuristics when reading, using numbers, or listening to or producing speech [4–6]. Adaptive disengagement should also be considered in the context of **learning by thinking**, which plays a crucial role in early cognitive development [9]. That is, high expected information processing cost may reduce the likelihood of the child experimenting with a given class of stimuli (e.g., numbers or language) through mental analogy and simulation in the absence of direct input, providing an additional obstacle to developing proficiency. Importantly, computational rationality is indifferent to diagnostic labels and to the broader neurotypical and neurodivergent distinction; the neurodivergent child is doing exactly what any rational agent would do: optimising their finite resources to maximise expected utility within a limited time horizon [7].

Adaptive disengagement as developmental niche construction

Collectively, adaptive disengagement behaviours attributable to neurocognitive constraints reflect a form of **developmental niche construction** that regulates pressures on the child because it is shaped to

Box 1. A normative Bayesian model of computational rationality

Agents infer a parameter such as the identity of a spoken or written word, μ, from an exemplar, x, where the posterior inference, $P(\mu \mid x)$, depends on perceptual experience, $P(x \mid \mu)$, and prior experience, $P(\mu)$:

$$P(\mu \mid x) \propto P(x \mid \mu) P(\mu) \qquad [I]$$

Learning is driven by the relative precision (inverse variance) of the perceptual experience, λ, and prior, λ_0. High perceptual precision supports effective learning (updating μ_0 to $\hat{\mu}$), while low precision leads agents to default to their priors:

$$\hat{\mu} = \mu_0 + \frac{\lambda}{\lambda + \lambda_0}(x - \mu_0) \qquad [II]$$

Reward, U, is inversely proportional to prediction error, $\epsilon = (\mu - \hat{\mu})^2$, and dependent, therefore, on perceptual precision. Agents can increase perceptual precision, λ, by increasing attention. However, the critical feature of computational rationality is that attention is bounded, as expressed by:

$$C = B \log_2\left(1 + \frac{S}{N}\right) \qquad [III]$$

where C is capacity, B is bandwidth, and $\frac{S}{N}$ is signal-to-noise ratio. Mutual information, $I(\mu; x)$, formalises how attention reduces uncertainty. The optimisation problem agents face balances reward procurement with attentional cost, κ:

$$\lambda^* = \arg\max_{\lambda} \quad U - \kappa I(\mu; x) \qquad [IV]$$

With high exogenous or endogenous noise, attentional disengagement and reliance on priors may be optimal. This formalism can be extended to policy selection with similar conclusions: noisy state knowledge results in the avoidance of action policies with low prior probability [7,8].

their abilities, needs, and preferences [10]. A consequence of this is that although disengagement behaviours may be optimal within a specific setting and short time horizon, they may not promote effective and generalisable long-term learning, and so may reinforce learning differences over time. Active disengagement or unconventional engagement with text, numeric stimuli, or speech, for instance, may contribute to the reinforcement of learning delays in dyslexia, dyscalculia, and DLD by precluding regular exposure to and practice with relevant stimuli (Figure 1C–E).

Formalising this idea, we recently showed that an active agent-based model with a precision deficit – a proxy for primary neurocognitive constraints, the nature of which was bracketed out – adaptively disengaged from subjectively noisy stimuli [11] (Figure 1C,D). This resulted in worse learning of stimuli affected by the precision deficit over time compared with a control model which had the same perceptual precision deficit, but which was programmed to engage equally with all of the information sources in its environment (Figure 1E). The capacity for variable engagement to shape a learning trajectory in this way has been described in terms of a Matthew effect [4] (because 'the rich get richer', and vice versa), and our treatment here in terms of computational rationality lends traction to this idea and highlights its transdiagnostic importance (Matthew effects have commonly been studied in dyslexia). However, although complementary, these frameworks are somewhat different. Literature citing the Matthew effect often centres on affective disengagement due to repeated failures to learn, in contrast to the idea developed here that adaptive disengagement may be an optimal policy.

There is an important link here with the **perceptual narrowing** literature, which indicates that infants gradually lose sensitivity to perceptual distinctions outside of their experience (e.g., to non-frequent language sounds) [12]. Our account argues that an analogous effect is seen because of the developmental niche shaped by optimal moment-to-moment decision-making under neurocognitive constraints. The parallel is that information outside of the child's niche – defined in terms of modes of passive learning, action policy selection, inter-personal experiences, and mental simulation – is subject to attenuated encoding in memory, itself explaining learning delay. This feedback cycle can be inferred from the formalism presented in Box 1 (see also Figure 1), where perceptual imprecision or processing constraints bias the rational agent to make inferences and select action policies with high prior probability, inhibiting exploration and learning [7,8]. Considering non-linear dynamics, saddle points, and the notion of sensitivity to initial conditions, a cycle like this may in principle be set in motion by relatively small perturbations in precision and capacity, in contrast to the gross, discrete deficits commonly assumed under dominant core-deficit hypotheses. This includes very subtle neurological variation attributable to a constellation of genetic and environmental risk factors and in itself difficult to reliably detect through neuroimaging and neurophysiological assessment. Resource rational decision-making may be an essential behavioural mechanism linking different forms and severities of neurological variability to common neurodevelopmental phenotypes.

Concluding remarks

In its search for core neurocognitive deficits, developmental science has overlooked the potential for adaptive disengagement behaviours to amplify individual differences and play a formative, transdiagnostic role in conditions including but not limited to dyslexia, dyscalculia, and DLD. Computational rationality builds on established frameworks examining decision-making under constraints and points to formal mathematical and computational tools that can help to determine how a child's behaviour – their choices, actions, and habits – shapes neurodevelopmental variation. In contrast to dominant core-deficit approaches, these formalisms are characteristically multivariate: they view behaviour and learning as the product of dynamic interactions between factors including perceptual integrity, processing bandwidth, policy selection, and developing long-term knowledge. This perspective enriches our understanding of probabilistic epigenesis and our capacity to respond to individual differences effectively when required. Future research should pursue the application of computational rationality to neurodevelopmental variation, validating existing formalisms developed to explain adult behaviour against child data.

Acknowledgments

S.J. is supported by a Royal Society Research Grant (RG\R1\241234). G.W. is supported by Economic and Social Research Council (ESRC) International Centre for Language and Communicative Development (LuCiD) (ES/S007113/1 and ES/L008955/1).

Declaration of interests

The authors have no conflicts of interest to disclose.

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

*Correspondence:
samuel.jones@bangor.ac.uk (S.D. Jones).

https://doi.org/10.1016/j.tics.2025.01.006

References

1. Astle, D.E. *et al.* (2022) Annual Research Review: the transdiagnostic revolution in neurodevelopmental disorders. *Child Psychol. Psychiatry* 63, 397–417
2. Astle, D.E. *et al.* (2023) Toward computational neuroconstructivism: a framework for developmental systems neuroscience. *Trends Cogn. Sci.* 27, 726–744
3. Gershman, S.J. *et al.* (2015) Computational rationality: a converging paradigm for intelligence in brains, minds, and machines. *Science* 349, 273–278
4. Snowling, M.J. *et al.* (2020) Defining and understanding dyslexia: past, present and future. *Oxf. Rev. Educ.* 46, 501–513
5. Mammarella, I.C. *et al.* (2021) No evidence for a core deficit in developmental dyscalculia or mathematical learning disabilities. *Child Psychol. Psychiatry* 62, 704–714
6. Ullman, M.T. *et al.* (2024) The neuroanatomy of developmental language disorder: a systematic review and meta-analysis. *Nat. Hum. Behav.* 8, 962–975
7. Gershman, S.J. and Burke, T. (2023) Mental control of uncertainty. *Cogn. Affect. Behav. Neurosci.* 23, 465–475
8. Lai, L. and Gershman, S.J. (2021) Policy compression: an information bottleneck in action selection. *Psychol. Learn. Motiv.* 74, 195–232
9. Lombrozo, T. (2024) Learning by thinking in natural and artificial minds. *Trends Cogn. Sci.* 28, 1011–1022
10. Constant, A. *et al.* (2018) A variational approach to niche construction. *J. R. Soc. Interface* 15, 20170685
11. Jones, S.D. *et al.* (2024) Rational inattention: a new theory of neurodivergent information seeking. *Dev. Sci.* 27, e13492
12. Krasotkina, A. *et al.* (2021) Perceptual narrowing in face- and speech-perception domains in infancy: a longitudinal approach. *Infant Behav. Dev.* 64, 101607