Developmental Review 75 (2025) 101182

Contents lists available at ScienceDirect

Developmental Review

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

Integrating attention, working memory, and word learning in a dynamic field theory of executive function development: Moving beyond the 'component' view of executive function

John P. Spencer [a,*], Aaron T. Buss [b], Alexis R. McCraw [b], Eleanor Johns [a], Larissa K. Samuelson [a]

[a] University of East Anglia, UK
[b] University of Tennessee, Knoxville, USA

ARTICLE INFO

Keywords:
Executive function
Working memory
Attention
Cognitive development
Neural networks
Dynamic field theory

ABSTRACT

Executive functions (EFs) are core cognitive abilities that enable self-control and flexibility. EFs undergo transformational changes between 3 to 5 years of age; critically, individual differences in these abilities are predictive of longer-term outcomes. Thus, a key question is how EFs change in early development. This question is complicated by evidence that EFs are supported by attentional, inhibitory, working memory, and task switching processes, 'component' abilities which themselves change over time. Thus, understanding the early development of EFs requires a framework for understanding how attention, working memory, and other abilities develop and how they are integrated to enable new EF skills. Here, we take a theory-based approach to this problem, building a neural process model that integrates multiple neurocognitive processes together and grounds these processes in perception–action dynamics. We then explore how EFs emerge from these integrated processes over development. In particular, we extend prior work showing how the concepts of dynamic field theory explain the emergence of EFs in the dimensional change card sort (DCCS) task by integrating our theory of EF with a new model of visual exploration and word learning (WOLVES). This integration (WOLVES 2.0) specifies how visual-spatial attention, visual working memory, auditory-visual word representations, and top-down attention mechanisms come together to enable EFs from 3 to 5 years. Our central hypothesis is that children learn autonomous self-control by using language to guide attention to key features of the world in context. We demonstrate this, showing how, for example, children's learning of individual colour words and the associations among colour words and the word 'colour' gradually enable dimensional attention. More generally, we use WOLVES 2.0 as a concrete framework to explore how the concept of executive functions can be moved beyond the 'component' view towards a developmental systems perspective.

Introduction

Executive functions (EFs) are core cognitive abilities that enable self-control and flexibility. These skills enable children to mentally play with ideas, to give a considered rather than an impulsive response, and to stay focused (Diamond & Lee, 2011). EFs are a set of

* Corresponding author at: School of Psychology, 0.09 Lawrence Stenhouse Building, University of East Anglia, Norwich NR4 7TJ, UK.
E-mail address: j.spencer@uea.ac.uk (J.P. Spencer).

https://doi.org/10.1016/j.dr.2024.101182
Received 27 June 2024; Received in revised form 17 November 2024;
Available online 20 December 2024
0273-2297/

'higher-level' cognitive abilities that sit on top of and coordinate 'lower-level' skills that include working memory, inhibition, and task-switching (Miyake et al., 2000). Conceptually, EFs are like the conductor of an orchestra – when everything works, you get music; when it doesn't, you get cacophony.

EFs undergo transformational changes around 3 to 5 years of age. For instance, once language comprehension develops and verbal instruction is possible, children can be instructed to switch between tasks (Brooks, Hanauer, Padowska, & Rosman, 2003; Perner & Lang, 2002), inhibit actions (Steelandt, Thierry, Broihanne, & Dufour, 2012), and maintain and use information held in working memory (Cheng, Kaldy, & Blaser, 2020; Simmering, 2012). One key task commonly used to probe EF in the early developmental period is the dimensional change card sort (DCCS) task (Frye, Zelazo, & Palfai, 1995). This task is commonly used from 3 years of age and upward. It involves sorting bivalent test cards according to one dimension during a pre-switch phase (e.g., colour) and then sorting cards according to the second dimension during a post-switch phase (e.g., shape). For instance, in Fig. 1, children are asked to sort test cards to trays that contain target cards showing which features go where in the task. The cards that children sort match both target cards along different dimensions. If instructed to sort by colour during the pre-switch phase, the blue star would be sorted to the left and the red circle would be sorted to the right. During the post-switch phase, children would be instructed to sort by shape which requires making the opposite pattern of responses (see 'post' in Fig. 1).

Evidence shows rapid change between 3 and 5 years with gradual change thereafter (Zelazo et al., 2013). Most 3-year-olds will sort correctly during the pre-switch phase but will continue to sort the test cards by this first dimension during the post-switch phase. One explanation is that these children experience attentional inertia and are unable to inhibit thinking about the objects' initially relevant attribute (Kirkham, Cruess, & Diamond, 2003). By 5 years of age, however, children can robustly switch how they sort cards during the post-switch phase (Müller, Dick, Gela, Overton, & Zelazo, 2006; Zelazo, Muller, Frye, & Marcovitch, 2003). One of the reasons why this task might be so effective in measuring EF is that multiple EF skills are reflected in the task: children must hold higher-level rules in working memory, they must inhibit the prepotent response to sort by the original rule, and they must switch from one rule to the other as the context changes. Note, however, that not all researchers conceptualise this task in the same way. For instance, some researchers (e.g., Zelazo et al., 2003) frame the DCCS task as primarily a cognitive flexibility task.

Previously, the prevailing theory of development in the DCCS task was the cognitive complexity and control (CCC) theory of Zelazo and colleagues (2003) which posits that the imbalance caused by reactivating the same rule over and over must be overcome through higher-order rule representations. For example, by nesting rules such as, "*if* I am playing the colour game, and *if* it is blue, *then* sort here; but *if* I am playing the shape game, and *if* it is a circle, *then* sort there," children acquire the representational flexibility needed to switch rules. According to CCC theory, the complexity of rule representations is related to the level of consciousness that children can attain. By taking time to reflect on the nature of the problem, children achieve higher levels of representation in which greater complexity of rule-representations can be supported. Mechanistically, children progress to higher levels of complexity of rule representation as the frontal cortex matures (Bunge & Zelazo, 2006).

Evidence from use of the DCCS task and other probes of EF demonstrate that EF can be reliably measured over development. Critically, evidence also shows that EFs are a key developmental skill. In particular, EFs are more predictive of school readiness than IQ (Blair & Razza, 2007), and EFs predict math and reading competence in the school years (Ahmed, Tang, Waters, & Davis-Kean, 2019; Bull, Espy, & Wiebe, 2008; Cortés Pascual, Moyano Muñoz, & Quílez Robres, 2019; Gathercole, Pickering, Knight, & Stegmann, 2004; Lee et al., 2012; Robson, Allen, & Howard, 2020; Santana, Roazzi, & Nobre, 2022; Spiegel, Goodrich, Morris, Osborne, & Lonigan, 2021; St Clair-Thompson & Gathercole, 2006). EFs remain critical throughout life, predicting career and marriage satisfaction and positive mental and physical health (Dunn, 2010; Eakin et al., 2004; Prince et al., 2007). Reversely, children 3–11 years with poorer EFs develop into adults who have worse health, earn less, and commit more crimes 30 and 40 years later than those with better EFs (controlling for IQ, gender, and social class; see Moffitt et al., 2011; Richmond-Rakerd et al., 2021).

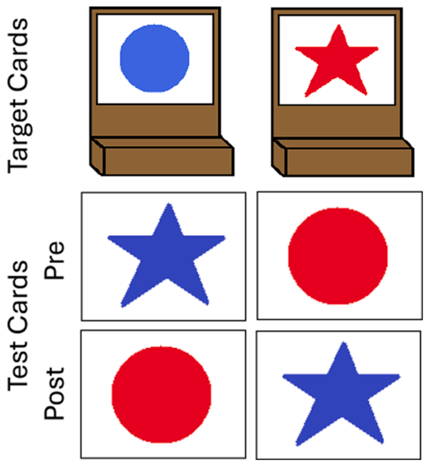

Fig. 1. Sample target and test cards used in the dimensional change card sort task. Test cards during the pre-switch phase are shown in the correct locations when playing the 'colour' game. Test cards during the post-switch phase are shown in the correct locations when playing the 'shape' game.

Given this long-term predictability, a key issue is how EFs change in early development. Simply put, if we could optimise EF early in development, for instance, ensuring that every child has good EF skills by the ages of 3–5 years, this could have a positive societal impact. This impact has been quantified by Heckman and colleagues (Heckman, 2006, 2011; Heckman, Savelyev, & Yavitz, 2010): The rate of return to society for investments in quality early childhood development for disadvantaged children is 12% through better outcomes in education, health, sociability, economic productivity and reduced crime. Thus, a key question is what do we know about the early development of EFs?

The literature on the early development of EF – so-called Early EF (EEF) – is complex. A central challenge here is that EFs have been largely conceptualised based on research with adults. Many definitions of the organisation of EF have been suggested, with over 32 separate definitions being presented within the literature (Goldstein & Naglieri, 2014). A common approach is to use confirmatory factor analysis (CFA) with a battery of tasks thought to tap different candidate skills, but also controlling for, for instance, differences in language ability that might influence performance (Friedman & Miyake, 2004; Wiebe, Espy, & Charak, 2008). Task performance scores are modelled as indicators of underlying latent variables creating a construct representing 'purified' EF (Wiebe et al., 2011). CFA of nine commonly used EF tasks with typical adults shows that measures of EF are structured around task-demands corresponding to working memory, shifting (cognitive flexibility), and inhibitory control. These are correlated but discrete in adulthood and can be measured separately (Miyake et al., 2000). Similar working memory, inhibition, and shifting three-factor models in adulthood have been reported in other studies as well (Friedman et al., 2008).

The challenge in early development is that these 'component' abilities themselves change over time. Thus, if we put this work in a developmental context, understanding EEF first requires understanding how working memory, inhibition, and switching skills develop over time on their own. Then we need to understand how they come together to enable complex EFs and how this integration works. And, of course, there's no guarantee that the structure revealed in studies with adults will be the same structure revealed in studies with young children. It is possible that some skills are important in early development and these skills become less important by adulthood. For instance, we might find that attentional skills are strongly related to EEF, but by the time EF consolidates into its adult state, attention no longer shows up in factor analytic studies with adults. We note that this raises one of the key questions of this special issue: how do we conceptualise EF and its relationship to a multi-component / multi-factor view?

To make this picture even more complicated, we don't always have great assessment tasks in early development, which highlights another key question of the special issue regarding the measurement of EF. A survey of the literature, for instance, reveals that there are at least 9 tasks commonly used across the literature to measure inhibitory control in early development. These include Go-No-Go (Simpson & Riggs, 2006), the Simon task (Davidson, Amso, Anderson, & Diamond, 2006), the Stroop task (Kochanska, Murray, & Harlan, 2000), the Day and Night task (Gerstadt, Hong, & Diamond, 1994), the Flanker task (Kerr-German, Tas, & Buss, 2023), the Piagetian A-not-B task (Thelen, Schöner, Scheier, & Smith, 2001), the anti-saccade task (Portugal, Bedford, Cheung, Mason, & Smith, 2021), the Gap-Overlap task (Hood & Atkinson, 1993), and the Early Childhood Inhibitory Touchscreen Task (ECITT; Holmboe et al., 2021). One issue is that these tasks tap into different senses of inhibitory control, with some, for instance, focusing on inhibiting a prepotent response, while others focus on oculomotor inhibition (Nigg, 2000). To make things even more complicated, few tasks can be used consistently over ages to enable longitudinal examinations. Instead, researchers mix tasks over time, for instance, using the A-not-B task in infancy and the ECITT in early childhood (Cuevas & Bell, 2010; Holmboe et al., 2021). These tasks place different demands on the child, and they differ considerably. Thus, there's no guarantee that such studies are measuring the same inhibitory process over time within-subjects. We note that this is just one skill (inhibitory control). Working memory is equally complex as people have looked at different types of working memory (e.g., verbal working memory: Bull et al., 2008; spatial working memory: Schutte & Spencer, 2009; visual working memory using span tasks, Berry et al., 2018, and change detection, Simmering, 2016).

Perhaps reflecting some of these complexities, recent studies have tried to use factor analytic methods to uncover the structure of EF in early development. Work with children from 3 to 6 years suggests a one factor model can adequately explain the structure of EF within this age group (Wiebe et al., 2011). Other evidence suggests that a two-factor model with working memory and inhibitory control as components can best fit the data when more tasks are included and age is range restricted (Miller, Giesbrecht, Müller, McInerney, & Kerns, 2012). Work by Howard et al. (2015) indicates that EF may go through a period of integration by school age, with performance in working memory, inhibitory control, and cognitive flexibility tasks at 3 years being largely unrelated to one another with stronger cross-task relationships by 4 years. Some of the inconsistency in this literature might reflect the task dependency of the latent variable approach, leading to proposals about EF that are more anchored to the dependent variables from particular tasks rather than the processes behind these tasks (Miller et al., 2012). Indeed, in some cases, factors such as shifting have been interpreted as absent when shifting was not explicitly measured (McAuley & White, 2011, for further discussion see Bardikoff & Sabbagh, 2017).

What is the way forward? Part of the challenge of understanding EF development arises from the reification of clustered measures from the CFA approach and mapping these onto 'components' of EF. That is, the component view of EF conflates the description of associated EF measures with an explanation for why those measures are associated. In resonance with this special issue, we contend that the field needs better theory to examine how *neurocognitive processes* give rise to EF skills. Specifically, if we're going to understand how working memory, inhibitory control, task-switching, and other skills *co-develop* as complex EFs emerge, we need theories that specify how these neurocognitive processes give rise to key EF-related behaviours. Ideally, such theories would be specified in neural terms, yielding insights into a key topic of this special issue: how is EF implemented in the brain.

Developmental changes in the prefrontal cortex correlate with improvements in EF (Diamond, 2002), and individual differences in prefrontal cortex maturity are proposed to underlie observed differences in behavioural performance in EF tasks (McKenna, Rushe, & Woodcock, 2017; Tamm, Menon, & Reiss, 2002). However, the neural underpinnings of EF involve more than the prefrontal cortex. For instance, 3-year-old children present with weak neural interactions within the frontal cortex and unrefined frontoparietal connectivity, but by around 4 years of age, children form stronger neural connections and develop a refined frontoparietal pathway that mediates early EF development (Buss & Spencer, 2018; Perone, Palanisamy, & Carlson, 2018). Other work suggests that neural measures may be predictive from an early age, with performance on the A-not-B task showing relationships to increased frontal activation at 8 months (Bell, 2012).

Our goal in the present paper is to build a neural theory of EF development that integrates key neurocognitive processes like working memory, inhibitory control, and task-switching abilities. We contend that such a framework is needed if we want to understand how EF skills interrelate, and how they combine to produce more complex EFs. Such a framework might also yield insights into how different tasks assess EF as well as how EF is implemented in the brain.

Dynamic field theory is an ideal framework for pursuing this goal. It is the only framework that has been used to quantitatively simulate a diverse array of findings from the literature on EF development using the DCCS task (Buss & Kerr-German, 2019; Buss & Spencer, 2014, 2018; Perone, Molitor, Buss, Spencer, & Samuelson, 2015; Perone, Plebanek, Lorenz, Spencer, & Samuelson, 2019). Our dynamic field model of EF development has also shed new light on EEF. For instance, the model has shown the complexities of bottom-up vs top-down interactions that underlie the performance of younger and older children across easy and hard versions of the task (Buss & Kerr-German, 2019; Buss & Spencer, 2014). The model has also generated novel predictions that have been tested empirically with both behavioural (Perone et al., 2015, 2019) and neural measures (Buss & Spencer, 2018), showing, for instance, the critical role played by binding of features to spatial positions between 3 and 5 years of age. Moreover, this framework is yielding new longitudinal insights as well. For instance, Lowery et al. (2022) and McCraw et al. (2024) have mapped how early dimensional label learning is related to later dimensional attention and EF skills.

Although this prior work has many positives, it is also limited. Our existing EF model only includes a limited sense of working memory; thus, it is not a productive framework to explore the co-development of, for instance, visual working memory and EF. Moreover, a key aspect of recent empirical findings is that the learning of dimensional labels underlies improvements in children's task switching abilities; however, the original model has no mechanism to learn words. Thus, here we integrate our prior model of EF development with a new model – Word-Object Learning via Visual Exploration in Space (WOLVES) – that instantiates visual exploration, visual working memory, feature attention, spatial attention, and word learning processes (Bhat, Spencer, & Samuelson, 2022).

We proceed as follows. In the next section, we provide a brief overview of the framework we use – Dynamic Field Theory. Next, we describe the DF model of EF development, what this model has achieved, but also its limitations. Then, we describe the WOLVES model of cross-situational word learning, what this model explains, and the neurocognitive processes it implements. We then ask if WOLVES 2.0 can capture the processes that underlie performance in the DCCS task, including developmental changes in DCCS performance between 3 and 5 years. To foreshadow our results, we show that improved dimensional labelling performance by the model – as evidenced by simulations of both comprehension and performance tasks – is associated with improved performance in the DCCS task. We conclude by looking toward the future, highlighting how WOLVES 2.0 can provide an integrated framework to probe key questions about the emergence of EF and its relations to other co-developing neurocognitive processes in early development.

Introduction to dynamic field theory (DFT)

We begin with a brief overview of dynamic field theory. DFT proposes that cognition arises from activation within dynamic cortical

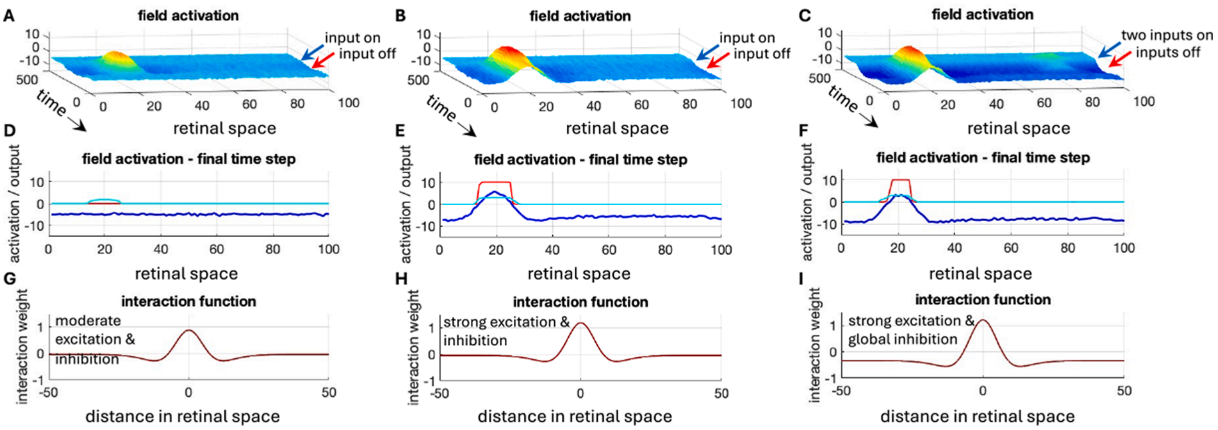

Fig. 2. The graphs in the top row (a, b, c) show how activation (z-axis) evolves through time (y-axis) in the dynamic field across locations in retinal space (x-axis) as inputs are turned on and off. The blue arrow in each panel shows when input is turned on; the red arrow shows when input is turned off. The panels in the middle row (d, e, f) show the state of the field activity at the last time step. The dark blue line shows the activation level (y-axis) over retinal space (x-axis), the red line shows which neurons are engaged in neural interactions (i.e., above zero activity), and the cyan line shows the strength of the memory trace. The graphs in the bottom row (g, h, i) show the rule governing how neurons talk to one another, with local excitation around the activated site (0) and surround inhibition to the left and right (at farther distances in retinal space). The simulation in the left column shows an 'input-driven' DF with moderate excitation and inhibition. The simulation in the middle column shows a 'self-sustaining' DF with strong excitation/inhibition where activity is sustained even though the input is turned off. The simulation in the right column shows a 'selective' DF with strong excitation and global inhibition. Here, the field 'selects' one input even though two inputs are presented (see panel c).

fields simulated using dynamic fields (DFs). DFs are collections or ‘populations’ of neurons, which are wired up in a way that mimics how neural populations in the brain are wired based on the perceptual, cognitive, or motor dimensions over which they are distributed. For instance, we might have a DF with neurons that represent retinal spatial position (see Fig. 2A; e.g., Markounikau et al., 2010). Here, neurons on the left side of the field would have receptive fields that ‘prefer’ inputs at ‘left’ spatial positions, while neurons on the right side of the field would have receptive fields that ‘prefer’ inputs at ‘right’ spatial positions. With this setup, ‘neighbouring’ neurons in the field (i.e., neurons that ‘code for’ similar spatial positions) would excite one another (local excitation), while neurons far apart in the field would inhibit one another (surround inhibition; see Fig. 2G-I). This allows activation ‘peaks’ in dynamic fields to represent the metric details of the input pattern. For instance, one could present a visual stimulus 20° to the left of midline. This would activate neurons that ‘prefer’ inputs on the left side of the retina, particularly those that really like the 20° location. The excited neurons would then activate their local neighbours, further driving up activation. At the same time, the excited neurons would inhibit neurons ‘far away’ in the field, preventing excitation from growing at other sites. The result will be a ‘peak’ of activation – a local above-zero ‘bump’ of activity centred at 20° in the field – which faithfully represents the presence of the visual stimulus at this location (see Fig. 2A).

Interestingly, by changing the strength of connections in a dynamic field, one can create different types of activation patterns through time. For instance, if local excitation and inhibition are moderate, peaks will be stable when inputs are present (e.g., visible), but decay back to a resting level (i.e., no peak) when the input disappears. We call such fields ‘*input-driven*’ (see Fig. 2A; Schöner et al., 2016). If, however, excitation and inhibition are strong, peaks can be *self-sustaining* and remain active even if input is removed, acting as a form of working memory to maintain information even when inputs are no longer available (Fig. 2B; Schöner et al., 2016; Spencer, 2020). Finally, with the addition of strong global inhibition (see Fig. 2I), DFs can be *selective or ‘winner-take-all’*, forming only one peak at the time. Consequently, if two inputs are present, the field will ‘pick’ one, forming a peak at one input location and suppressing activation at the other input location (Fig. 2C).

DF architectures can be constructed by coupling fields together, that is, by having one DF pass activation to another DF and vice versa. This must be done carefully as each DF has receptive fields ‘tuned’ to a particular dimension. For instance, we might have one DF which is sensitive to retinal spatial position and another that is sensitive to colour (i.e., hue). In this case, how would we pass activation from the spatial field to the colour field as these dimensions don’t have any a priori relationship (i.e., blue objects can be on the left or right; the same for red objects)? One solution is that the spatial and colour fields can be joined up into a two-dimensional DF where we

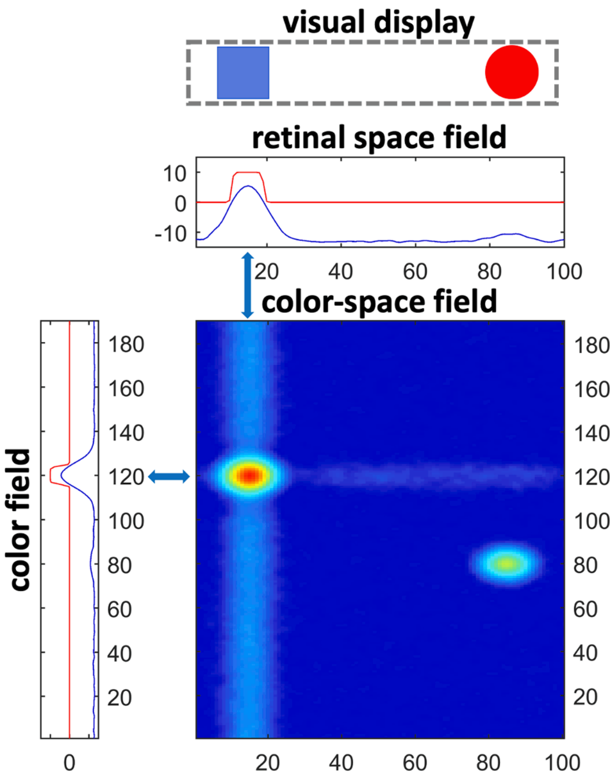

Fig. 3. An example of a DF model with a two-dimensional color-space field which shares activation with a retinal space field (top) and a color field (left). Input from the visual display is passed directly into the two-dimensional field, mapping the location of the input along both the spatial dimension (15 degrees; x-axis) and the color dimension (hue value 120 = ’blue’; y-axis). The intensity of activation in the 2D field is captured by the color value with ‘hotter’ colors showing more intense activation. Activation in the color-space field then projects activation to the retinal space field (top) and the color field (left), building a peak in those fields (see activation profile in blue; y-axis). The red curve in the space and color fields shows which neurons are above 0 activation (i.e., ‘active’ or ‘on’). These peaks then pass activation back to the 2D field, further amplifying the activation pattern in the color-space field.

have neurons ‘tuned’ to all possible combinations of space and colour (see the ‘colour-space’ field in Fig. 3). Here, the presentation of a blue object on the left would lead to a peak in the ‘blue-left’ region of a colour-space field while the presentation of a red object on the right would lead to a peak in the ‘red-right’ region of the colour-space field (Fig. 3). Interestingly, we can also pass activation to and from this two-dimensional field to separate retinal space and colour fields as shown in Fig. 3 (see bi-directional blue arrows). This

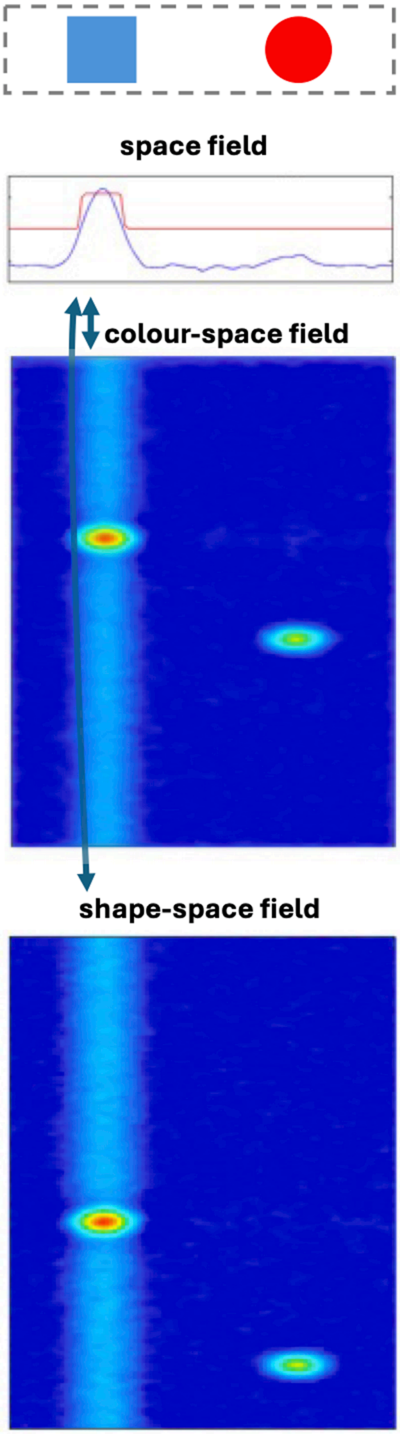

Fig. 4. An example object representation model with a spatial field, a colour-space field, and a shape-space field. This model is being shown a blue square on the left and a red circle on the right (see top display panel). The model is attending to the blue square as indicated by the peak in the spatial field. This peak passes activation to both feature-space fields on the left side (see blue arrows), ‘binding’ the blue and square features together.

enables the neural architecture to represent that the object is on the 'left' (in the space field), it is 'blue' (in the colour field), and the 'blue object is on the left' (in the colour-space field). Note that we don't always have to represent information in multi-dimensional DFs; indeed, there are good reasons not to do this as very high-dimensional DFs have a lot of neurons (more than in the human brain). Thus, we have proposed some rules for how to join dimensions up to create large neural architectures using special 'binding' dimensions. To date, these include spatial dimensions (as in Fig. 3) as well as more abstract binding dimensions such as words or labels. Because WOLVES uses these binding dimensions, we return to this issue below.

So far, we have discussed how DFs capture patterns of neural activation from second-to-second through time. But we also need a way for the patterns of activity to be carried forward over longer periods of time. For instance, how can a pattern of activity on trial 1 impact a future pattern of activity on trial 2 or trial 12? Or, more to the point of this special issue, how can a pattern of activity be learned over days, weeks, months, and years of experience? DFT uses a variant of Hebbian learning to capture such effects. In particular, 'memory traces' can form in DFs when strong peaks build. In this case, the peak boosts activation in a memory trace which feeds back on the field activity, strengthening local excitation in that region of the field. For instance, if the model were asked to encode and respond to the blue item on the left, a memory trace might form which makes the model faster to respond to blue things on the left on future trials (because the memory trace boosts local excitation in this region of the field). This results in a 'pre-shaping' effect, facilitating recognition of familiar inputs. By building up memory traces over long periods of time, we can simulate changes over learning and development. More recent models have also implemented more canonical versions of Hebbian learning in dynamic field architectures (Tekulve & Schöner, 2022). We'll return to this issue in the Discussion.

Overview of the DF model of EF development

In previous work, we built a DF model of the development of EF that replicated children's behaviour on the DCCS task. Here, we present a brief overview of this prior work as it sets the stage for our integration of this EF model with WOLVES. The EF model is composed of two primary systems. One system forms object representations by binding visual features to spatial locations; a second system forms associations between dimensional labels like 'colour' and 'shape' and visual features. The object representation system contains a spatial planning field, a colour-space field, and a shape-space field. The spatial planning field is reciprocally coupled to each of the feature-space fields. An object representation in the model would reflect a pattern of activation peaks across these fields at a common spatial location. Fig. 4 shows a neural architecture with a spatial planning field, a colour-space field, and a shape-space field.

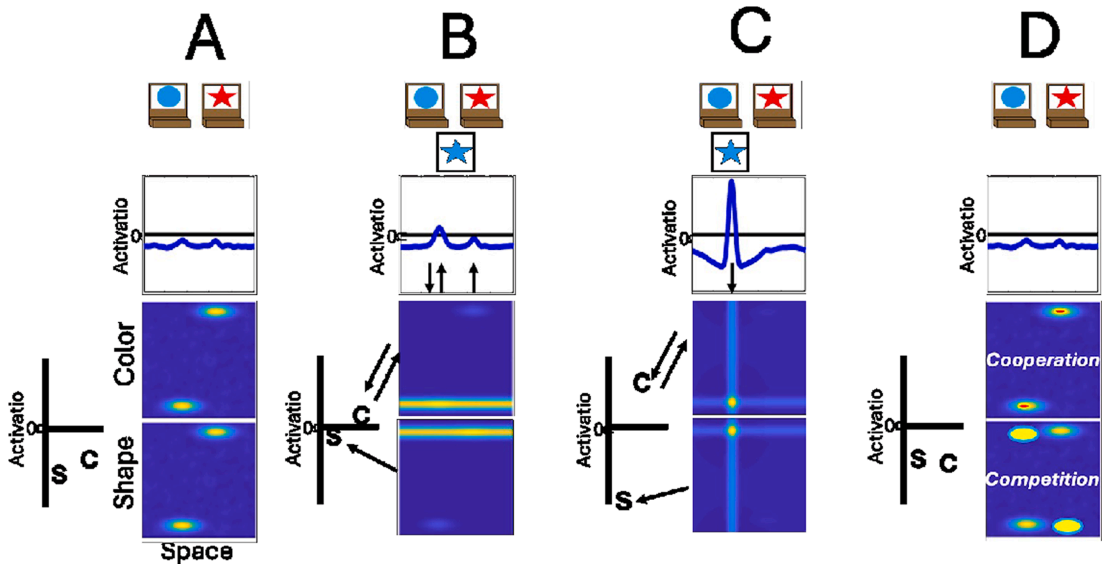

Fig. 5. The DF model of EF development. The neural architecture in each panel has an object representation system (see spatial field at top, colour-space field in middle, and shape-space field at bottom) along with nodes corresponding to "colour" and "shape" labels (see activation plot at left of each panel where "C" is the activation state of the "colour" node and "S" is the activation state of the "shape" node). A) Initial state of model when told to play "colour" game. At this point, the "C" node is given an input which boosts its baseline activation above that of the "S" node. B) Model is shown a blue star (see horizontal ridges in colour-space and shape-space fields). The inputs to the colour-space and shape-space fields are reaching the activation threshold and are sending output to the spatial field (note the arrows pointing upward to the spatial field) and to the "C" and "S" nodes (note the arrows pointing from the feature-space fields to the label nodes). The "C" node is breaching the activation threshold (0) and is sending activation to the colour-space field (note the arrow pointing back to the colour-space field). C) Model sorts the blue star to the left due to the boost provided by the "C" node to the colour-space field. D) State of inputs to the model after sorting cards during the pre-switch phase with 'cooperation' in the colour-space field and 'competition' in the shape-space field. See text for additional details. Note that the feature-binding process is mediated by activation sent from the spatial WM field to the feature-space WM fields (note the downward pointing arrow from the spatial field in Panels B and C).

Excitatory activation along the spatial dimension is reciprocally shared with both the colour-space field and the shape-space field. Note that the feature-space fields do not pass activation between one another. With this architecture, the model can 'bind' features through their shared spatial positions, knowing not just that there is a blue object on the left, but that the object is a blue square.

This model also uses dimensional labels. As illustrated in Fig. 5, the dimensional label system is composed of a simplified population of neurons, one for the label "colour" and the other for the label "shape" (see left 'activation' axis in Fig. 5A). These label units are reciprocally coupled to the feature-space fields in the object representation system. For example, when the "colour" label is activated, it sends a global excitatory boost to the colour-space field and vice versa. This coupling gives rise to a form of dimensional attention by which the model can selectively enhance processing of task-relevant feature dimensions. Developmental improvements in performance are a result of changes in the model's knowledge of labels which is reflected in the strength of associations between labels and visual features. Younger children have weaker associations between, for instance, the colour label and the colour dimension. Older children have stronger associations, with more selective associations between the colour label and the colour dimension.

Fig. 5A shows the model at the start of the pre-switch phase. The target cards (see display at the top of the figure) are presented as a pattern of spatially localized feature inputs to the feature-space fields. At the leftward location, there are inputs for blue in the colour-space field and circle in the shape-space field. At the rightward location, there are inputs for red in the colour-space field and star in the shape-space field. Next, the model is instructed to sort by colour (Fig. 5B). This is reflected by the input to the colour neuron ("C") which increased its baseline level of activation higher than the shape neuron ("S"). A test card is presented to the model by providing feature inputs for the combination of features on the test card, in this case a blue star. The feature inputs are presented as ridges that activate the feature across the spatial dimension (see yellow horizontal ridges in Panel B). For blue, this ridge overlaps with the target input at the leftward location. For star, this ridge overlaps with the target input at the rightward location. This spatial conflict is resolved via the input from the label system. Because the "colour" neuron has more activity due to the instruction to sort by "colour", the "colour" neuron reaches activation threshold sooner and inhibits the "shape" neuron. The "colour" neuron also boosts global activation within the colour-space field. Consequently, activation builds more quickly at the leftward location in the spatial field (see emerging peak in Fig. 5B at the leftward location in the spatial field). In Panel C, the model has made a decision to sort this test card to the leftward location reflected by the binding of the object's features to the leftward location (see peaks in Fig. 5C at the leftward location in the spatial field, the colour-space field, and the shape-space field with strong activation of the 'colour' node).

Panels B and C illustrate the dimensional attention of the model with the arrows highlighting patterns of interactions between model components. In these panels there are reciprocal arrows between the "colour" node and the colour-space feature field whereas there is only a unidirectional arrow between the "shape" node and the shape-space feature field. Thus, both feature-space fields have built activation, which is being sent to their associated dimensional labels, but only the "colour" node has been activated above threshold (0) due to the inhibitory competition between label nodes. In this case, only the colour-space feature field is boosted. The impact of this form of dimensional attention is to prioritize information conveyed by inputs and memory traces within the colour-space field. For example, in Panel B, there is activation at both spatial locations being sent to the spatial WM field (note the two arrows pointing upward). The colour-space field has greater activation at the left and the shape-space field has more activation at the right. Output from the colour-space field is stronger than the output from the shape-space field, leading the model to bind the features on the test card to the leftward location (note the downward arrows from the spatial WM field).

One of the key aspects of the DCCS task is the conflict introduced between the sorting instructions during the post-switch phase and behaviour during the pre-switch phase. As the model makes decisions and sorts cards over the course of the pre-switch phase, memory traces build up that strengthen the association between visual features and spatial locations. For example, when sorting a blue star to the left, the model will form associations that strengthen blue at the left location in the colour-space field and star at the left location in the shape-space field. Similarly, the model will form associations between red at the right location in the colour-space field and circle at the right location in the shape-space field. This leads to a state of 'cooperation' in the colour-space field because the target card inputs match the associations built over the pre-switch trials (see Fig. 5D). By contrast, there is 'competition' in the shape-space field because the model remembers sorting stars to the left and circles to the right, but it is seeing a circle on the left and a star on the right.

When instructed to switch rules and sort by shape in the post-switch phase, the model gets a boost to the "S" node (see Fig. 5D); however, for this boost to be sufficient to help the model sort by shape, this boost must overcome the competition in the shape-space field, and this must happen quickly enough to override the cooperation in the colour-space field. Together, this means that a model with weak dimensional attention, that is, a weak boost to the shape node, is unlikely to follow the post-switch rules.

Buss and Spencer (2014) presented a series of simulations that provided a unifying explanation of findings from the literature. To simulate developmental differences between these age groups, two sets of parameters were used that reflected different degrees of associations between labels and visual features. The model of 3-year-olds was composed of a population of parameters with a weak dimensional boost and weak coupling between each dimensional node and the respective feature dimension. By contrast, the model of 4-year-olds was composed of a population of parameters with a stronger dimensional boost and stronger coupling between each dimensional node and the respective feature dimension. Critically, performance of the model is influenced by a range of factors beyond the strength of associations between labels and visual features, especially in early development when the associations between labels and features is weak. Across the different versions of the DCCS task that were simulated, the performance of these groups of models depended on the interaction between memory traces and the task-structure provided by the target cards. For example, in a total-change condition, all of the features involved in the task are changed for the post-switch phase. That is, if children sorted by colour with red stars and blue circles during the pre-switch phase, then they would be instructed to sort by shape with green squares and yellow triangles during the post-switch phase. In this case, most 3-year-olds have little difficulty switching rules.

In total, the model replicated the behaviour of 3- and 4-year-olds across 16 different conditions of the DCCS task, predicted performance of 3-year-olds in new conditions that altered the spatial configuration of the target cards between pre- and post-switch phase

(Buss & Spencer, 2014), predicted associations in performance between the DCCS task and other measures of dimensional attention that do not use explicit instructions to attend to dimensions (Buss & Kerr-German, 2019), and predicted the influence of pre-exposure to the post-switch dimension outside of the context of the DCCS task (Perone et al., 2015, 2019). Notably, the model shows how executive function can arise from a neurocognitive system in a way that does not require explicit executive function components to be built in. Here, we note a distinction between neural processes and psychological constructs. Although it is true, for instance, that the model includes inhibitory neural processes, these are not isomorphic with the psychological construct of inhibitory control. In particular, there is not an explicit inhibitory control module in the model that gives rise to the behaviour associated with inhibitory control – successfully inhibiting a bias to perseveratively respond by colour in the post-switch phase. Instead, the model relies on a system of basic components that operate with generic neural processes such as local excitation and lateral inhibition. These neural components are integrated in such a way to enable performance in EF tasks. The object representation system was designed to build object representations in a neurally-realistic fashion by binding visual features to spatial locations. Here, the basic components are neural populations that encode visual features such as colour, shape, and space. Likewise, the label representation system was designed to form associations between labels and visual features across cortical fields tuned to these dimensions. By combining these neural processes into an integrated neural system, new functions emerge from the model such as card sorting behaviours. In this way, executive function is a property of the system, not a component of it. We expand on these ideas in the Discussion.

Inspired by the model, subsequent research has begun to explore the relationship between dimensional label representations and EF. A recent longitudinal study examined whether neural function during tasks measuring the comprehension and production of dimensional labels would predict performance on the DCCS task and whether this prediction was specific to tasks requiring dimensional attention (Lowery et al., 2022). At 33-months, children performed simple comprehension (“show me the red one”) and production (“what colour is this”) tasks with canonical shapes and colours while neural activity was measured. At 45 months of age, children performed the DCCS task and the flanker task. The flanker task is a standard measure of attentional control which requires spatial selective attention (respond to the central item and ignore flanking items) and does not contain dimensional conflict as in the DCCS task. Results from this study showed that activation in the left frontal cortex during label production tasks at 33 months was predictive of performance on the DCCS task at 45 months. Measures of neural function during label production tasks were not related to flanker task performance. Thus, evidence suggests that dimensional label learning impacts measures of EF that require attention to visual dimensions. Given that the DF model of EF development does not have any particular grounding in word learning, these data motivated us to consider the WOLVES model which specifies how words are learned.

Overview of WOLVES

A key goal of the present report is to integrate our DF model of EF development with a new model of early word learning: Word-Object Learning via Visual Exploration in Space (WOLVES). There are two motivating factors for pursuing this goal. First, WOLVES integrates many of the component processes involved in EF, including attention and working memory. Second, dimensional label learning is a key precursor to EF skills in early development; thus, there is evidence that word learning and EF are developmentally related skills. Before pursuing this integration, we first present a brief overview of WOLVES.

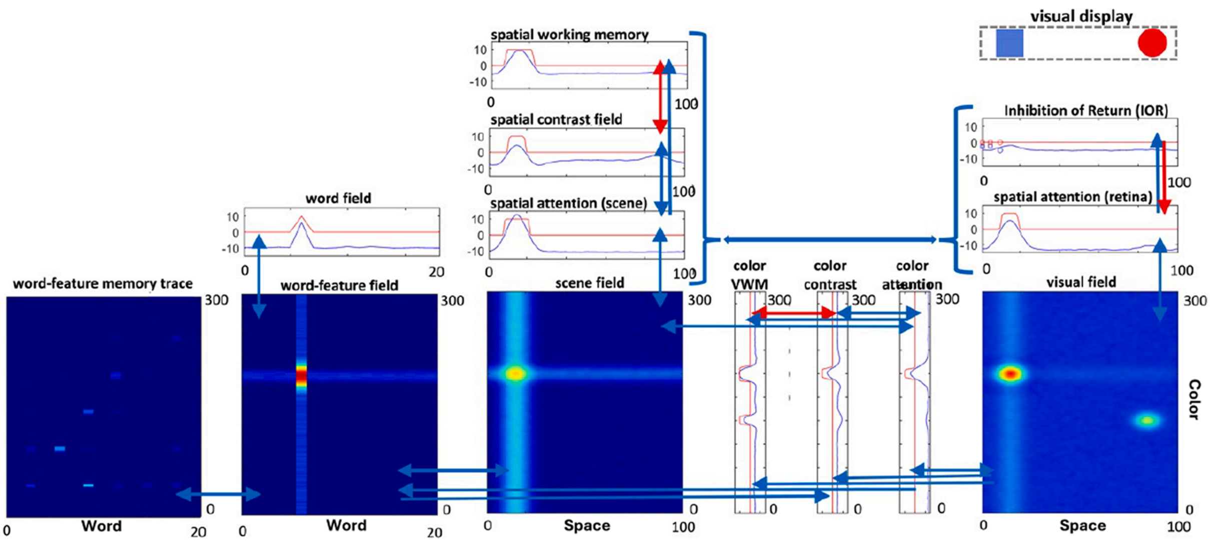

Fig. 6. The WOLVES model neural architecture. The one-dimensional (1D) and two-dimensional (2D) dynamic fields (DFs) in the model are responding to the visual display in the top right. Arrows represent uni/bidirectional connectivity (blue: excitatory, red: inhibitory). 1D fields show activation profile in blue and above-threshold neural activity in red. 2D fields visual and scene fields show activation representing specific colours in specific spatial locations, with higher activation in ‘hotter’ colours. The 2D word-feature and word-feature memory trace field show binding of specific words and the object colour feature. Some working memory and memory trace fields are not shown for simplicity.

WOLVES was developed by Bhat et al. (2022) to examine the multiple processes supporting learning of novel mappings between words and objects by infants, children and adults. A visual representation of WOLVES is depicted in Fig. 6. With the exception of the visual display in the upper right, each box in Fig. 6 depicts a dynamic field. Blue arrows indicate excitatory connections, red arrows inhibitory. The figure depicts fields corresponding to lower-level cortical fields on the right and higher-level cortical fields to the left. Thus, information flows through the model in a right-to-left manner, corresponding to flowing from lower-level visual areas in the back of the brain, to higher level areas in the frontal cortex.

Beginning with the fields on the right, the model captures aspects of processing related to visual exploration in space (Perone & Spencer, 2013; Schneegans, Spencer, & Schöner, 2016). Initially, visual inputs (for example the blue square and red circle in the visual display of Fig. 6) are presented to a visual field (to the far right in Fig. 6), encoding features like colour and shape, and where those features are in the visual field (like a retina). This information is then passed along two pathways: a dorsal pathway for spatial information (see horizontal 1-dimensional fields in Fig. 6) and a ventral pathway for feature information (vertical 1-dimensional fields in Fig. 6). Along the dorsal pathway, WOLVES represents the location of objects in the world, including spatial working memory for crucial locations within the current scene. The ventral pathway represents object features (like colour), including visual working memory for the specific features of presented stimuli. A scene attention field re-combines this information into a scene representation of what is where (Treisman & Gelade, 1980) using focal attention.

The left side of Fig. 6 depicts the portion of the model that handles learning of mappings between words and objects (Samuelson, Smith, Perry, & Spencer, 2011). Auditory inputs are presented to the word field, which identifies the labels presented during a task. Words and object features are then integrated in the word-feature (binding) field. Across repeated presentations, memory traces of these bindings build up in a long-term memory trace layer (far left). Over time, these become able to support recognition of word-object mappings and, via top-down connectivity, guide attention to familiar items when a word is presented.

Fig. 6 illustrates neural activation in WOLVES at a moment when the model recalls a word-object mapping and directs attention to the corresponding object in the visual field. To illustrate how WOLVES works, we will walk through three autonomous cycles of action in the model: the visual exploration in space cycle, the word-object learning cycle, and the top-down attention cycle.

When a visual display is presented to the model, the **visual exploration in space cycle** starts. For example, when the visual display in Fig. 6 is presented, activation for each element is generated. The peak on the left is at the intersection of the colour blue and the left location, while the peak on the right is at the intersection of the colour red and the right location. The blue and red features are then passed to the colour contrast field which detects visual novelty (e.g., any feature that fails to match the contents of working memory). All inputs are novel at the start of each visual exploration in space cycle, but at other times, peaks in visual working memory can suppress peaks in the contrast field. This differentiates what is 'known' (e.g., has a peak in visual working memory) from what is 'novel' (e.g., has a peak in the contrast field).

Peaks in the contrast field pass activation to the 'winter-take-all' attention field, which supports only a single peak—the focus of attention. This attention peak amplifies activation in the visual field, selecting this peak as the focus of attention and boosting spatial attention. In Fig. 6, WOLVES is attending to the blue square on the left with the colour attention field highlighting the 'blue' feature. Attended features are consolidated in working memory. This occurs along the ventral pathway in visual working memory (VWM) for features (colour in Fig. 6), and along the dorsal pathway in spatial working memory (SWM). Consolidated features are passed from working memory fields to the two-dimensional scene field which binds the spatial and feature information together at the level of the spatial scene. This representation is typically anchored to the body or an external frame of reference, and not coded in a retinal position. The inhibition of return field detects the consolidated item, suppressing spatial attention and releasing it from focus. The model is then free to attend to another novel item, detected by the feature contrast field, and the cycle repeats.

The **word-object learning cycle** begins when a peak is created in the word field by the presentation of an auditory word. This peak passes activation to the word-feature field. If the model is attending to a visual object at the same time, the inputs will combine. This binds the word to visual features (see red hot spot of activation in the word-feature field) and leaves a trace in the memory trace layer (far left). These traces build slowly over time and decay even more slowly. With repeated presentations, the model learns word-object mappings. This can be seen as the light blue dots in Fig. 6, indicating word-object mappings that have previously been encoded by the model.

The third and final cycle in WOLVES is the **top-down attention cycle**. This cycle engages when the model has a strong memory trace for a word-object mapping. Presenting the word causes a peak in the word-feature field at the location of the strong memory trace. This peak passes a 'top-down' signal from the word-feature field to the contrast field. This signal activates the associated feature, directing attention to the object via interactions with the colour attention field. This is how WOLVES directs attention to a visual object when hearing a known word.

These cycles emerge over multiple timescales as neural activation propagates in the model as it performs a task. On a real timescale of milliseconds and seconds, the model autonomously shifts attention between objects in the visual field, recognizes words and binds visual features to words. Learning occurs over a longer timescale, affecting each cycle differently. Visual habituation emerges as strong memory traces alter the visual exploration in space cycle, causing the model to swiftly release fixation from 'known' items and spend more time exploring novel items. The word-object learning cycle benefits from learning cascades as WOLVES repeatedly builds correct word-object co-occurrences. Once these memory traces are strong, they can block the formation of new incorrect associations. Learning also impacts the top-down attention cycle, as strong memory traces direct attention to labelled objects.

To summarise, on a real timescale of milliseconds and seconds, the model autonomously alternates attention between visible objects in the visual field, recognizing words and establishing connections between visual features and words. This real-time behaviour aligns with participants' looking behaviour, enabling the model to be embedded in the same scenarios as participants using identical visual and auditory inputs as those presented to participants with the same timings. For example, Bhat et al. (2022) used WOLVES to capture data from 7 adult and 5 child studies of Cross Situational Word Learning (CSWL). In CSWL studies, participants are presented with multiple words and objects together such that it is not clear which word labels which object on an individual trial. The correct mapping only becomes clear when a given word-object pair repeatedly co-occur over multiple trials. Bhat et al. presented WOLVES with the exact experimental timings and protocols presented to participants in the simulated studies and showed it fit data better than two competitor models and generalized better to three “held-out” experiments. Further, because WOLVES produces data corresponding to participants’ looking behaviour, it captures more data. Finally, Bhat and colleagues proposed the first developmental account of CSWL, showing how memory processes change from infancy to adulthood and how this impacts CSWL. Further, WOLVES sheds light on the processes supporting CSWL, showing how visual exploration and selective attention in CSWL are dependent on and also indicative of learning and how learning is driven by the real-time synchrony of words and gaze and dynamically constrained by memory processes.

WOLVES 2.0: Toward an integrated theory of EF development and word learning

The goal of this section is to describe WOLVES 2.0, that is, a modified version of WOLVES that implements key aspects of EF development. As noted previously, this goal was motivated by two factors. First, WOLVES already integrates attention, working memory, and word learning – ingredients that may be important for the early development of EF. Second, data from Lowery, Buss, and colleagues (Lowery et al., 2022) show that the learning of dimensional labels precedes and predicts developmental improvements in EF skills as evidenced by performance in the DCCS task.

The specific goals for our modelling efforts were as follows. Our first goal was to modify WOLVES – creating WOLVES 2.0 – so it could sort cards successfully when placed in the DCCS task. To preview results on this front, one new connection was needed in WOLVES to enable DCCS performance. Second, we wanted to examine if WOLVES could capture developmental changes in DCCS performance, that is, would WOLVES perseverate with ‘weak’ learning of dimensional labels and sort correctly with ‘strong’ learning of dimensional labels. To probe dimensional label learning, we had a third goal: to place WOLVES in comprehension and production tasks and examine if the model could capture developmental improvements in these tasks as well. Thus, taken together, our first 3 goals were to move systematically toward a new theory of EF development based on the idea proposed by Buss and colleagues (Buss & Kerr-German, 2019; Buss & Spencer, 2014; Lowery et al., 2022) that improvements in dimensional label learning predict improvements in DCCS performance. The fourth goal was to quantitatively reproduce findings from all three tasks (DCCS, comprehension,

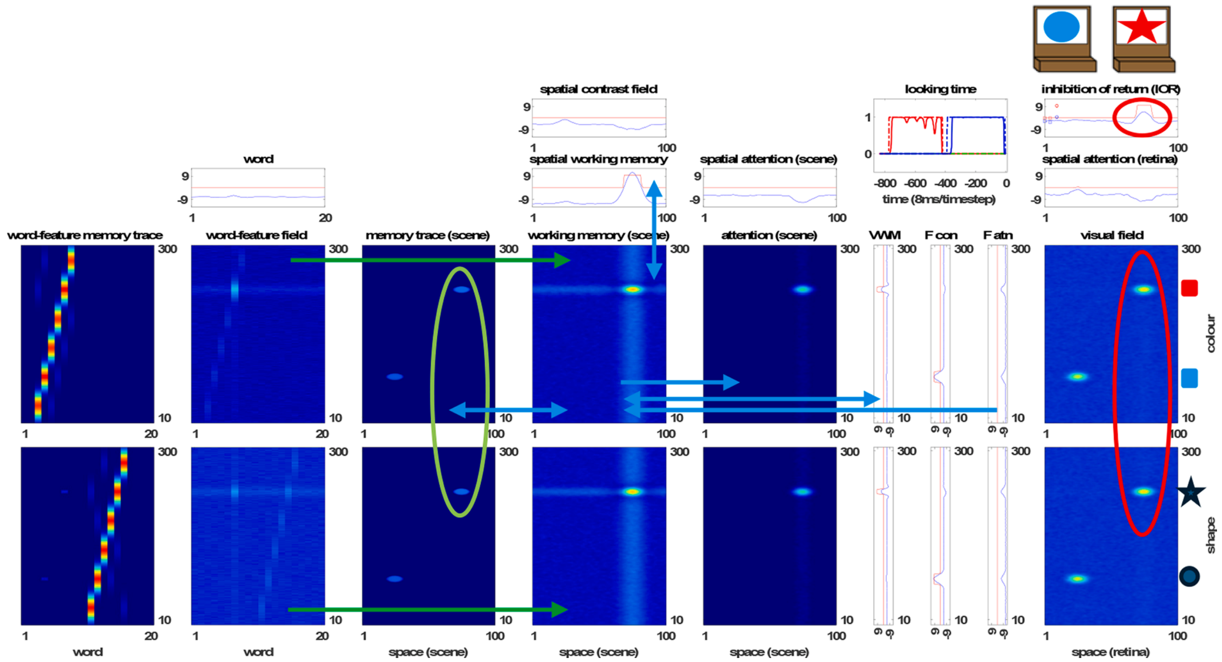

Fig. 7. The WOLVES 2.0 architecture. The one-dimensional (1D) and two-dimensional (2D) dynamic fields (DFs) are responding to the visual display in the top right. Model connectivity is as in Figure 6. Here, we also show the scene working memory fields (column 4 from the left) and scene memory traces (column 3 from the left). These were included in the original WOLVES model, but not shown for simplicity. As these fields are critical to the functioning of WOLVES 2.0, we show them here, along with their connectivity (see blue excitatory arrows; note that connectivity is only shown for the colour feature pathway but is identical along the shape pathway). 1D fields show activation profile in blue and above-threshold neural activity in red. 2D fields show activation in ‘hotter’ colours. The dark green arrows show the one addition to the WOLVES architecture to create WOLVES 2.0. Red ovals highlight activation patterns in the inhibition of return field and the visual fields (see text for details). Green oval highlights the representation of the red star in the model’s scene memory trace (see text for additional details). Finally, the ‘looking time’ plot shows the past history of the model’s looking behaviour; red curve indicates looking to the left, blue curve indicates looking to the right.

production) including developmental changes in performance. As a fifth goal, we then extended this model to another key set of findings from the literature, simulating data from Yerys and Munakata (2006) that also probed the role of labels in the DCCS task. Finally, we asked if WOLVES 2.0 would still robustly reproduce developmental findings from cross-situational word learning tasks (Bhat et al., 2022).

What would success with goals 1–6 achieve? Success would mechanistically link EF development in the DCCS task with dimensional label learning while also embedding word learning and decision-making processes together into an integrated framework that allows researchers to study how attention, working memory, inhibitory control, and word-mediated task switching processes work together as a system in early development – ideas we articulate in detail in the Discussion.

Can WOLVES 2.0 show task-switching in the DCCS task? Our simulation work started with a programming task: to embed WOLVES in the DCCS task. This is a key advantage of neural process models – because such models capture the real-time neural processes that, in our case, underlie autonomous visual exploration, attention, working memory, and word learning, it's relatively straightforward to explore how well they generalise to other tasks. All this requires is placing the model in the new task and exploring what it does (Bhat, Samuelson, & Spencer, 2023).

Explorations of WOLVES led to some initial insights. First, in our prior EF model (Buss & Spencer, 2014), inputs were fed directly into the object representation system, creating an input-driven representation of which features were where in the task space. At the time, this was a placeholder for a more elaborated model that actively explored a visual environment. WOLVES is such a model. Thus, when we embedded WOLVES in the DCCS task, it immediately started looking back and forth between the target cards, building a scene representation by creating peaks in its scene working memory. Those peaks then built up longer-lasting connection weights in the scene memory trace.

An example of the model is shown in Fig. 7. This figure is comparable in layout to Fig. 6 but with a few additions. First, we show two additional layers: the scene working memory and the scene memory trace (along with the associated connectivity into and out of these layers; see blue excitatory arrows). These layers were part of the original WOLVES model; however, the word learning behaviours WOLVES captured did not depend on these layers so we did not show them for simplicity. Here, these layers play an important role in DCCS performance, so we have added them to the model depiction. The second change relative to Fig. 6 is that we are showing both the feature pathways for colour (top horizontal row of fields in Fig. 7) and for shape (bottom horizontal row of fields in Fig. 7). The final change is highlighted by the dark green arrows from the word-feature fields to the scene working memory fields. The connections were not present in WOLVES. As we discuss below, they constitute the only change moving from WOLVES to WOLVES 2.0.

The model in Fig. 7 has visually explored the blue circle on the left and the red star on the right across several looks back and forth (see 'looking time' plot). It is just finishing a look to the red star on the right (see red oval in visual field and red circle highlighting the inhibition of return peak which is suppressing retinal spatial attention). The result of this visual exploration is that the model has built a memory of the blue circle on the left and the red star on the right (see green oval in the scene memory trace highlighting the memory of the red star on the right). Note that WOLVES has two spatial reference frames – a retinal frame where the objects appear on the retina (see fields on the right column of Fig. 7) and a scene frame which captures the objects in space (see fields in the middle three columns of

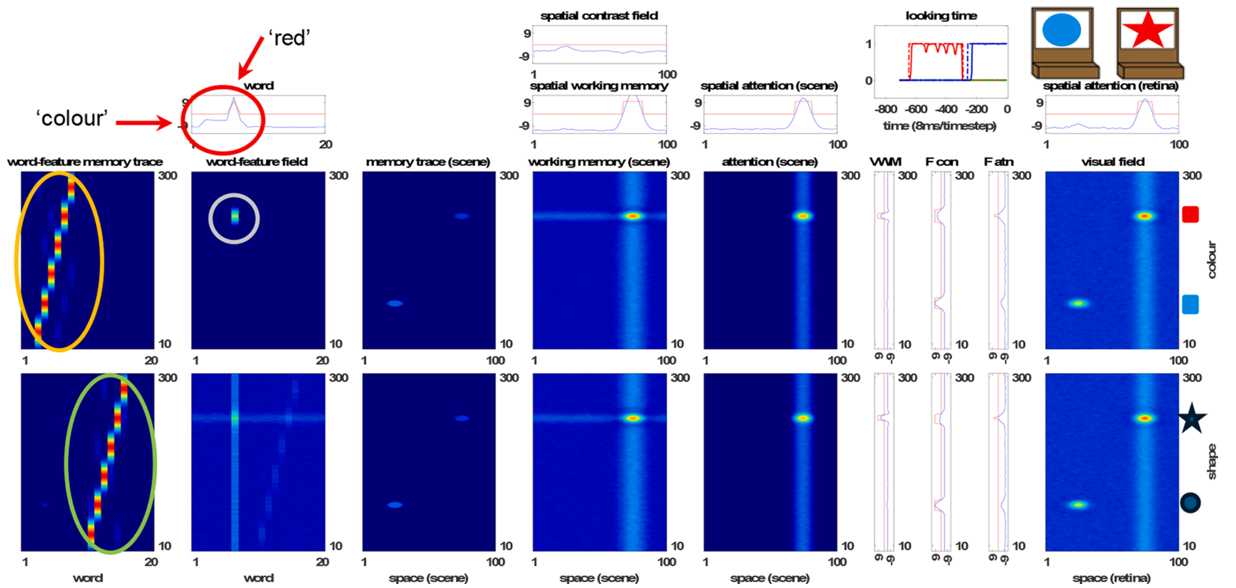

Fig. 8. The 'old' WOLVES 2.0 model during the instruction phase of the DCCS task. The experimenter directs the model's attention to the red star on the right (see visual and spatial attention fields on the right side of the figure) while giving two auditory inputs that 'red' things go here in the 'colour' game (see red oval highlighting activation patterns in the word field). Orange oval highlights the learned pattern of mappings from colours (hue values) to colour labels (e.g., 'red'). Green oval highlights the mappings from shape values to shape labels (e.g., 'star'). Grey oval highlights a peak in the word-colour field that maps the 'red' label to the red hue value shown on the red star target card in the task space.

Fig. 7). For simplicity, we typically keep these in alignment; however, in robotic implementations of related models, we have shown how the concepts in WOLVES can be extended to include multiple reference frames with transformations from one frame to the next (Richter, Lins, & Schöner, 2021).

One consequence of having a visual exploratory model like WOLVES is that we also need to build in all phases of the task including the instruction phase and demonstration trials that occur prior to the pre-switch phase. We show one part of the instruction phase in Fig. 8. Here, we are instructing the model that it is playing the 'colour' game and in the 'colour' game 'red' things go 'here' (e.g., to the right). We instantiate the 'colour' instruction by activating the full set of colour labels the model knows. This is captured by the flat boost across all the colour labels in the red oval in the word field. Here we are assuming that children have learned that the word 'colour' is associated with 'red', 'green', 'blue' and so on. Note that we have yet to model the learning of this association. For now, we simply build this knowledge into the model and assume the association between a 'colour' node and the individual colour labels can be learned (Sandhofer & Smith, 1999; Verdine, Lucca, Golinkoff, Hirsh-Pasek, & Newcombe, 2016). The activation profile in the word field also shows a peak at the 'red' label (due to the verbal input). This sends a vertical 'ridge' into the two word-feature fields at the 'red' label (see, e.g., the vertical ridge in the word-shape field). Because the model is attending to the red star to the right (due to the experimenter directing attention to this location; see peaks in the visual field at the right location), the model maps the word 'red' to the red hue value in the word-colour field, building a peak at this site (see grey circle).

The orange oval shows the strength of the memory trace for each mapping between hue value and label. Notice that the 'red' label has a strong Gaussian centred on the red hue value. This captures the model's prior learning that 'red' maps to a range of hue values around a prototypical red value. Again, in the current report, we don't simulate this learning; that will be a topic for future work. The green oval shows the representation of the shape labels. Notice that the 'star' label has a strong representation of the star feature value. This explains why the model only forms a weak mapping of the word 'red' to the star feature in the word-shape field where the vertical and horizontal ridges meet. The strong mapping between 'star' and star feature is partially blocking the model from forming a mapping between 'red' and star feature (for discussion of this 'blocking' property of WOLVES, see Bhat et al., 2022).

The next step in the instructions phase is a demonstration trial shown in Fig. 9. Here, the experimenter shows the child one of the test cards and sorts this card to one side. For instance, in Fig. 9, we are showing the model a red circle card and placing it to the right in the 'colour' game (see 'colour' boost in the word field). This was the second demonstration in this simulation. In the first demonstration, we showed the model that blue stars go to the left. This left a memory of this event in the scene memory trace highlighted by the green oval.

Fig. 10 shows the scene memory trace after the full instruction phase. Note that there is a cooperative pattern of weights in the colour-space scene memory trace in the top panel (see red oval): the model has always seen that red things go to the right. By contrast, there is a competitive pattern in the shape-space scene memory trace in the bottom panel (see orange oval): the model sees a star on the right but a star was sorted to the left on the demonstration trial. Note that this is the same pattern of cooperation and competition described in Buss and Spencer (2014). The difference here is that WOLVES builds this scene representation via autonomous visual exploration in the context of task instructions.

The next phase in the DCCS task is the pre-switch phase. This is where the one modification to the WOLVES architecture comes into play: the added excitatory connection from the word-feature fields to the scene working memory (see dark green arrows in Fig. 11).

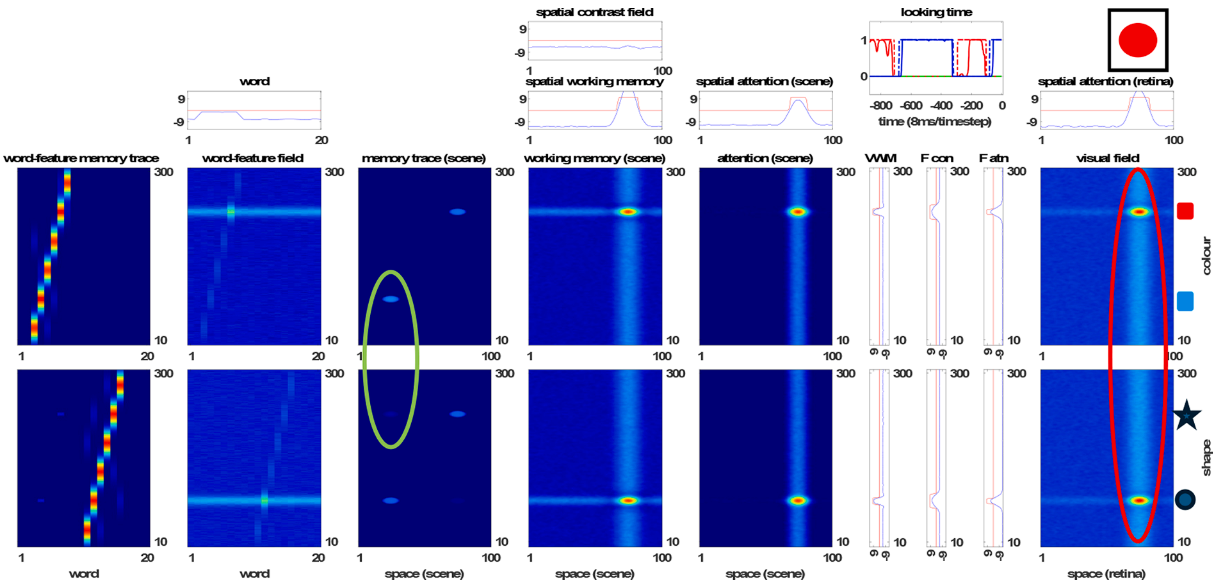

Fig. 9. The 'old' WOLVES 2.0 model during the demonstration phase of the DCCS task. The model is watching as the experimenter sorts a red circle to the right (see red oval in the visual fields). The green oval highlights the model's memory of seeing a blue star placed to the left by the experimenter on the previous demonstration.

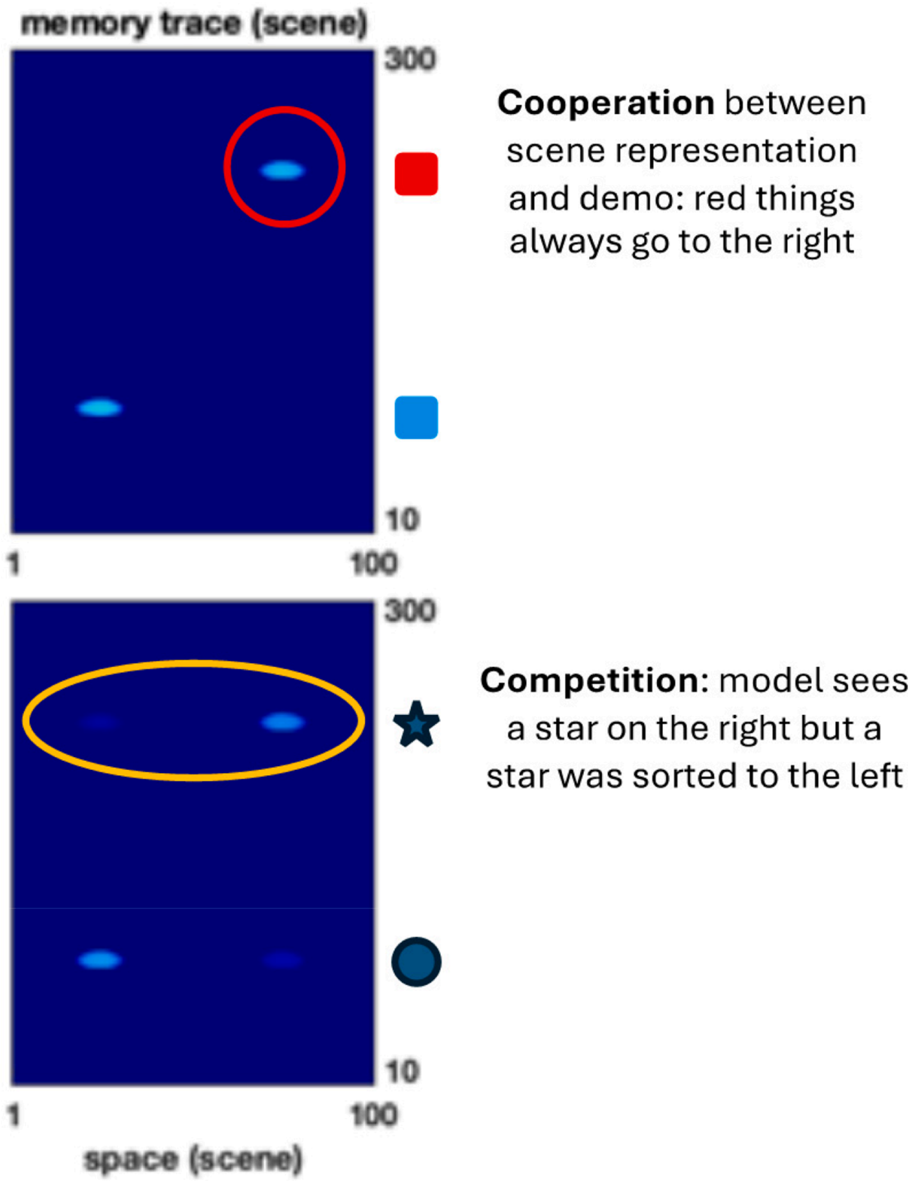

Fig. 10. The state of the model's scene memory trace following the instruction phase. The model has a cooperative pattern in the colour-space memory trace because it sees a red target card on the right and red things were always sorted to the right (see red oval). The model has a competitive pattern in the shape-space memory trace because it sees a star to the right but stars were sorted to the left (see orange oval).

Conceptually, we want the model to decide where to sort test cards based on its representation of word-feature mappings (sorting, for instance, 'red' items to the right). To enable this, we need the model's representation of words (captured in the left two columns of the model architecture) to pass activation to the scene working memory field. That way, word-related activation patterns can operate on the scene representation to influence decisions about whether a test card should be sorted to the left or the right. In our previous work with WOLVES, this type of interface with a scene representation was not needed in any of the tasks we simulated.

Fig. 11 shows WOLVES 2.0 – that is, WOLVES with one new connection added – on the first pre-switch trial. To ask the model to sort a test card, we implemented an approach described by Schneegans and colleagues in their model of visual working memory and change detection (Schneegans et al., 2016): we decoupled the retinal and scene reference frames. This allows the model to look at the test card at the centre of the retinal frame and think about where the card should go in the scene reference frame. We implemented this by setting the following connections to zero: (1) the input from retinal spatial attention to scene spatial attention, (2) the input from the visual field to spatial working memory, (3) the input from the visual field to the spatial contrast field, and (4) the input from the inhibition of return field to the scene spatial attention field. Thus, in Fig. 11, the model is looking at the blue star at the centre of the retinal frame and passing those features to the rest of the model along the feature pathway (see horizontal ridges). What we want is for

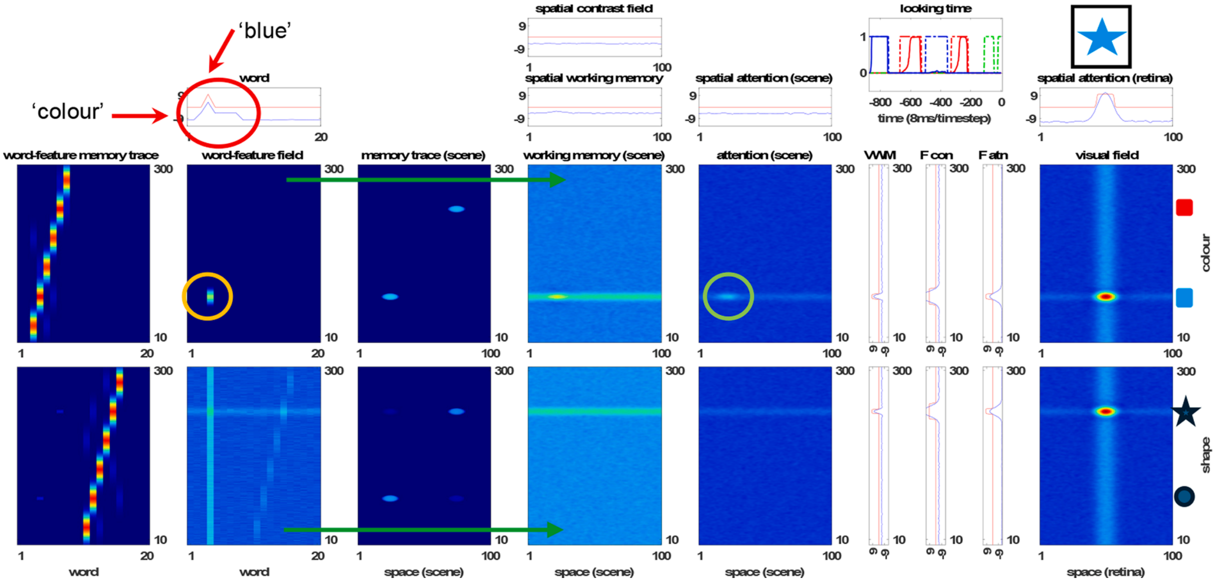

Fig. 11. The 'old' WOLVES 2.0 model sorting a blue star card to the left on the first pre-switch trial. The model sees the blue and star features presented on the test card in the centre of the retinal space. These feature inputs are passed into the word-feature fields where they match previously learned feature-labels mappings (see word-feature memory traces). Consequently, the model recalls that 'blue' matches the blue hue because it is playing the 'colour' game (see 'colour' boost in the red oval). The peak in the word-colour field (see orange oval) passes activation to the scene working memory field, boosting activation at the left location. Thus, the model decides to sort the blue star to the left (see green oval).

the model to decide where to sort the card in the scene reference frame and to do this by following the experimenter's instructions to sort by 'colour'. Thus, we have boosted all the colour labels in the word field (see red oval); we also gradually raise the resting levels of the word-feature and the word fields. This moves these fields closer to the activation threshold. As a consequence, the model has built a peak at the 'blue' colour label in the word field (see red oval) and a peak at the intersection of the blue hue value and the 'blue' label in the word-colour field (see orange oval). The word-colour peak passes activation to the scene working memory field which also receives input from the scene memory trace. This boosts activation at the left location in the scene working memory supported by the memory trace that blue things go to the left. This, in turn, increases activation at the left location in the scene attention field (see green oval). As the simulation continues to unfold, this bias toward the left grows stronger and the model ultimately builds a full pattern of peaks on

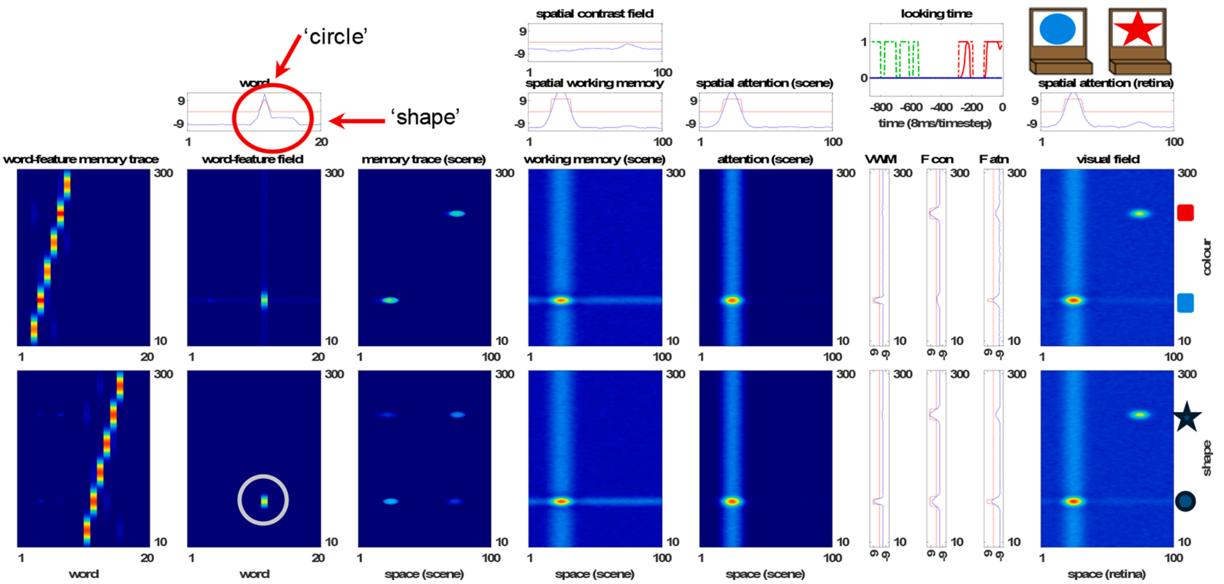

Fig. 12. The 'old' WOLVES 2.0 model on a post-switch instruction trial. The experimenter is pointing to the blue circle on the left and indicating that 'circles' go here in the 'shape' game (see red oval). As a consequence of these inputs, the model builds peak in the word-shape field at the intersection of the circle feature and the 'circle' label (see grey oval).

the left side of the scene working memory field and the left side of the scene attention field, thereby deciding to sort the blue star to the left. Related patterns of activation continue across all 6 pre-switch trials, with the model correctly sorting red circle test cards to the right and blue star test cards to the left.

Next, we tell the model it's time to switch and play the 'shape' game. This requires another instruction phase as shown in Fig. 12. Here, the experimenter is pointing to the blue circle target card on the left (see visual field), while instructing the model to play the 'shape' game in which 'circles' go to the left. The 'shape' label boosts the set of associated shape names the model learned previously (see word-feature memory trace), and the 'circle' input boosts this label in particular (see peak in the word field highlighted in the red oval). As a consequence of these inputs, the model builds a peak in the word-shape field at the intersection of the 'circle' label and the circle feature value (see grey oval).

The final phase of the task involves the 6 post-switch trials. The first of these trials is shown in Fig. 13. Here, the model is given a red circle to sort in the retinal frame (see visual field), the retinal and scene reference frames are spatially decoupled, and the model is told to play the 'shape' game (see flat boost to the shape labels in the red oval). As in the pre-switch trials, the visual features of the test card are passed across the feature pathway (see horizontal ridges in Fig. 13) while the resting levels of the word-feature and word fields are boosted. The model eventually recognises that the circle feature maps to the 'circle' label and it builds a 'circle' peak in the word field (see red oval) and a peak at the intersection of the 'circle' label and circle feature value in the word-shape field (see orange oval). This peak then passes activation into the scene working memory field, boosting the activation of the circle feature. When combined with the scene memory trace input, this builds a peak at the left location in both the scene working memory field and the scene attention field (see green oval). This bias is sufficient to build a full pattern of peaks at the left location in scene working memory and scene attention, sorting the red circle to the left. Thus, by attending to the 'shape' of the target card, the model overcomes the pre-potent response to sort the red circle to the right as it did during the pre-switch trials.

Can WOLVES 2.0 capture developmental changes in DCCS performance with 'weak' dimensional label learning? Next, we explored whether the model would show perseverative responding with 'weak' dimensional label learning as proposed by Buss and Spencer (2014). To implement weak dimensional label learning in WOLVES 2.0, we implemented two changes. First, we decreased the strength of the 'colour' and 'shape' boosts to the word field. This implements the idea that 'colour' more weakly pre-activates the set of colour labels due to a weaker association among these labels. Similarly, 'shape' more weakly pre-activates the set of shape labels, again, due to a weaker association among these labels. Conceptually, this reflects that early in learning, dimensional labels such as 'colour' are only weakly associated with specific feature labels such as 'red' or 'blue' and that these associations strengthen as learning progresses. As noted above, we do not implement the learning of these associations here; that will be a topic for future work. For now, our goal was to examine whether 'weak' dimensional label learning was sufficient to capture developmental changes in DCCS performance.

The second change we made in the model was to weaken the learning of the individual colour and shape labels as well. Recall that in Fig. 8, we highlighted that the old DCCS model had a strong Gaussian memory trace anchored at the intersection of the 'red' label and the red hue value in the word-colour field. We also included very weak memory traces for erroneous label-feature mappings. For the young model, we decreased the strength of the correct label-feature mappings and increased the strength of the incorrect label-

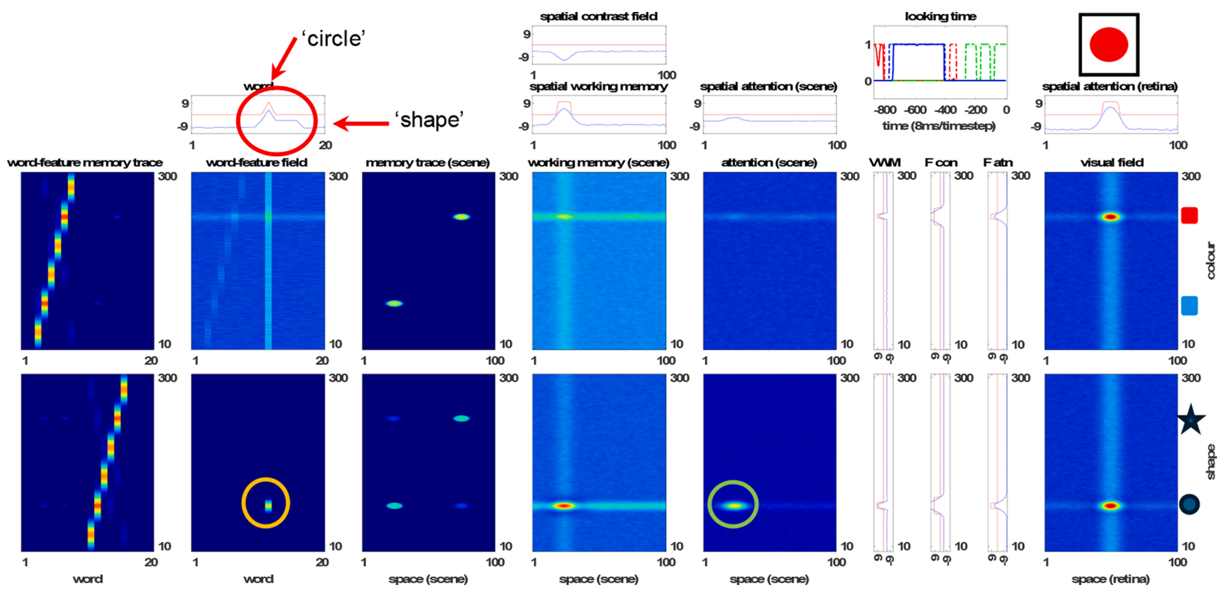

Fig. 13. The 'old' WOLVES 2.0 model correctly sorting a red circle to the left on the first post-switch trial. The model sees the red circle in the centre of the retinal space. The circle feature builds a peak in the word-shape field (see orange oval) which then activates the 'circle' label in the word field which is pre-activated by the 'shape' boost (see red oval). Once the model recalls 'circle', this boosts the activation from the word-shape peak sending stronger input into the scene working memory field. Consequently, the model decides to sort the red circle to the left (see green oval).

feature mappings. Conceptually, this reflects that early in development, specific feature labels such as ‘red’ or ‘blue’ are only weakly associated with the correct hue and that the messiness of learning these labels across different learning opportunities are likely to build erroneous, albeit weak, associations between labels and the wrong features. The result is shown in the word-colour memory trace (see orange oval) and the word-shape memory trace (see green oval) in Fig. 14. Note that the diagonal pattern – the correct label-feature mappings – is still evident, but there is strong competition for some label-feature combinations reflecting mis-matches that this particular individual model had previously learned. Again, we highlight that we didn’t train this model to create these long-term representations, but we know from Bhat et al. (2022) that WOLVES is quite good at learning individual word-feature mappings across situations by tracking cross-situational statistics and, like children, the model can sometimes make incorrect mappings. Thus, we are confident the model can learn individual colour and shape labels by simply repeatedly pairing these words with a variety of objects.

What is the consequence of weak dimensional label learning and weaker label-feature memory traces when we put this ‘young’ version of WOLVES 2.0 in the DCCS task? There are no substantive differences in performance initially: the ‘young’ WOLVES model visually explores the target cards, building a scene representation that blue stars go to the left and red circles go to the right.

Performance on the pre-switch trials is also comparable to the ‘old’ model. Fig. 14 shows the model sorting a red circle card to the right on the first pre-switch trial. As before, the model first builds a peak in the word-feature field at the intersection of the red hue value and the memory trace in the word-colour field (see grey oval). This peak in the word-colour field then passes activation to the word layer, activating the’red’ label (see red oval) which had been pre-activated by the ‘colour’ boost. The peak in the word field then amplifies the peak in the word-colour field. This sends a horizontal ridge into the scene working memory field, boosting this field closer to threshold at the red hue value. This boost intersects with the scene memory trace that red things get sorted to the right, leading to the emergence of a peak in the scene working memory field which then activates the associated site in scene attention (see white oval). Consequently, the model sorts the red circle to the right. The remaining pre-switch trials continue in a similar manner as the model sorts red circles to the right and blue stars to the left.

Although the pre-switch phase for the ‘young’ model is similar to how the ‘old’ model sorts cards, there are some differences. For instance, in Fig. 15, we show a simulation of a ‘young’ model on the first pre-switch trial as before. Note, however, that this model has some confusion about the mapping between the blue hue value and the label ‘blue’ (see orange oval). Instead of just a single memory trace, the model has two memory traces for this hue value – one which maps the blue hue to ‘blue’ and one which maps the blue hue to ‘purple’. During the pre-switch trial, the model actually builds a peak in the word-colour field at the intersection of the blue hue value and the ‘purple’ label (see grey oval). This drives the model to recall ‘purple’ in the word field (see red oval), which amplifies the peak in the word-colour field as before. This then passes activation to the scene working memory field which builds a peak at the intersection of the blue hue value and the left location. Thus, the model correctly sorts the blue star to the left (see white oval), even though it recalled the wrong colour label. Errors like this will be less and less frequent as the hue-label mappings refine over learning.

What about during the post-switch trials: does the ‘young’ model perseverate, and if so, why? Fig. 16A shows the first post-switch trial for the model from Fig. 14. As with the ‘old’ model, the ‘young’ model gets a boost for all the shape labels (see red oval). Nevertheless, because this boost is weak (due to a weak association between ‘shape’ and the individual shape labels), the model builds

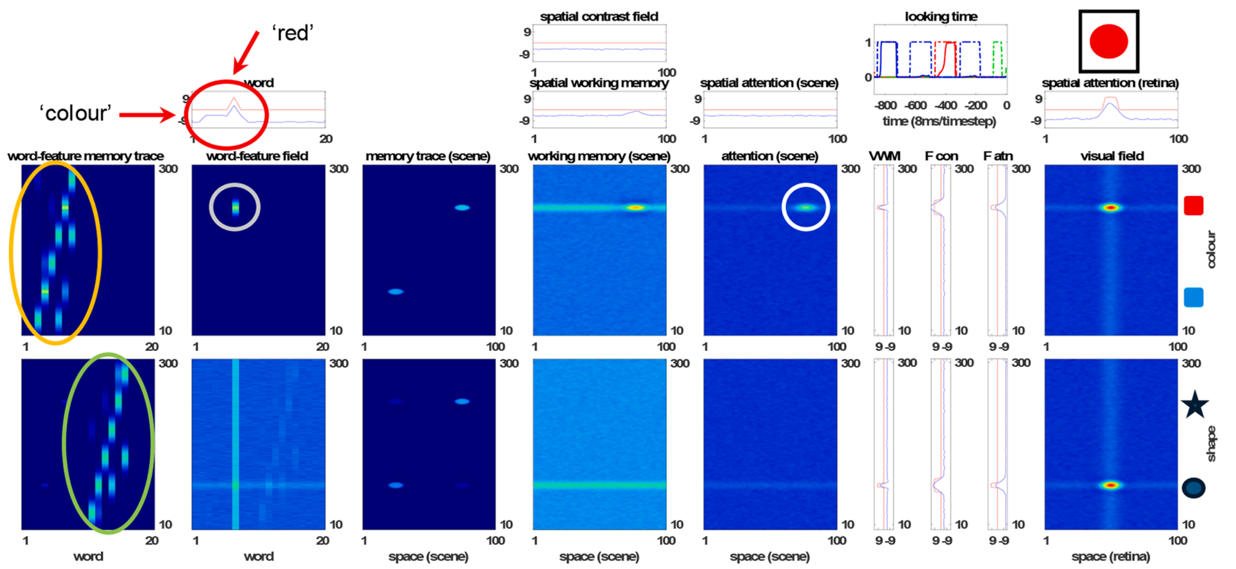

Fig. 14. The ‘young’ WOLVES 2.0 model sorting a red circle to the right location on the first pre-switch trial. The model sees the red circle at the centre of the retinal frame. These features build a peak in the word-colour field at the red hue value and the ‘red’ label (see grey oval). This peak helps the model recall ‘red’ when the ‘colour’ labels are preactivated (see red oval). The ‘red’ peak amplifies the peak in the word-colour field which sends activation into the scene working memory field. This leads to the emergence of a peak at the right location in the scene attention field (see white oval). Note the weaker word-feature mappings for this young model for both the colour labels (orange oval) and the shape labels (green oval).

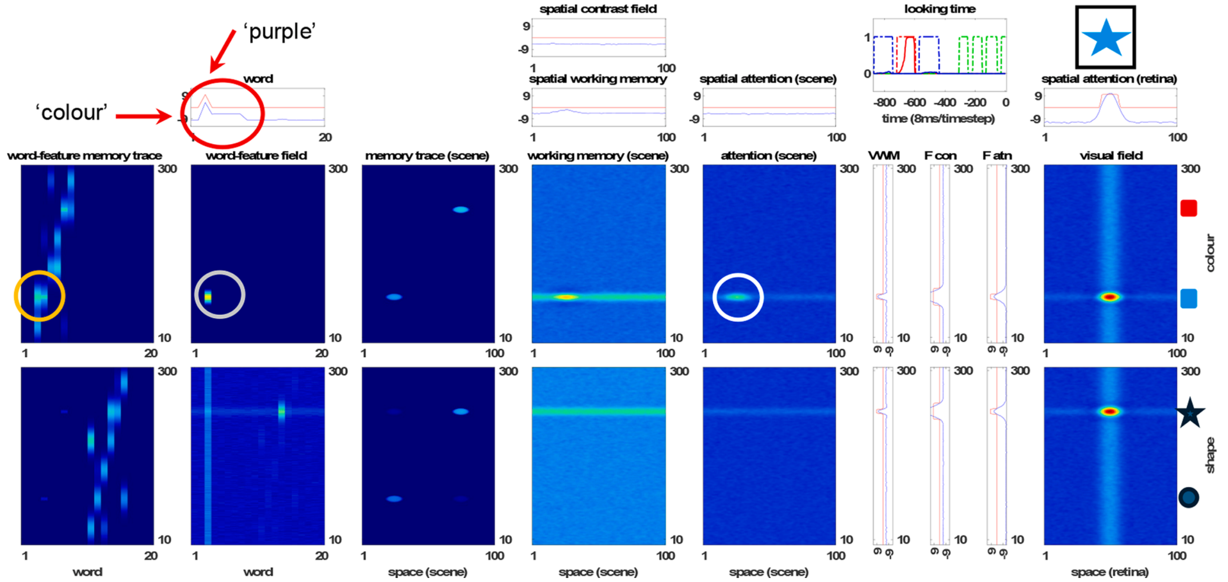

Fig. 15. The 'young' WOLVES 2.0 model correctly sorting a blue star to the left (see white oval) on a pre-switch trial even though it erroneously recalls the label 'purple' (see red oval). The error occurs because of a mis-matching word-colour memory trace (see orange oval), that is, the model has yet to figure out if the blue hue value maps to 'blue' or 'purple'. In this case, the model maps the blue hue to the 'purple' label (see grey oval) but still sorts the card correctly.

a peak in the word-colour field (see orange oval) at the blue hue value *before* building a peak for one of the shape labels. This 'blue' peak in the word-colour field is relatively strong given that the model sorted blue cards on 3 of the 6 pre-switch trials. Consequently, this peak passes activation to the scene working memory which intersects with a strong scene memory trace that blue things go to the left. This combination of inputs is sufficient to cause the model to sort the blue star to the left, perseveratively sorting by colour. These dynamics are consistent with behavioural data showing that children who fail the DCCS task show weaker implicit priming of dimensional attention (Benitez, Vales, Hanania, & Smith, 2017). We have used earlier instantiations of the DCCS model to demonstrate that such differences in priming also arise from the weak dimensional label representations that are associated with perseveration in the DCCS task (Buss & Kerr-German, 2019).

Interestingly, Fig. 16B shows this same model a few time steps later. Note that the model has decided to sort the blue star to the left, but it has recalled that the shape on the test card is a 'star'. Thus, the model is thinking 'star' but sorted the card by its hue value. This mis-match in behaviour is reminiscent of some behaviours young children show in the DCCS task. For instance, if asked where 'stars' go in the 'shape' game, children will often correctly point to the right location. If they are then given a blue star to sort, however, they will sort it to the left, again, perseveratively sorting by colour. The difference between these situations in the model reflects the sequence of events in these two cases. In the first, the labels are provided (i.e., 'star' and 'shape') and then the child is asked to point to the star. The 'young' model has no problem doing this as 'star' sends a vertical ridge into the word-shape field at the same time as there is a horizontal ridge indicating that a matching feature is in the task space. By contrast, in the second case on the DCCS trial, the model (and child) must decide where to sort the blue star card based on its memory of past sorting decisions and its internal ability to recall the 'blue' and 'star' labels in context.

In summary, WOLVES 2.0 does show perseverative responding in the DCCS task when we implement the developmental proposal from Buss and Spencer (2014) with young children showing weak dimensional label learning and older children showing strong dimensional label learning. We note as in our prior work, that sorting behaviour in this model is an emergent property of many factors coming together. For instance, the details of the scene representation matter, how the model visually explores the target cards matter, the labels used in the task matter, and so on. Critically, these demonstrations that WOLVES can show developmentally appropriate behaviours embeds our account in a richer neural process model that specifies how many 'component' processes like attention and working memory work together to create sorting decisions in the DCCS task.

Do changes in the strength of dimensional label learning yield differences in comprehension and production tasks? Our next question was whether WOLVES 2.0 would show appropriate behaviours in both the comprehension and production tasks. We begin by focusing on the comprehension task. This task is similar to many of the original behaviours studied using the WOLVES model. In many cross-situational word learning tasks (CSWL), there is a learning phase followed by a test phase. In the test phase, a single word is presented and the participant must pick the correct referent from a set of multiple items, that is, the model must show comprehension of the word.

Fig. 17 shows WOLVES 2.0 embedded in this type of task. Because we are using familiar / previously-learned labels, we did not implement a learning phase; instead, we move directly to comprehension test trials. In Fig. 17A, we present the model with the word 'purple'. This word input builds a peak in the word field (see red oval) which then projects a vertical ridge into the word-feature fields,

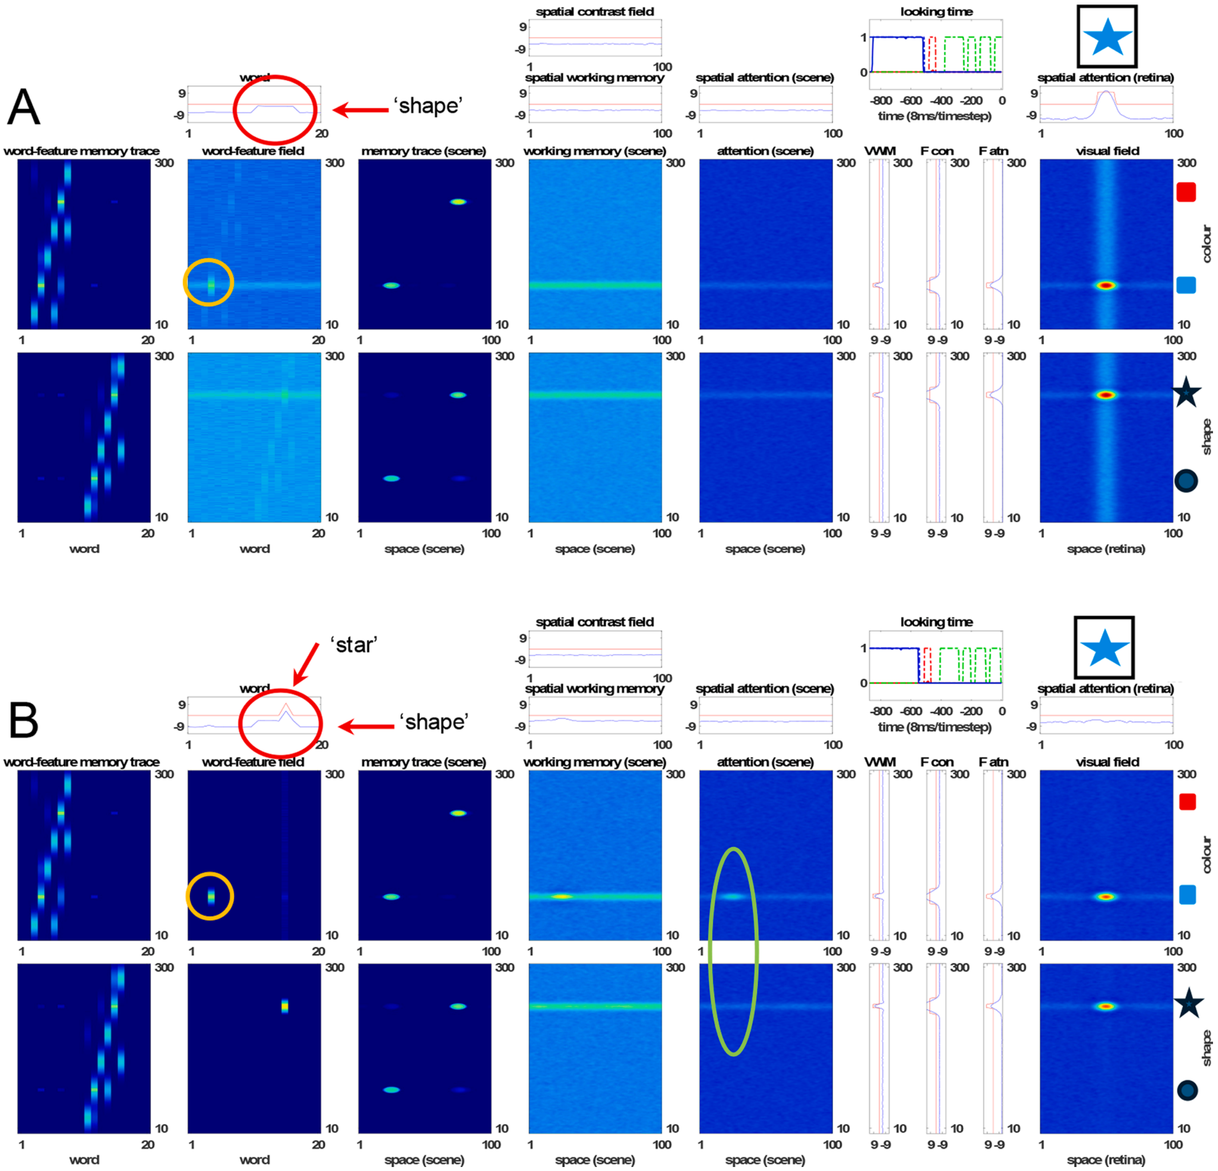

Fig. 16. The 'young' WOLVES 2.0 model making a perseverative error on the first post-switch trial. A) The model see a blue star at the centre of the retinal frame and recalls that the blue hue value mapped to the 'blue' label on previous trials (see orange oval). This sends a boost to the scene working memory, creating a bias to sort the card to the left. Note that the model builds a peak in the word-colour field even though it has boosted the 'shape' labels in the word field (see red oval). B) A few time steps later, the model correctly identifies that the test card is a 'circle' (see red oval); however, it has already decided to sort the card to the left in the scene attention field (see green oval). Thus, the model makes a perseverative error.

building a weak peak at the recalled hue value (i.e., the centre of the purple category; see orange oval). Next, we present a set of 6 objects as in experiments by Buss and colleagues (Lowery et al., 2022). As can be seen in the green oval, the model quickly builds a peak at the purple hue value in the feature contrast field, highlighting this feature value. This reflects the top-down input from the word-feature fields to the feature contrast fields. Fig. 17B shows the model a few timesteps later. Now the peak in the word-colour field has sharpened (see orange oval) and a peak has emerged at the purple hue value in the feature attention field (see green oval). Fig. 17C shows the model a few time steps later. Now there is a robust peak at the purple hue value in the feature attention field as well as in the visual field (see green oval). This then passes activation to the retinal spatial attention field (see black oval) and the model binds the hue and shape features of the same object together, identifying the object at location 75 as the 'purple' item.

As expected, the 'young' model makes more errors in this comprehension task. Fig. 18 shows one example. Here, we give the model the label 'green' (see red oval). The young model successfully recalls the correct hue value (see orange oval) and attends to this feature value (see green oval). However, it also simultaneously attends to one of the shape features (see black oval). Because the 'green' peak in the word-colour field is relatively weak (due to a weaker memory trace for this mapping), this simulation ends up picking based on the

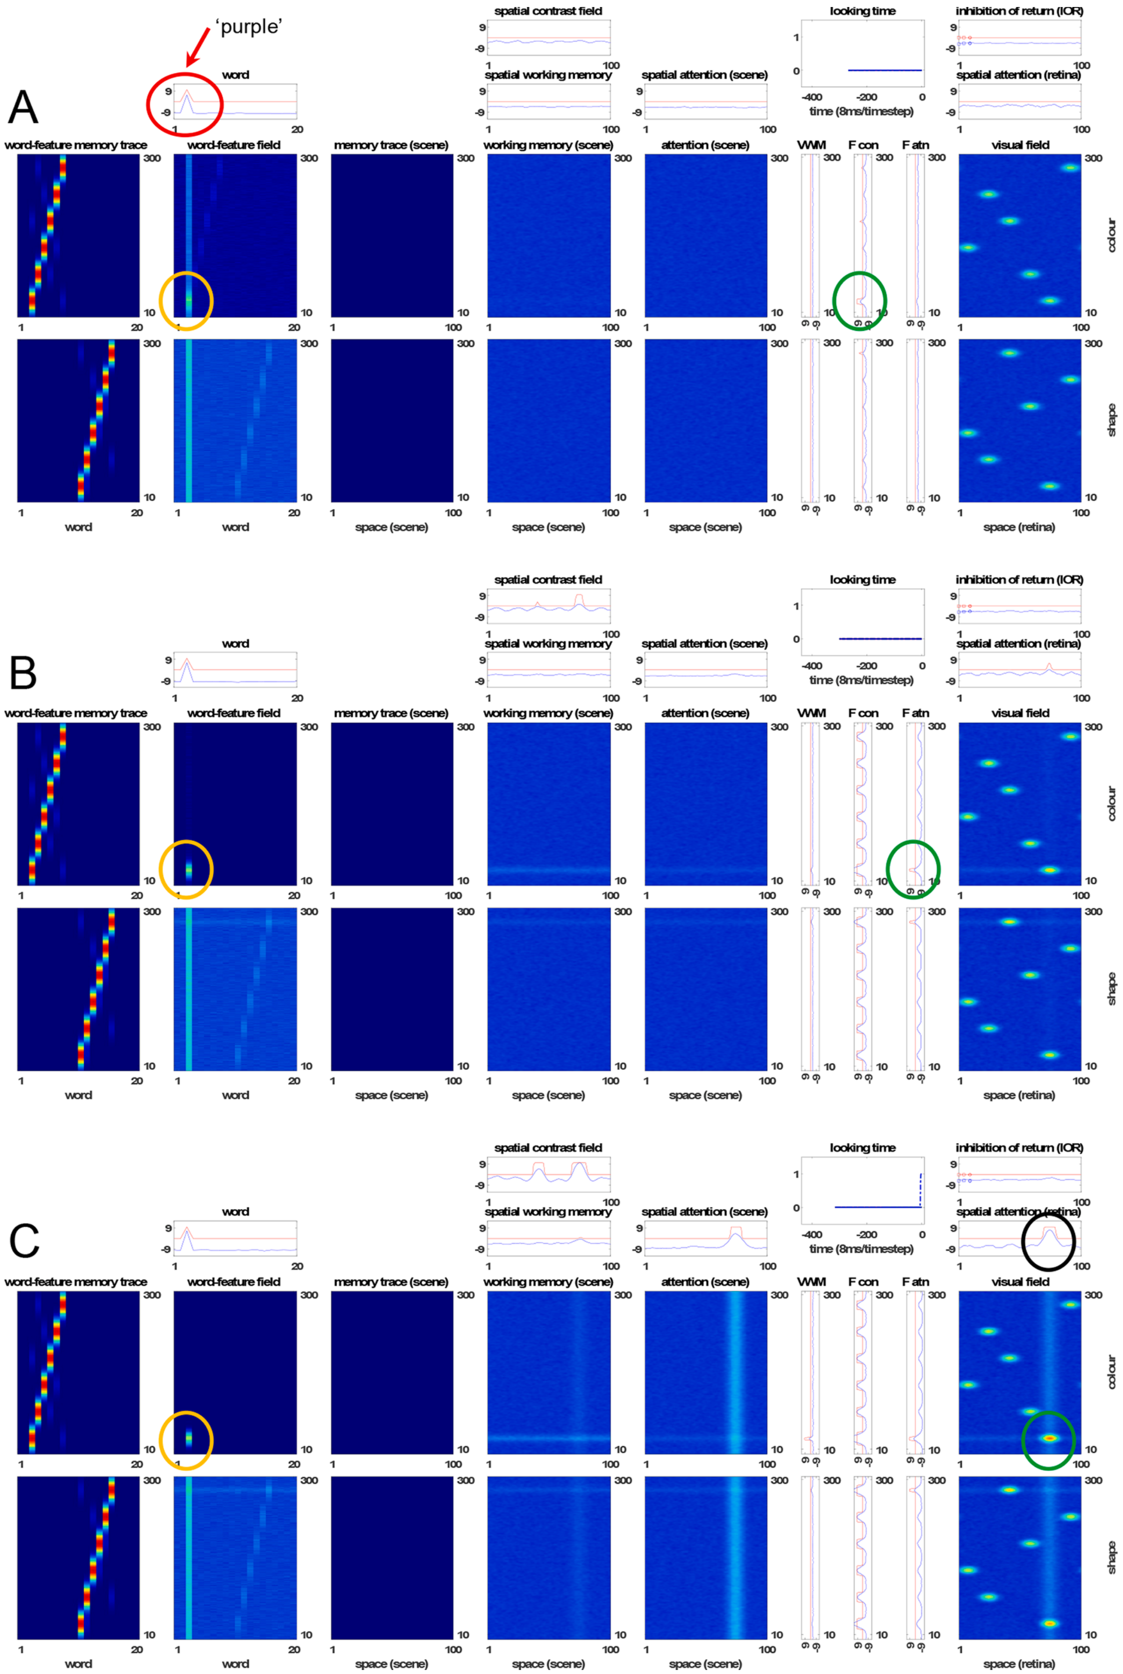

(caption on next page)

Fig. 17. The 'old' WOLVES 2.0 model selecting the 'purple' item correctly on a comprehension trial. A) At the start of the trial, the model is told to find the 'purple' item (see red oval). This primes the purple hue value in the word-colour field (see orange oval) and boosts the activity of this hue value in the feature contrast field (see green oval). B) A few time steps later, the model has sharpened the peak in the word-colour field (see orange oval) and built a peak at the purple hue value in the feature attention field (see green oval). C) A few time steps later, the model has spatially attended to the purple item in the task space (see black oval), correctly selecting this as the 'purple' object (see green oval).

shape feature which happens to be the blue item. In summary, then, the 'young' model makes more errors in comprehension than the 'old' model due to weaker memory traces for, in this case, the hue-label mappings.

What about in the production task? Fig. 19 shows a simulation of the 'old' model in this task. The model is presented with an object to look at in the visual field and it is asked to name the 'colour' of the object (see 'colour' boost in the red oval). As in the DCCS task, the horizontal feature ridge generated as the model attends to the object features intersect with the word-feature memory traces. This produces a peak in the word-colour field (see orange oval) which then creates a peak in the word field at the correct label (in this case, 'orange'), much like the dynamics in the DCCS task where the model recalls the label for the feature of the object it is looking at. Thus, WOLVES 2.0 can readily perform both comprehension and production tasks (for related findings, see (Samuelson, Spencer, & Jenkins, 2013).

Can we quantitatively simulate children's performance in the DCCS, comprehension and production tasks, including changes over development? The qualitative simulations above demonstrate that WOLVES 2.0 can capture the right types of behaviours in the DCCS, comprehension, and production tasks. Here, we push this modelling work to a quantitative level of detail. Our goal was to simulate quantitative data from 3- to 5-year-olds using a set of 'young' and 'old' parameters. Ideally, the 'young' and 'old' parameters would only differ in the strength of the dimensional label boost and the strength of the word-feature mappings for the individual colour and shape labels. We hoped to keep other parameter changes to a minimum, acknowledging that some task-specific modulation of parameters might be necessary.

To establish the initial model parameters, we ran simulations using a graphical user interface (as shown in the model figures) and systematically varied parameters to move the model's performance toward the target behaviours. This work established a good value for the strength of input from the word-feature fields to the scene working memory fields (1.1) – the one architectural change in WOLVES 2.0 – and the values for the dimensional label boost for the 'young' (5.0) and 'old' (5.4) model. We also set the base strength for the word-feature memory traces for the'young' (0.375) and 'old' (0.8) models. This reflects the strength of the memory trace for each learned label (e.g., 'red' and 'circle'). We also modulated the array of competing memory traces, that is, the strength of erroneous label-feature mappings (e.g., mapping 'blue' to the purple hue value). For the 'young' model, the 6 competing memory traces varied in strengths of 0.1, 0.1, 0.2, 0.3, 0.4, and 0.45, with these values randomly assigned to specific label-feature mappings. For the 'old' model, the 6 competing memory traces all had a strength of 0.1.

During the interactive modelling step, we also identified 8 parameter changes to existing connections in WOLVES were needed relative to the parameters from Bhat et al. (2022; see Table 1, section 2): (1) we corrected one parameter from WOLVES, setting the input from the word-feature fields into the word field to have a width of 0; (2) we increased the input from scene working memory to

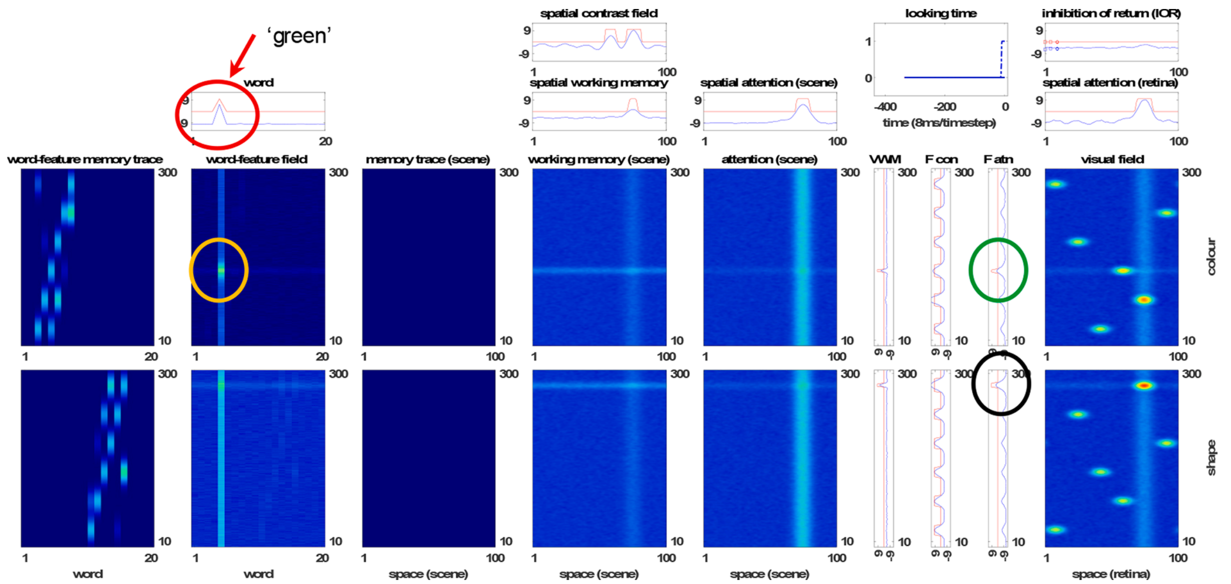

Fig. 18. The 'young' WOLVES 2.0 model making an error on a comprehension trial. The model correctly activates the 'green' label, boosting the green hue value in the word-colour field (see orange oval). However, because the model has a weak word-colour memory trace, this only weakly boosts the green hue value in the feature attention field (see green oval). This peak competes with an erroneously attended shape feature (see black oval) and the model selects the wrong object.

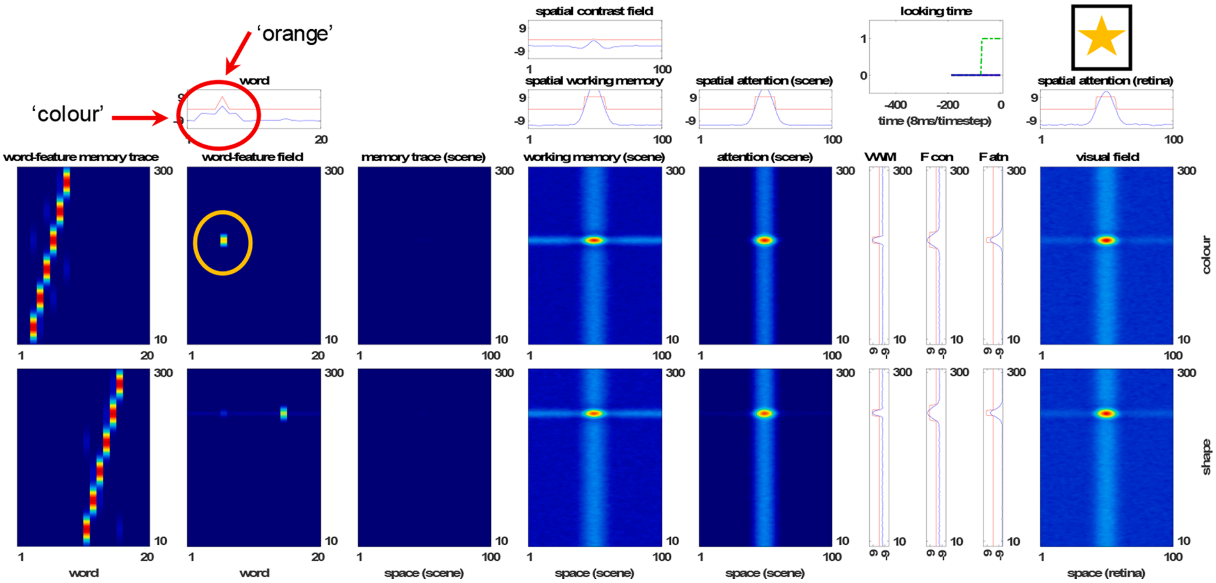

Fig. 19. The 'old' WOLVES 2.0 model correctly producing the label 'orange' in the production task. The model is shown an orange star and is asked to name the 'colour' of the object (see red oval). The features of the object are passed into the word-feature fields, building a peak at the orange hue value (see orange oval). This boosts activation in the word field and the model correctly activates the 'orange' label.

Table 1
Table of parameter changes for WOLVES 2.0. See text for details.

Parameters Created for the New Tasks (WOLVES 2.0)			
Element		**Sigma (σ)**	**Amplitude (a)**
word-feature field to scene WM (wf → wm_c)		5	1.1
Parameters Changed for All Tasks			
Element	**Component**	**Old Value**	**New Value**
word-feature to word field (wf → word)	σ_{exc}	5	0
scene WM to SWM (wm_c → wm_s)	a_{exc}	0.3	0.6
feature attention to word-feature field (atn_f → wf)	a_{exc}	1	1.4
word memory trace to word field (hword → word)	a_{exc}	0.1	0.2
scene memory trace to scene WM (hwm_c → wm_c)	a_{exc}	1	0.4
VWM to VWM (wm_f → wm_f)	a_{exc}	24	22
scene WM to scene WM (wm_c → wm_c)	a_{glob}	0	−0.01
scene WM to scene attention (wm_c → atn_c)	a_{exc}	5.5	9.5
Parameters Changed for DCCS, Production, and Comprehension			
Element	**Component**	**Old Value**	**New Value**
word-feature memory trace to word-feature field (hwf → wf)	a_{exc}	4	1
word-feature field to word field (wf → word)	a_{exc}	0.05	0.2
word field to word-feature field (word → wf)*	a_{exc}	4.25	2.2

*This parameter was not modified for the Comprehension task.
Note that parameter names in parentheses reflect the names used in the simulator, which main name refers to the labels used in the figures. Duplicate fields for colour and shape dimensions had identical parameters.

spatial working memory to ensure the model would spatially bind features in scene working memory given the strong feature conflict in the DCCS task; (3) we increased the strength of input from feature attention to the word-feature fields as we needed to drive words in a bottom-up way in the DCCS task (vs. the direct word input in cross-situational word learning); (4) we increased the input from the word memory trace to the word field to prevent erroneous 'other' words from becoming active too frequently; (5) we decreased the input from the scene memory trace to scene working memory to allow words to contribute to the 'what-where' mappings in scene working memory; (6) we decreased the self-excitation in the feature working memory fields as with the longer simulation durations we had a few cases where peaks in working memory sustained for a long time due to strong memory trace input; (7) we added some global inhibition to the scene working memory fields to make this field more selective and force decisions in the DCCS task; (8) we increased the input from scene working memory to scene attention to compensate for the reduced input from scene working memory caused by the increase in global inhibition from change 7.

Once we arrived at the parameter changes in Table 1, the model was performing the DCCS task relatively well for both ‘young’ and ‘old’ parameter sets. We then ran batches of simulations ($N = 100$) to quantitatively compare the model’s performance to children’s performance from the DCCS, Comprehension, and Production tasks. At this point, the model was doing well in the DCCS task and the Production task but was making too many errors in Comprehension. Given the overlap between Comprehension and cross-situational word learning (which requires comprehension of a learned word), we tried using the stronger input from the word field into the word-feature fields used in Bhat et al. (2022; from 2.2 to 4.25). This created a stronger top-down attentional bias from the word-feature fields to the feature contrast fields. With this change, the model performed well in the Comprehension task as well.

Table 1 (section 3) summarises the 3 parameter changes modulated across tasks. In particular, two parameters were changed for the DCCS, Production and Comprehension tasks relative to the original CSWL values: (1) the 3 new tasks used a weaker input from the word-feature memory trace into the word-feature fields to accommodate the stronger longer-term memory traces built into the model to reflect long-term learning of colour and shape labels; and (2) the 3 new tasks used stronger word-feature field input to the word field to enable peaks in the word-feature fields to drive the formation of label peaks in the word field. Finally, the input from the word to the word-feature fields was weaker in DCCS and Production (2.2) and stronger in Comprehension and CSWL (4.25).

As a final step, we generated variability in the values of parameters that are important for developmental changes in performance of the model. We used probability density functions to define a population of values for the strength of word-feature mappings and the strength of the dimensional label boost (see Table 2). These distributions of parameter values reflect individual differences in the learning of dimensional labels between 3 and 5 years of age. For each run of the model in each task, we randomly sampled a value from these distributions over batches of 100 simulations in each task for both the “young” and “old” model parameter distributions.

Fig. 20 shows the fit between the model results and the behaviour of children. Fig. 20A shows results from the DCCS task. The behaviour of the model was categorised in the same fashion as the behavioural data from children. Specifically, the model was deemed to have passed the pre-switch phase if it sorted at least 5 out of 6 cards correctly. During the post-switch phase, the model was categorised as passing if it sorted 5 out of 6 trials correctly and as failing if it sorted 1 or fewer trials correctly during the post-switch phase. Otherwise, the model was deemed to show mixed responding (i.e., sorting 2, 3, or 4 trials correctly during the post-switch). As can be seen in Fig. 20A, the “young” model showed a high rate of perseveration similar to 3-year-olds and the “old” model showed a high rate of correct switching similar to 4- and 5-year-olds. Note that empirical rates of switching/failing were taken from the rates summarised from the literature in Buss and Spencer (2014).

The fit between the model and children’s behaviour in the production and comprehension tasks is plotted in Fig. 20B. Note that the behavioural data are summarised from McCraw et al. (2024). The model does a good job capturing developmental changes in performance of children on these tasks. The model also captures the differences in performance between the comprehension and production tasks. Specifically, the “young” model performed better on the comprehension task relative to the production task similar to children. The model also makes patterns of errors that mimic that of children. Specifically, the model always generates a label from the correct dimension during the production task, similar to children. This is due to the dimensional boost to the relevant feature labels when cued with the label’s “shape” or “colour”. Thus, the model builds a peak for a label within the correct dimension, even if it is not the correct label.

Together, these results show that changes in the model’s dimensional label knowledge produce changes in behaviour on the DCCS task that mimic the developmental changes observed in children. Moreover, these changes in dimensional label knowledge produce changes in behaviour on comprehension and production tasks that mimic patterns shown by children. We note here that these modelling efforts are only a first step. This first step is useful in that it shows there is nothing in the model that prevents a good fit to the integrated data set. This was not guaranteed given the constraints provided by the DF model. Moreover, these are quite different tasks with different processing demands, so there was no guarantee a priori that the model would fit behaviour across these tasks. Success with this first step sets the stage for future work where we use the model to explain new data and make new predictions. One exciting possibility is to use this model as a platform to generate predictions about how to optimise the dimensional label learning process such that EF development can be facilitated. We return to these issues in the Discussion.

Can we capture other label- and feature-based manipulations that have been shown to impact DCCS performance? In the initial publication of the DF model that simulated performance and development on the DCCS task, Buss and Spencer (2014) pointed to one set of findings from Yerys and Munakata (2006) that was beyond the scope of what the initial model could capture. Yerys and Munakata (2006) used the partial partial-change (PPC) version of the DCCS (Zelazo et al., 2003) in which the test cards match the target cards only along the features that are relevant for the pre-switch rules. For example, if children are sorting by shape, then the target cards could be a red star and a blue circle and the test cards could be a green star and a yellow circle. During the post-switch phase, the test cards were changed so that they matched the target cards along opposite dimensions (in this example, a red circle and a

Table 2
Parameter values used to create distributions of ‘young’ and ‘old’ models.

Values used to define probability density functions for random parameter values.

		Mean	SD	Skew	Kurtosis
Word-Feature Mapping	Young	0.375	0.12	3.6	14
	Old	0.8	0.12	−3.6	14
Dimensional Label Boost	Young	5.0	0.15	3.6	14
	Old	5.4	0.12	−3.6	14

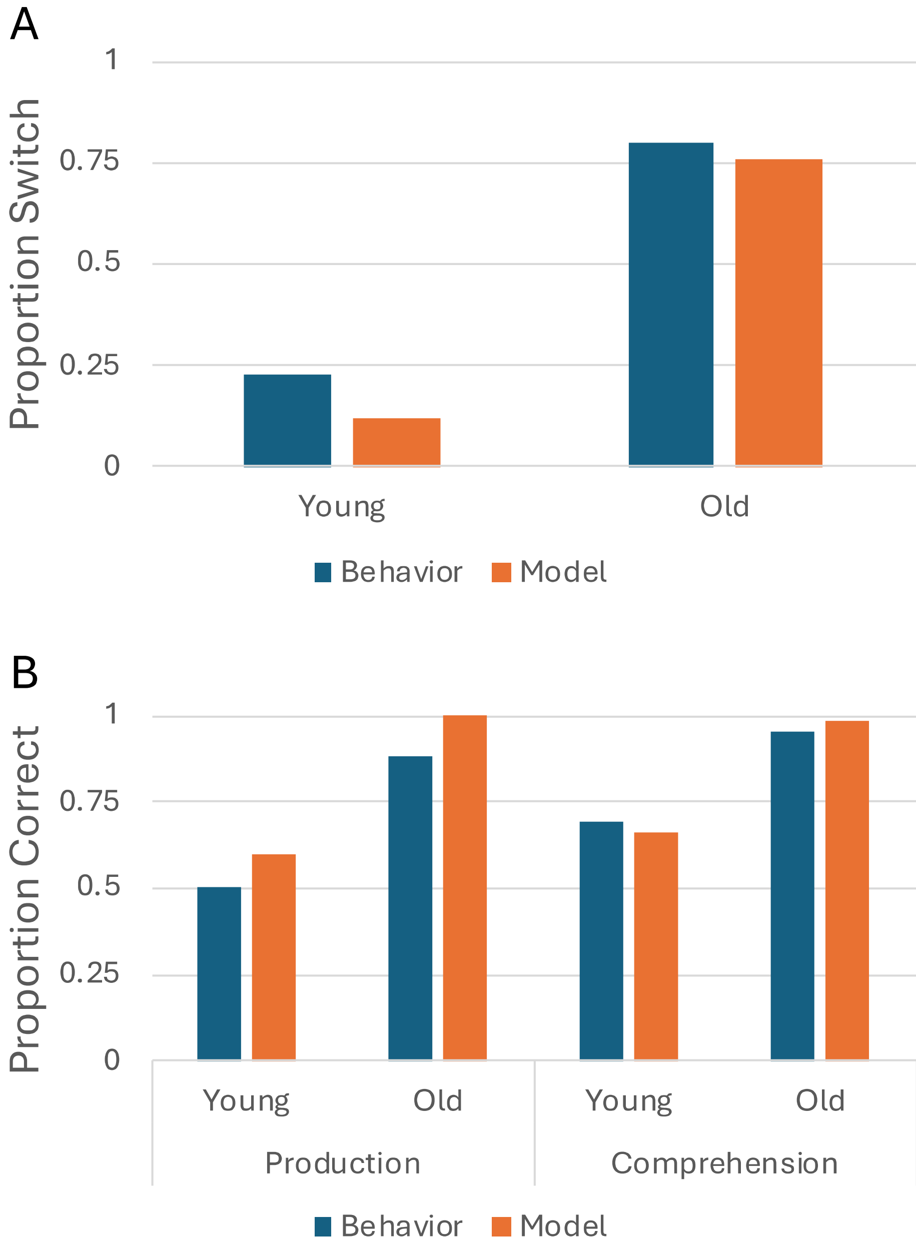

Fig. 20. Simulations of the 'young' and 'old' WOLVES 2.0 model compared to children's performance. Model results are from 100 simulations per condition. A) Switching behaviour in the DCCS task. B) Proportion correct responding in the production and comprehension tasks.

blue star).

Yerys and Munakata (2006) administered the basic PPC condition and compared performance across two other conditions. In an *uninformative* condition, dimensional and feature labels were not used during pre-switch instructions. Instead, the experimenter described the task as a "sorting game" and said "these go here" while showing a test card that matched the target card along the relevant dimension (e.g., holding a test card with a picture of a flower next to a target card with a picture of a flower). Standard instructions using labels for dimensions and features were administered during the post-switch phase. In a *novel feature* condition, the features that were relevant for the rules used during the pre-switch phase were unfamiliar to children and were given novel labels. For example, children would be told they are going to play the "shape" game, and that "daxes" go to the location with a rounded blob shape and "gubs" go to the location with a jagged blob shape. Yerys and Munakata (2006) found that children switched at a low rate in the basic PPC conditions but switched rules at a significantly higher rate in both the uninformative and novel feature conditions.

We implemented these conditions in the model without any changes to the model itself. The PPC condition is straightforward to implement by changing the features on the pre-switch test cards. This was done across all three conditions as administered by Yerys and Munakata (2006). To implement the uninformative condition, we simply did not provide any label inputs during the pre-switch instructions. To implement the novel feature condition, we shifted the features of the pre-switch dimension by 10 neural units so that they did not overlap with the feature-label associations that were built into the word-feature field. We also instructed the model using the typical dimensional boost, but we used completely novel labels (i.e., activating a word for which the model had no prior learning). We ran batches of the basic PPC task, the uninformative task, and the novel feature task as described above. Initial simulations of the novel feature task resulted in low pre-switch performance because the model did not build a response during the trial interval (e.g., in our initial test of this condition, only 68% of the models passed the pre-switch phase). This makes sense because the novel feature and label inputs do not overlap with the model's knowledge of features and dimensions. We re-ran this condition and increased the trial length from 7,000 timesteps (as was the case for every other DCCS simulation) to 10,000 timesteps. This boosted performance in the pre-switch phase considerably (see Fig. 21).

Simulation results are shown in Fig. 21 next to empirical results from Yerys and Munakata (2006). Across conditions, the model was impacted by the task manipulations in a similar fashion as children. The young model switched at a very low rate during the basic PPC task but improved in both the uninformative and novel feature conditions. The model data quantitatively match the empirical observations quite well, showing good generalisations to these novel task conditions.

Does WOLVES still capture quantitative changes in cross-situational word learning? Given that we successfully simulated findings from the DCCS, Comprehension, and Production tasks, we next examined if WOLVES 2.0 could still quantitatively simulate the original findings from the cross-situational word learning (CSWL) domain. To probe this question, we first ran individual simulations to examine how WOLVES behaved in the CSWL task with the DCCS parameters. Generally, the model behaved reasonably; however, we had to reset the 3 parameters shown in Table 1 (section 3) back to their original values. With this reset, WOLVES 2.0 successfully simulated key findings from the original set of CSWL tasks. Thus, the 9 other parameter changes needed to simulate the DCCS, Production, and Comprehension tasks (see Table 1, sections 1 and 2) proved to not impact CSWL performance substantially.

For instance, Fig. 22A shows the model performing the canonical CSWL paradigm from Smith and Yu (2008; see also Yu & Smith, 2011). In this paradigm, two objects are presented on each training trial along with two novel words. On any given trial, it is ambiguous which object each word refers to; however, this can be determined across trials as the same word-object pairings are repeated over trials. After 30 training trials, participants are shown two objects and hear a single word. As can be seen in Fig. 22A, toddlers look longer to the target item and less often at a distractor, showing robust learning. WOLVES 2.0 quantitatively simulates these behaviours. Fig. 22B shows looking times on each trial. Again, WOLVES 2.0 does a good job reproducing an appropriate looking time. Note that WOLVES 2.0 was anchored to the performance of 3- to 5-year-olds above, so we might expect shorter looking times for WOLVES 2.0 as the model is 'older' than the children simulated (and looking times generally decrease in this task over development as speed of processing increases). Consistent with this, WOLVES 2.0 learned on average 4.91 words out of 6, while toddlers and the original WOLVES model with toddler parameters learned 3.5–4.0 words (see Bhat et al., 2022).

Fig. 22C shows simulations of a second example CSWL task from Vlach and DeBrock (2019). In this task, 47- to 58-month-old children were presented with 12 word-object mappings that were either grouped together (Massed Condition) or distributed (Interleaved Condition). Results from a prior study with younger children showed better learning in the massed condition (Vlach & Johnson, 2013), while older children showed better learning in the Interleaved condition. WOLVES captured this developmental transition (see Bhat et al., 2022). In Fig. 22C, we show that WOLVES 2.0 with 'old' parameters also shows better learning in the interleaved condition.

A final example simulated a study by Suanda et al. (2014). This study investigated the role that contextual diversity – defined as the degree to which multiple word-object mappings tend to co-occur – plays in 5- to 7-year-old children's cross-situational word learning. The hypothesis was that if children learn word-object mappings by tracking the co-occurrences of words and objects, they should be

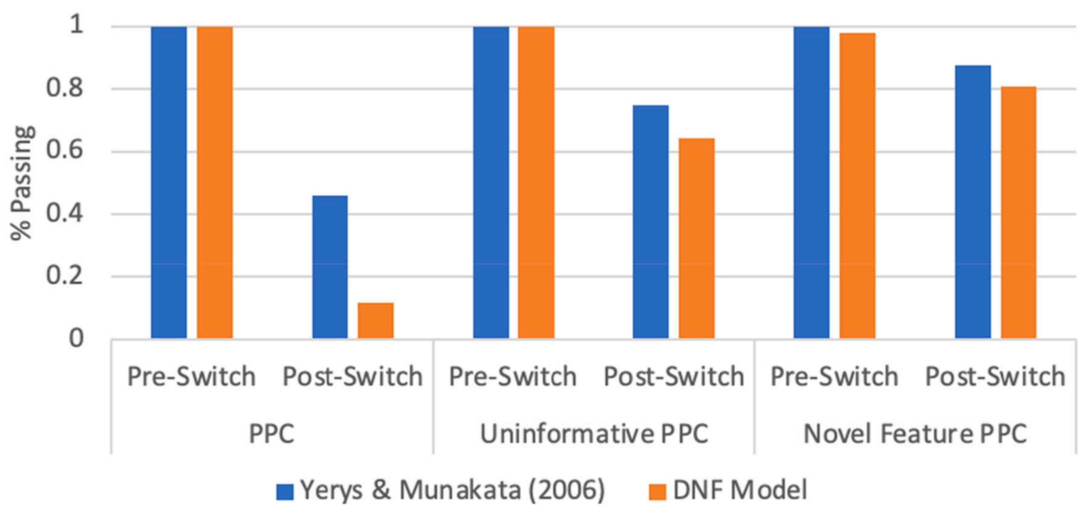

Fig. 21. Performance of the young WOLVES 2.0 model compared to children's performance across the three conditions described by Yerys & Munakata (2006). Results are from 100 runs per condition. PPC = Partial-partial change.

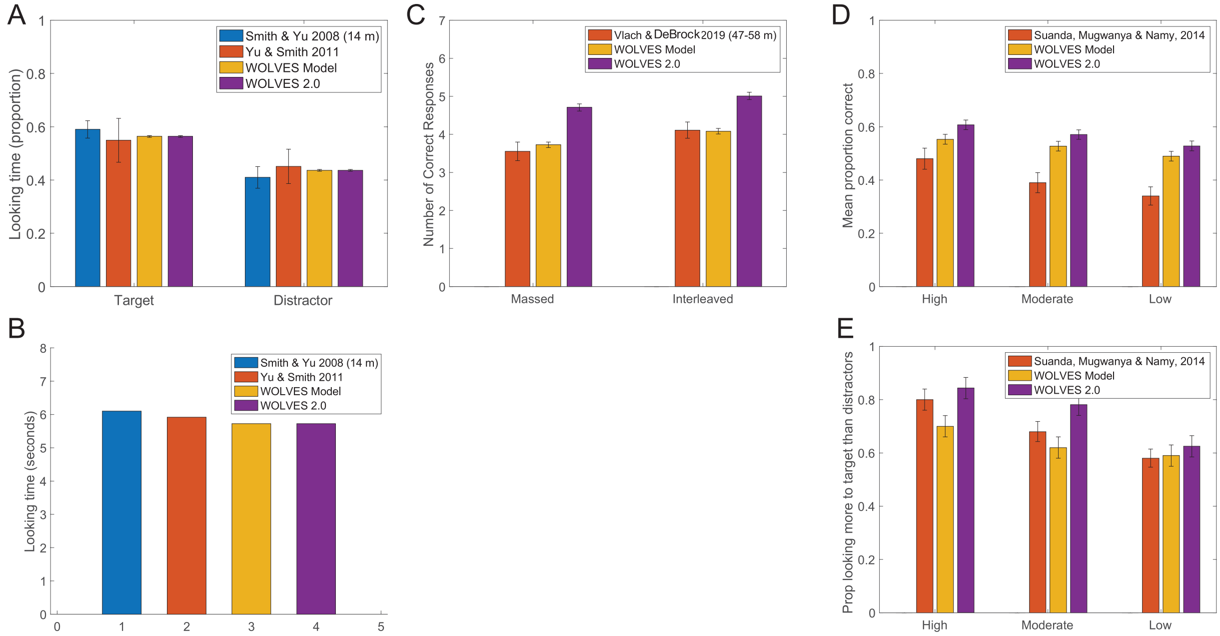

Fig. 22. The WOLVES 2.0 model simulating findings from the cross-situational word learning literature. Simulations of the CSWL task from Smith & Yu (2008) and Yu & Smith (2011) showing (A) mean looking time in seconds and (B) proportion looking to the target vs. the distractor. The yellow bars show the original WOLVES model performance and the purple bars show WOLVES 2.0. C) Simulations of the task by Vlach and Johnson (2018). D) and (E) show simulations of the task by Suanda et al. (2014) for the mean proportion correct and the proportion of models looking more to the target than the distractors. See text for additional details.

less successful in situations with lower contextual diversity where there are higher cross-correlations between words and objects, than in situations with higher contextual diversity. To examine this, Suanda et al. (2014) presented children eight word-object mappings to learn in conditions of either high, medium, or low contextual diversity. For example, in the High condition, Word 1 (W1) co-occurred with its referent (P1) on all four trials in which it occurred. W1–P1 was accompanied by W2–P2 on one of those trials, W3–P3 on a different trial, W4–P4 on another trial, and W5–P5 on yet another trial, resulting in maximal contextual diversity. After 16 learning trials, children were tested on eight four-alternative forced-choice test trials, one per target word. On each test trial, a target referent was presented along with three foils from the set of objects that had never co-occurred with the target during the learning phase. Children were presented with a word and were asked to indicate which of the four pictures the word referred to. Fig. 22D and 22E show results from WOLVES 2.0 as compared to WOLVES and children. WOLVES 2.0 does a good job reproducing the key patterns from the data with, for instance, a decrease in the proportion correct as contextual diversity decreases. In summary, then, WOLVES 2.0 performs comparably to WOLVES in three canonical cross-situational word learning tasks that have been used with young children. Taken together with simulations from prior sections, this shows that WOLVES 2.0 can capture findings from all four tasks probed in this report: DCCS, Comprehension, Production, and Cross-situational Word Learning.

Discussion

The present paper has focused on a key challenge for the early executive function field – to understand the development of EF early in development when neurocognitive processes such as working memory and inhibitory control are also undergoing rapid and qualitative changes. We argued that to understand these complex co-developmental relationships, we needed a framework that specifies how neurocognitive processes including attention, inhibitory control, and working memory are integrated with higher-level cognitive control systems in a way that can be anchored to children's performance in key EF tasks. We suggested that DFT might be up to this challenge and provide a way to anchor theory to neural processes, opening the door toward understanding EF development at the levels of brain and behaviour.

We then proposed a modified theory of EF development, building on two prior theoretical advances: a dynamic field model of EF development (Buss & Kerr-German, 2019; Buss & Spencer, 2014, 2018) and a dynamic field theory of cross-situational word learning, Bhat et al.'s (2022) WOLVES model. In the latter case, we noted that WOLVES already integrates attention, working memory, and word learning. Thus, extending this model into the EF domain might provide a productive framework for understanding EF development. We were also motivated by empirical evidence suggesting that dimensional label learning plays a key role in EF development in the DCCS task (Lowery et al., 2022; McCraw et al., 2024). Thus, we asked if WOLVES might shed light on how children's learning of dimensional labels gives rise to dimensional attention skills, and how such dimensional attention abilities are manifest in the canonical DCCS task.

Simulation results showed that WOLVES 2.0 can reproduce quantitative details of both 'young' and 'old' children's performance in the DCCS task including variations in label usage from Yerys and Munakata (2006), the dimensional label comprehension task, the dimensional label production task, and cross-situational word learning. We achieved these quantitative fits with modest parameter changes across tasks, modulating 3 parameters in a task-specific manner. These 3 parameters influence how strongly the model relies on longer-term memory traces of word-feature mappings (i.e., its representation of how words map to objects), as well as the strength of activation passed between the word-feature and word fields. The latter parameters can be conceptualised as shifting the balance between top-down, word-driven attention (the input from the word field to the word-feature fields) relative to more bottom-up feature-driven attention to words (the input from the word-feature fields to the word field). From this perspective, the DCCS and Production tasks require strong bottom-up feature-driven activation of words, while Comprehension and CSWL require more top-down, word-driven attention.

The other key parameter changes needed to achieve quantitative fits to the data implemented our developmental hypothesis about how dimensional label learning impacts EF. Here, we modified the strength of the dimensional label boost when the model was instructed to attend to 'colour' versus 'shape', with 'young' models showing a weak dimensional label boost and 'old' models showing a stronger dimensional label boost. As noted, this was a placeholder for a more elaborated theory of how children learn dimensional labels. Here, we assume children have learned the mapping between the word 'colour' and individual colour labels. The dimensional boost in the model, then, reflects that word 'colour' pre-activating the set of individual colour labels. In addition, we used different word-feature mappings for the 'young' and 'old' models, with 'young' models having weaker memory traces for the individual colour and shape labels (e.g., 'red' and 'circle') with stronger representations of mis-matching representations as well (e.g., mis-remembering that the blue hue value maps to 'purple'). Again, these developmental changes were implemented 'by hand'; however, in this case, we have a strong foundation as WOLVES is a model of word learning. Thus, we are confident that WOLVES can learn these individual word-feature mappings, although it remains a key modelling task to demonstrate this in practice.

What have we achieved with WOLVES 2.0 and what challenges remain? There are several strengths of the model presented here. WOLVES 2.0 can learn words via autonomous visual exploration in environments with visual objects and auditory labels. Thus, we can present the model with inputs that match the properties of real objects in real environments (e.g., features in space) and also present isolated words to mimic, for instance, input from caregivers. The model will then visually explore the presented object-word pairings and, through cross-situational learning, identify the correct object-word mappings. These learned mappings can then be used in other contexts to, for instance, find the 'blue circle', or produce the label of the queried object. Moreover, the model can use words to operate on its scene representation, matching an object highlighted by the caregiver to find another object in the scene based on the object's colour or shape.

This ability of the model to generalise across different tasks extends beyond prior models of EF in important ways. For instance, the connectionist model of the DCCS task (Morton & Munakata, 2002) replicates some patterns of performance from specific variants of the DCCS task, but this model only performs the DCCS task. Thus, it cannot formally link performance across tasks in the way that WOLVES 2.0 can. Further, WOLVES 2.0 can be trained in ways that reflect real-world experience, and its learned associations can be probed in multiple ways that reflect that ways we probe children's knowledge in laboratory settings. This can give us insights into the mechanisms that underlie how children develop EF, not just in controlled laboratory settings, but also in more naturalistic real-world environments. Thus, WOLVES 2.0 provides an important bridge between theories of EF anchored to, for instance, the DCCS task, and how EF might develop in the wild. For instance, dimensional label learning teaches children not just what things are, but *how* to think about them.

Although WOLVES 2.0 has many strengths, it also has limitations. We have yet to specify how dimensional labels are learned (e.g., 'colour'). This is a key next step for the development of the theory. This requires tracking word-word associations to learn the statistical relationship that the word 'colour' is regularly paired with individual colour labels ('red', 'blue', etc.). In recent models, we have been incorporating connectionist style nodes into dynamic field models, an approach which is a good fit for the word-word association challenge (Tekulve & Schöner, 2022). We also note on this front that the memory traces used in WOLVES 2.0 may not be optimal. In recent models, we have used more standard forms of Hebbian learning which allow synaptic connections to have more stable patterns over longer timescales. For instance, Hebbian weights only update when the nodes linking the connectivity pattern are active, allowing synapses to remain stable when not being actively used. By contrast, the entire memory trace pattern either actively builds or decays in WOLVES 2.0 whenever the associated field is engaged. Thus, when we activate 'blue' and build a peak in the word-feature field, we decay the representation of 'red' a bit (although slowly). More standard forms of Hebbian learning implement local connectivity patterns that are protected from one another such that the learning of 'blue' does not necessarily interfere with the connectivity pattern representing 'red' (Sabinasz & Schöner, 2023).

We note that a central issue raised by our model is the relationship between EF and word learning in early development. Vygotsky (1962) proposed that children learn to self-regulate by internalising the language used by adults to scaffold self-regulatory behaviour. This view is consistent with our model which uses abstract words like 'colour' and 'shape' as neurocognitive 'tools' to focus processing on task-relevant information. We also note that other EF tasks such as the Flanker task or the Simon task involve controlling attention to visual stimuli, but do not rely on shifting attention to different dimensions 'on the fly' using dimensional labels. Thus, the DCCS task presents a unique window into EF development which relies upon children's ability to use language to guide object processing. Of course, EF is more than just one measure provided by the DCCS task, and language has been shown to be associated with various other aspects of EF as well. Future work developing the model could explore how processes like object category labels, directional labels (e.g., left, right, above, below), or self-regulation language (e.g., "no", "don't") impact children's ability to engage EF skills across more diverse tasks and contexts.

Another limitation of WOLVES 2.0 is the implementation of auditory processing. We have side-stepped this problem by providing isolated auditory inputs directly into the word field. Dynamic fields have been used to capture phonological properties of auditory processing, so it is possible that the auditory representation of words can be unpacked using DFT (Shaw & Tang, 2022). We have also simplified the spatial reference frames in WOLVES 2.0 by keeping the retinal and scene frames in alignment. More recent robotic models have shown that this is not necessary and that the concepts of DFT can be used to, for instance, build a scene representation in a robotic model that looks around a task space and operates on real objects with a robotic arm. This requires coordinating multiple frames of reference and transforming information back and forth from retinotopic to head- or body-centred reference frames (Richter et al., 2021).

A final limitation of WOLVES 2.0 is the complexity of the model. For instance, extending WOLVES to the DCCS task required expert-level knowledge of WOLVES as well as the framework of DFT. This limits how readily other researchers can innovate with this class of models, and also raises concerns about replicability – would other researchers given the same tools arrive at the same implementation of WOLVES 2.0? On this front, it is useful to emphasise that we have made all modelling code available on OSF (https://osf.io/mzw2x/) with tutorial videos on the DFT website (https://www.dynamicfieldtheory.org) so other researchers can readily reproduce our findings. We also note that expert-level knowledge of WOLVES is not required to make new discoveries. For instance, researchers could use the code we have provided to tweak the DCCS task to probe what WOLVES 2.0 does in new situations. This is largely just a programming task. One could change the number of trials, the instructions given to the model, the feature details (e.g., similarity), the trial sequence, and so on. One could even explore how modest variations in parameters – akin to individual differences – yield a range of behavioural outcomes. This is just one example of how new researchers could use WOLVES as a tool to make novel predictions. In summary, while innovation with WOLVES can be limited by its complexity, we have tried to reduce these barriers.

General Discussion: Reflections on the special issue

In this final section, we reflect on several key questions raised in the call for this special issue on EF development.

To what extent are the component constructs of EF separable and their developments interrelated? We have focused on the collection of EF skills, reified as components, that have been motivated by the latent variable approach of Miyake and colleagues (2000). The component view of EF is the prevailing perspective of EF in the literature. As we have iterated throughout this paper, however, the DF model is an integrated neural architecture built via generic principles of neural population dynamics. In this framework, EFs are best characterised as emergent properties of the integrated neural system rather than abilities which arise from the coordination of separable psychological constructs (see Perone, Simmering, & Buss, 2021). As noted previously, for example, inhibitory neural processes in the model give rise to behaviours that can be labelled as 'inhibitory control' behaviours (e.g., avoiding perseveration), even though there is no inhibitory control module in the model. An open question for future work is how the component structure identified in latent variable analyses relates to the DF model. One way to probe this would be to use WOLVES 2.0 to simulate a full battery of tasks used in the latent variable approach. This would help clarify how the latent variables map to parts of the WOLVES architecture.

With this type of future work in mind, what are the components of EF in WOLVES 2.0? One candidate 'component' is attention. Our model implements four types of attention: (1) spatial attention in both a retinal and a scene-based frame of reference, (2) feature-based attention, (3) top-down, word-driven attention, and (4) dimensional attention implemented 'by hand' via a boost to a set of dimensional labels. We also modulated the strength of the input from the word-feature field into the word field to implement a stronger 'bottom-up' bias to the word-feature patterns, which could be described as 'attending' more to the bottom-up information.

It is useful to ask in this context, what is 'attention' in a dynamic field model? In one sense, attention is not a thing in DFT, it is a state. Specifically, the 'attention' fields in WOLVES 2.0 are labelled this way because they are 'winner-take-all' fields that force 'attentional selection', that is, one peak is selected at a time. As discussed in Schöner, Spencer, and the DFT Research Group (2016), this plays a critical role in binding features to their spatial positions and binding features via words. Concretely, to pass information from, for instance, a retinal frame of reference to a scene frame of reference, we need to select a focus of attention in the retinal frame and rebind this information in the scene frame. 'Winner-take-all' dynamics allows this to happen accurately. All four types of attention listed above share these 'winner-take-all' dynamics, although the dimensional attention boost was not implemented neurally. The final example of 'attention' above (i.e., modulating the strength of the input from the word-feature field into the word field) contrasts with the others in that it is based on the modulation of a parameter in a particular field. If we were to implement this neurally, we would need some form of gain modulation on the connectivity pattern between, for instance, the word-feature fields to the word field. As this is a different type of neural implementation, perhaps it is clearest to avoid calling this a type of 'attention'.

If we think about another component in WOLVES 2.0 – working memory – we again see multiple working memory fields that serve different functions: (1) spatial working memory, (2) visual working memory (i.e., working memory for features), and (3) scene working memory. We have labelled these fields as 'working memory' fields as they have the potential to show self-sustaining activation patterns in the absence of input – the likely neural basis for working memory in the brain (Constantinidis & Steinmetz, 1996; Wei, Wang, & Wang, 2012). Thus, once again, working memory in DFT is a state of the system: working memory fields with strong self-excitation and strong self-inhibition have the potential to sustain activation patterns even after the input is removed. We might also think about the dimensional attention boost as having a working memory flavour in that the rules for the game in the DCCS task need to be actively maintained during each phase of the task.

Interestingly, however, the tasks we simulated do not really place many demands on working memory. There is no need to actively hold the target cards in mind as they are visible in the task space. Similarly, the test card is visible during the sorting trial. We would certainly need some form of verbal working memory to remember the instructions (e.g., sort by 'colour'), but as our model did not fully implement dimensional attention in a neural way, this form of working memory was not needed. Thus, from the perspective of working memory as a state, it's less clear what role working memory plays in the suite of tasks probed here.

That said, the fields we have labelled as 'working memory' in the model do, of course, play key roles in the tasks we simulated. For instance, the scene working memory plays a key role in building a scene representation of which features go where, and the model uses the scene representation to make sorting decisions. Interestingly, however, we did not operate the scene working memory in a 'working memory mode' with sustaining activation patterns. Rather, we relied on re-activation of patterns from long-term memory 'on-the-fly' as the model made decisions by combining word-feature associations with its scene representation. This is related to recent debates in the working memory literature about whether working memory is driven by self-sustaining activation or re-activation from long-term memory (Schneegans & Bays, 2017). DFT shows that both forms of 'working memory' can emerge from the same neural architecture depending on the task demands and the state of the neural field being probed.

What about inhibitory control in the model? How is that implemented? If we focus on the DCCS task, inhibitory control is usually conceptualised as inhibition of a prepotent response, that is, inhibiting the pull toward a colour-based response during a post-switch trial when asked to sort by 'shape'. In WOLVES 2.0, this type of inhibition does not occur – there is no explicit process that shuts down the 'colour' system when playing the 'shape' game. Nevertheless, conceptually this type of inhibition may be involved in a more detailed model that activates a 'shape' node, holds this activation in working memory, and inhibits the activity of a 'colour' node (similar to how the PDP model works, see Morton & Munakata, 2002). In this case, would this be an example of working memory, inhibitory control, or both?

This leads us to our central point: once we have an articulated neural process model like WOLVES 2.0, the question is less about broad categories of psychological components like attention and working memory, and more about the details of how the neurocognitive processes in the model work together. A major problem with the concept of EF 'components' that has come to dominate research on EF is that they are largely defined by the tasks used to measure them. This is evident in the simulations we presented here. For instance, the DCCS task does not probe an isolated EF component in our model separate from the comprehension, production, or CSWL tasks. Rather, each task involved the engagement of an integrated system of attention, working memory, and word learning processes. In the DFT framework, we can define a neurocognitive system and examine how behaviour is organised across different task demands. The organisation of behaviour around goals can be thought of as a product of coupled perception–action systems that can dynamically influence one another. Through the interaction of these systems, new behaviours can emerge that were not built into any individual system. These new behaviours emerge through quantitative changes driven by learning that then impact how these perception–action systems interact with one another.

Indeed, when viewed through the lens of WOLVES 2.0, it becomes clear that EF is not a thing; rather, executive function is an emergent property of the system. For instance, when the 'old' model shows 'good EF', it is inhibiting a prepotent response in a word-feature field to respond based on colour via a top-down, word-driven bias. This helps the model recall its shape label ('circle') which biases the system to focus on the circle feature when deciding – based on its scene representation – that circles go to the right. If we think about development in this context, then good EF will emerge when children have a good understanding of dimensional labels, a good verbal working memory to keep the dimensional label in mind, strong top-down, word-driven attention to the shape features, a good scene working memory to build a robust scene representation, and a well-coordinated system that allows word-feature decisions to impact sorting decisions. While complicated, having an articulated and detailed model gives us a robust tool to understand the reciprocal back-and-forth between developmental strengths/weaknesses in particular neurocognitive processes and how these ultimately influence EF performance in specific tasks.

And, of course, this is just the tip of the iceberg. We've focused on dimensional label learning, but we know there are parallel, co-developments in the other parts of the system during the window from age 3 to age 5. For instance, VWM improves during this period (Simmering, 2016) as does spatial working memory (Schutte & Spencer, 2009). We contend that WOLVES 2.0 provides a strong theoretical framework to ground such issues. For instance, WOLVES 2.0 contains the same architecture for visual working memory used to simulate changes in working memory in infancy (Perone et al., 2011) and childhood (Simmering, 2016). Similarly, WOLVES contains all the ingredients used to simulate changes in A-not-B performance in infancy (Thelen et al., 2001). Thus, in future work, we should be poised to capture co-developmental changes in A-not-B performance, visual working memory performance, and later performance in the DCCS task using a single, integrated model.

Can one component be impaired while the others are spared? As is evident from the discussion above, there are many relevant 'components' in our theory of EF – each field could be considered a component. Consequently, there are many possible ways in which a developing EF system like WOLVES could show 'impairments' or, more optimistically, areas of 'strength'. Manipulating parameters in WOLVES 2.0 can give us novel insights into individual differences as subtle differences in parameter values in one component of the model – for example the strength of self-excitation, which would impact VWM – may led to differences in behaviour in, for instance, the DCCS task. We could manipulate any piece of the neural architecture and make novel predictions about how a localised strength or impairment should impact behaviour in predictable ways.

To what extent do the neural processes of EF overlap and does this change with age? One of the motivating issues raised in the introduction was to use dynamic field theory to gain insights into the development of EF at the levels of both brain and behaviour. Indeed, the strong claim made by DFT is that these are not separate levels; rather, behaviour emerges from the type of neural population dynamics captured in DFT. This view was formalised by Buss and colleagues (Buss, Magnotta, Penny, Schöner, & Spencer, 2020) where we adopted an integrative cognitive neuroscience approach to directly simulate brain activity from a DF model of visual working memory, generating hemodynamic predictions that we tested with fMRI. In this way, we did not have a separate model that predicted behavioural patterns at one level and brain patterns at another. Rather, one integrated model captured brain and behaviour simultaneously.

A related approach was used in our prior work looking at the development of EF in children using fNIRS (Buss & Spencer, 2018).

Here, we used our DF model of EF development to predict children's performance in 'easy' and 'hard' versions of the DCCS task along with patterns of brain activity from frontal and posterior cortical fields. The novel prediction of the model was that 3-year-olds would show frontal activity in 'easy' versions of the DCCS task, even though such children have an 'immature' frontal cortex. This frontal activity arose in the model from bottom-up activation in posterior cortical fields reflecting the type of cooperativity shown in Fig. 10. We also predicted that 3- and 4-year-olds would show enhanced posterior activation in an 'easy' DCCS task reflecting stronger down-top attention from a dimensional labelling system. Both predictions were supported.

What about with WOLVES 2.0: what neural insights might this model offer regarding the neural components of EF? The strong claim is that the fields implemented in WOLVES 2.0 should be open to direct neural measurement. Concretely, using the approach developed by Buss et al. (2020), it should be possible to simulate brain hemodynamics from the model directly and measure these predicted hemodynamics using fNIRS. Buss and colleagues (Lowery et al., 2022) have already used fNIRS to examine children's neural dynamics in the DCCS, comprehension, and production tasks longitudinally. Thus, future work will probe if we can use the integrative cognitive neuroscience approach with WOLVES 2.0 to model these data directly. The strong claim is that the fields instantiated in WOLVES 2.0 'live' directly in the brains of children and can be measured using task-based brain imaging methods.

What can the different EF literatures learn from one another? Our response to this question is clear: to understand the development of EF, particularly in early development where 'components' are co-developing, we need strong integration across literatures. The model we proposed here gives a starting point for such work. Thus, we encourage other researchers to step in, run simulations, and test predictions.

To make this vision a reality, however, we recognise the need to make complex theories like DFT more accessible and more usable. We are actively working on this. We have a new simulation platform called Cedar that makes it a bit easier to work with DF models (see www.dynamicfieldtheory.org). We are also developing ways to input real task stimuli into the model. This involves using CNNs as the front-end perceptual input to the model so researchers could use images of their own stimuli as input to the model, rather than scripting Gaussian inputs as we showed in our simulations. We have demonstrated that this is possible, but we acknowledge that more work is needed to make the ideas presented here both understandable for non-modelers and also useful for non-modelers.

What similarities/differences are there in how EF components are measured throughout development? A key challenge in the literature, especially in the early EF literature, is that researchers use different tasks at different ages. Consequently, there is no guarantee that each 'inhibitory' task measures 'inhibition' in the same way or measures the same type of inhibition. The same is true for measures of attention and working memory. Thus, we need to either develop standardised measures that can be used at multiple ages, or we need to understand exactly how the different tasks used relate to one another. WOLVES 2.0 is the type of theory that might facilitate such work as it effectively generalises across tasks, as we have shown here. Indeed, more generally, DFT has shown a strong ability to integrate findings from multiple tasks as reflected in early work using a model of spatial working memory to generalise across spatial memory and discrimination tasks (Simmering & Spencer, 2008). More recent work has generalised a model of visual working memory across preferential looking and change detection tasks, revealing why estimates of visual working memory capacity differs across tasks and over development (Perone et al., 2011; Simmering, 2016).

Conclusions

We presented here a unique and rigorous theoretical framework that integrated diverse research including early word learning, dimensional label learning, and EF. Although these are not lines of work that are typically discussed together, we demonstrated that label learning, in particular dimensional label learning, provides the first quantitatively defined learning mechanism that can explain the development of EF. This theoretical framework redefines how we think about EFs, framing them as an emergent property of a complex, neurocognitive system. In doing so, this perspective offers a new way of conceptualising the relationship between different measures of EF and, more importantly, offers a new perspective on efforts to improve EF through learning. While the majority of research on EF training has failed to produce generalised effects on developmental outcomes (Niebaum & Munakata, 2023) by focusing on training EF components themselves, our work here motivates new lines of work looking beyond EF itself to examine the impact of label learning on EF development.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work supported by funding to JPS from the Leverhulme Trust under grant RPG-2021-350.

References

Ahmed, S. F., Tang, S., Waters, N. E., & Davis-Kean, P. (2019). Executive function and academic achievement: Longitudinal relations from early childhood to adolescence. *Journal of Educational Psychology, 111*(3), 446–458. https://doi.org/10.1037/edu0000296

Bardikoff, N., & Sabbagh, M. (2017). The differentiation of executive functioning across development: Insights from developmental cognitive neuroscience. *New Perspectives on Human Development, 47–66*. https://doi.org/10.1017/CBO9781316282755.005

Bell, M. A. (2012). A Psychobiological Perspective on Working Memory Performance at 8 Months of Age. *Child Development, 83*(1), 251–265. https://doi.org/10.1111/j.1467-8624.2011.01684.x

Benitez, V. L., Vales, C., Hanania, R., & Smith, L. B. (2017). Sustained selective attention predicts flexible switching in preschoolers. *Journal of Experimental Child Psychology, 156*, 29–42. https://doi.org/10.1016/J.JECP.2016.11.004

Berry, E. D. J., Waterman, A. H., Baddeley, A. D., Hitch, G. J., & Allen, R. J. (2018). The limits of visual working memory in children: Exploring prioritization and recency effects with sequential presentation. *Developmental Psychology, 54*(2), 240–253. https://doi.org/10.1037/DEV0000427

Bhat, A. A., Samuelson, L. K., & Spencer, J. P. (2023). Formal theories clarify the complex: Generalizing a neural process account of the interaction of visual exploration and word learning in infancy. *Child Development, 94*(6), 1491–1510. https://doi.org/10.1111/CDEV.14023

Bhat, A. A., Spencer, J. P., & Samuelson, L. K. (2022). Word-Object Learning via Visual Exploration in Space (WOLVES): A neural process model of cross-situational word learning. *Psychological Review, 129*(4), 640–695. https://doi.org/10.1037/rev0000313

Blair, C., & Razza, R. P. (2007). Relating effortful control, executive function, and false belief understanding to emerging math and literacy ability in kindergarten. *Child Development, 78*(2), 647–663.

Brooks, P. J., Hanauer, J. B., Padowska, B., & Rosman, H. (2003). The role of selective attention in preschoolers' rule use in a novel dimensional card sort. *Cognitive Development, 18*(2), 195–215.

Bull, R., Espy, K. A., & Wiebe, S. A. (2008). Short-Term Memory, Working Memory, and Executive Functioning in Preschoolers: Longitudinal Predictors of Mathematical Achievement at Age 7 Years. *Developmental Neuropsychology, 33*(3), 205–228. https://doi.org/10.1080/87565640801982312

Buss, A. T., & Kerr-German, A. (2019). Dimensional attention as a mechanism of executive function: Integrating flexibility, selectivity, and stability. *Cognition, 192*, Article 104003. https://doi.org/10.1016/J.COGNITION.2019.06.015

Buss, A. T., Magnotta, V. A., Penny, W. D., Schöner, G., & Spencer, J. P. (2020). How do neural processes give rise to cognition? Simultaneously predicting brain and behavior with a dynamic model of visual working memory. *Psychological Review, 128*(2), 362–395.

Buss, A. T., & Spencer, J. P. (2014). The emergent executive: A dynamic field theory of the development of executive function. Monographs of the Society for Research in Child Development, 79(2). DOI: 10.1002/mono.12096.

Buss, A. T., & Spencer, J. P. (2018). Changes in frontal and posterior cortical activity underlie the early emergence of executive function. *Developmental Science*. https://doi.org/10.1111/desc.12602

Cheng, C., Kaldy, Z., & Blaser, E. (2020). Coding of featural information in visual working memory in 2.5-year-old toddlers. *Cognitive Development, 55*, Article 100892. https://doi.org/10.1016/J.COGDEV.2020.100892

Constantinidis, C., & Steinmetz, M. A. (1996). Neuronal activity in posterior parietal area 7a during the delay periods of a spatial memory task. Retrieved from *Journal of Neurophysiology, 76*(2) http://jn.physiology.org/content/76/2/1352.short.

Cortés Pascual, A., Moyano Muñoz, N., & Quílez Robres, A. (2019). The Relationship Between Executive Functions and Academic Performance in Primary Education: Review and Meta-Analysis. *Frontiers in Psychology, 10*, Article 449759. https://doi.org/10.3389/fpsyg.2019.01582

Cuevas, K., & Bell, M. A. (2010). Developmental progression of looking and reaching performance on the A-not-B task. *Developmental Psychology, 46*(5), 1363–1371. https://doi.org/10.1037/a0020185

Davidson, M. C., Amso, D., Anderson, L. C., & Diamond, A. (2006). Development of cognitive control and executive functions from 4 to 13 years: Evidence from manipulations of memory, inhibition, and task switching. *Neuropsychologia, 44*(11), 2037–2078.

Diamond, A. (2002). Normal development of prefrontal cortex from birth to young adulthood: Cognitive functions, anatomy, and biochemistry. In D. T. Stuss, & R. T. Knight (Eds.), *Principles of Frontal Lobe Function* (pp. 465–503). New York: Oxford University Press.

Diamond, A., & Lee, K. (2011). Interventions shown to aid executive function development in children 4 to 12 years old. *Science, 333*(6045), 959–964.

Dunn, J. R. (2010). Health behavior vs the stress of low socioeconomic status and health outcomes. *The Journal of the American Medical Association, 303*(12), 1199–1200.

Eakin, L., Minde, K., Hechtman, L., Ochs, E., Krane, E., Bouffard, R., & Looper, K. (2004). The marital and family functioning of adults with ADHD and their spouses. *Journal of Attention Disorders, 8*(1), 1–10.

Friedman, N. P., & Miyake, A. (2004). The Relations Among Inhibition and Interference Control Functions: A Latent-Variable Analysis. *Journal of Experiment Psychology: General, 133*(1), 101–135. https://doi.org/10.1037/0096-3445.133.1.101

Friedman, N. P., Miyake, A., Young, S. E., DeFries, J. C., Corley, R. P., & Hewitt, J. K. (2008). Individual Differences in Executive Functions Are Almost Entirely Genetic in Origin. *Journal of Experimental Psychology: General, 137*(2), 201–225. https://doi.org/10.1037/0096-3445.137.2.201

Frye, D., Zelazo, P. D., & Palfai, T. (1995). Theory of mind and rule-based reasoning. *Cognitive Development, 10*(4), 483–527.

Gathercole, S. E., Pickering, S. J., Knight, C., & Stegmann, Z. (2004). Working Memory Skills and Educational Attainment: Evidence from National Curriculum Assessments at 7 and 14 Years of Age. *Applied Cognitive Psychology, 18*(1), 1–16.

Gerstadt, L., Hong, Y. J., & Diamond, A. (1994). The relationship between cognition and action: Performance of children 33-7 years old on a Stroop- like day-night test. *Cognition, 53*, 129–153.

Goldstein, S., & Naglieri, J. A. (2014). Handbook of executive functioning. *Handbook of Executive Functioning, 1–567*. https://doi.org/10.1007/978-1-4614-8106-5/COVER

Heckman, J. J. (2006). Skill formation and the economics of investing in disadvantaged children. *Science (New York, N.Y.), 312*(5782), 1900–1902. https://doi.org/10.1126/science.1128898

Heckman, J. J. (2011). The Economics of Inequality: The Value of Early Childhood Education. *American Educator, 35(1), 31–36*, 47.

Heckman, J. J., Savelyev, P., & Yavitz, A. (2010). A New Cost-Benefit and Rate of Return Analysis for the Perry Preschool Program : A Summary. In A. J. Reynolds, A. Rolnick, M. Englund, & J. A. Temple (Eds.), *Cost-Effective Programs in Children's First Decade a Human Capital Integration* (pp. 366–380). New York: Cambridge University Press.

Holmboe, K., Larkman, C., de Klerk, C., Simpson, A., Bell, M. A., Patton, L., & Dvergsdal, H. (2021). The early childhood inhibitory touchscreen task: A new measure of response inhibition in toddlerhood and across the lifespan. *PLOS ONE, 16*(12). https://doi.org/10.1371/journal.pone.0260695

Hood, B. M., & Atkinson, J. (1993). Disengaging visual attention in the infant and adult. *Infant Behavior and Development, 16*(4), 405–422. https://doi.org/10.1016/0163-6383(93)80001-O

Howard, S. J., Okely, A. D., & Ellis, Y. G. (2015). Evaluation of a differentiation model of preschoolers' executive functions. *Frontiers in Psychology, 6*, 285. https://doi.org/10.3389/fpsyg.2015.00285

Kerr-German, A., Tas, A. C., & Buss, A. T. (2023). A multi-method approach to addressing the toddler data desert in attention research. *Cognitive Development, 65*, Article 101293. https://doi.org/10.1016/J.COGDEV.2022.101293

Kirkham, N. Z., Cruess, L., & Diamond, A. (2003). Helping children apply their knowledge to their behavior on a dimension-switching task. *Developmental Science, 6*(5), 449–467. https://doi.org/10.1111/1467-7687.00300

Kochanska, G., Murray, K. T., & Harlan, E. T. (2000). Effortful control in early childhood: Continuity and change, antecedents, and implications for social development. *Developmental Psychology, 36*(2), 220–232. https://doi.org/10.1037/0012-1649.36.2.220

Lee, K., Ng, S. F., Pe, M. L., Ang, S. Y., Hasshim, M. N. A. M., & Bull, R. (2012). The cognitive underpinnings of emerging mathematical skills: Executive functioning, patterns, numeracy, and arithmetic. *British Journal of Educational Psychology, 82*(1), 82–99. https://doi.org/10.1111/j.2044-8279.2010.02016.x

Lowery, K., Nikam, B., & Buss, A. T. (2022). Dimensional label learning contributes to the development of executive functions. *Scientific Reports, 12*(1), 11008. https://doi.org/10.1038/s41598-022-14761-2

Markounikau, V., Igel, C., Grinvald, A., & Jancke, D. (2010). A Dynamic Neural Field Model of Mesoscopic Cortical Activity Captured with Voltage-Sensitive Dye Imaging. *PloS Computational Biology, 6*(9). https://doi.org/10.1371/journal.pcbi.1000919

McAuley, T., White, D., & a.. (2011). A latent variables examination of processing speed, response inhibition, and working memory during typical development. *Journal of Experimental Child Psychology, 108*(3), 453–468. https://doi.org/10.1016/j.jecp.2010.08.009

McCraw, A., Sullivan, J., Lowery, K., Eddings, R., Heim, H. R., & Buss, A. T. (2024). A dynamic field theory of executive function: Identifying early neurocognitive markers. *Monographs of the Society for Research in Child Development, 89*(3).

McKenna, R., Rushe, T., & Woodcock, K. A. (2017). Informing the structure of executive function in children: A meta-analysis of functional neuroimaging data. *Frontiers in Human Neuroscience, 11*, Article 227082. https://doi.org/10.3389/FNHUM.2017.00154/BIBTEX

Miller, M. R., Giesbrecht, G. F., Müller, U., McInerney, R. J., & Kerns, K. A. (2012). A Latent Variable Approach to Determining the Structure of Executive Function in Preschool Children. *Journal of Cognition and Development, 13*(3), 395–423. https://doi.org/10.1080/15248372.2011.585478

Miyake, A., Friedman, N. P., Emerson, M. J., Witzki, H., Howerter, A., & Wager, T. D. (2000). The unity and diversity of executive functions and their contributions to complex "Frontal Lobe" tasks: A latent variable analysis. *Cognitive Psychology, 41*(1), 49–100. https://doi.org/10.1006/COGP.1999.0734

Moffitt, T. E., Arseneault, L., Belsky, D., Dickson, N., Hancox, R. J., Harrington, H., & Caspi, A. (2011). A gradient of childhood self-control predicts health, wealth, and public safety. *Proceedings of the National Academy of Sciences of the United States of America, 108*(7), 2693–2698.

Morton, J. B., & Munakata, Y. (2002). Active Versus Latent Representations: A Neural Network Model of Perseveration, Dissociation, and Decalage. *Developmental Psychobiology, 40*(3), 255–265.

Müller, U., Dick, A. S., Gela, K., Overton, W. F., & Zelazo, P. D. (2006). The role of negative priming in preschoolers' flexible rule use on the dimensional change card sort task. *Child Development, 77*(2), 395–412.

Niebaum, J. C., & Munakata, Y. (2023). Why Doesn't Executive Function Training Improve Academic Achievement? Rethinking Individual Differences, Relevance, and Engagement from a Contextual Framework. *Journal of Cognition and Development, 24*(2), 241–259. https://doi.org/10.1080/15248372.2022.2160723

Nigg, J. T. (2000). On Inhibition/Disinhibition in Developmental Psychopathology: Views from Cognitive and Personality Psychology and a Working Inhibition Taxonomy. *Psychological Bulletin, 126*(2), 220–246. https://doi.org/10.1037/0033-2909.126.2.220

Perner, J., & Lang, B. (2002). What Causes 3-year-olds ' Difficulty on the Dimensional Change Card Sorting Task ? I. *Infant and Child Development, 11*, 93–105.

Perone, S., Molitor, S. J., Buss, A. T., Spencer, J. P., & Samuelson, L. K. (2015). Enhancing the Executive Functions of 3-Year-Olds in the Dimensional Change Card Sort Task. *Child Development, 86*(3). https://doi.org/10.1111/cdev.12330

Perone, S., Palanisamy, J., & Carlson, S. M. (2018). Age-related change in brain rhythms from early to middle childhood: Links to executive function. *Developmental Science, 21*(6). https://doi.org/10.1111/desc.12691

Perone, S., Plebanek, D. J., Lorenz, M. G., Spencer, J. P., & Samuelson, L. K. (2019). Empirical Tests of a Brain-Based Model of Executive Function Development. *Child Development, 90*, 210–226. https://doi.org/10.1111/cdev.12885

Perone, S., Simmering, V. R., & Buss, A. T. (2021). A dynamical reconceptualization of executive-function development. Perspectives. *Psychological Science, 16*(6). https://doi.org/10.1177/1745691620966792

Perone, S., Simmering, V. R., & Spencer, J. P. (2011). Stronger neural dynamics capture changes in infants' visual short-term memory capacity over development. *Developmental Science, 14*, 1379–1392.

Perone, S., & Spencer, J. P. (2013). Autonomous visual exploration creates developmental change in familiarity and novelty seeking behaviors. *Frontiers in Psychology, 4*, 648.

Portugal, A. M., Bedford, R., Cheung, C. H. M., Mason, L., & Smith, T. J. (2021). Longitudinal touchscreen use across early development is associated with faster exogenous and reduced endogenous attention control. Scientific Reports 2021 11:1, 11(1), 1–12. DOI: 10.1038/s41598-021-81775-7.

Prince, M., Patel, V., Saxena, S., Maj, M., Maselko, J., Phillips, M. R., & Rahman, A. (2007). No health without mental health. *Lancet, 370*(9590), 859–877.

Richmond-Rakerd, L. S., Caspi, A., Ambler, A., d'Arbeloff, T., de Bruine, M., Elliott, M., & Moffitt, T. E. (2021). Childhood self-control forecasts the pace of midlife aging and preparedness for old age. *Proceedings of the National Academy of Sciences of the United States of America, 118*(3), Article 2010211118. https://doi.org/10.1073/pnas.2010211118

Richter, M., Lins, J., & Schöner, G. (2021). A Neural Dynamic Model of the Perceptual Grounding of Spatial and Movement Relations. *Cognitive Science, 45*(10). https://doi.org/10.1111/COGS.13045

Robson, D. A., Allen, M. S., & Howard, S. J. (2020). Self-regulation in childhood as a predictor of future outcomes: A meta-analytic review. *Psychological Bulletin, 146*(4), 324–354. https://doi.org/10.1037/bul0000227

Sabinasz, D., & Schöner, G. (2023). A Neural Dynamic Model Perceptually Grounds Nested Noun Phrases. *Topics in Cognitive Science, 15*(2), 274–289. https://doi.org/10.1111/TOPS.12630

Samuelson, L. K., Smith, L. B., Perry, L. K., & Spencer, J. P. (2011). Grounding word learning in space. *PloS One, 6*(12).

Samuelson, L. K., Spencer, J. P., & Jenkins, G. W. (2013). A Dynamic Neural Field Model of Word Learning. *Theoretical and Computational Models of Word Learning: Trends in Psychology and Artificial Intelligence, 1–27*. https://doi.org/10.4018/978-1-4666-2973-8.CH001

Sandhofer, C. M., & Smith, L. B. (1999). Learning color words involves learning a system of mappings. Retrieved from *Developmental Psychology, 35*(3), 668–679 http://www.ncbi.nlm.nih.gov/pubmed/10380858.

de Santana, A. N., Roazzi, A., & Nobre, A. P. M. C. (2022). The relationship between cognitive flexibility and mathematical performance in children: A meta-analysis. *Trends in Neuroscience and Education, 28*, Article 100179. https://doi.org/10.1016/J.TINE.2022.100179

Schneegans, S., & Bays, P. M. (2017). Restoration of fMRI Decodability Does Not Imply Latent Working Memory States. *Journal of Cognitive Neuroscience, 29*(12), 1977–1994. https://doi.org/10.1162/jocn_a_01180

Schneegans, S., Spencer, J. P., & Schöner, G. (2016). Integrating "What" and "Where": Visual Working Memory for Objects in a Scene. In G. Schöner, J. P. Spencer, & T. D. R. Group (Eds.), *Dynamic Thinking—A Primer on Dynamic Field Theory* (pp. 197–226). New York: Oxford University Press.

Schöner, G., Spencer, J. P., & Group, T. D. R. (2016). *Dynamic Thinking: A primer on Dynamic Field Theory*. New York: Oxford University Press.

Schutte, A. R., & Spencer, J. P. (2009). Tests of the dynamic field theory and the spatial precision hypothesis: Capturing a qualitative developmental transition in spatial working memory. *Journal of Experimental Psychology. Human Perception and Performance, 35*(6), 1698–1725. https://doi.org/10.1037/a0015794

Shaw, J. A., & Tang, K. (2022). A dynamic neural field model of leaky prosody: Proof of concept. In *Proceedings of the Annual Meetings on Phonology* (p. 10). https://doi.org/10.3765/AMP.V10I0.5442

Simmering, V. R. (2012). The development of visual working memory capacity during early childhood. *Journal of Experimental Child Psychology, 111*(4), 695–707. https://doi.org/10.1016/j.jecp.2011.10.007

Simmering, V. R. (2016). Working memory in context: Modeling dynamic processes of behavior, memory, and development. *Monographs of the Society for Research in Child Development, 81*(3), 1–166. https://doi.org/10.1111/mono.12249

Simmering, V. R., & Spencer, J. P. (2008). Generality with specificity: The dynamic field theory generalizes across tasks and time scales. *Developmental Science, 11*, 541–555.

Simpson, A., & Riggs, K. J. (2006). Conditions under which children experience inhibitory difficulty with a "button-press" go/no-go task. *Journal of Experimental Child Psychology, 94*(1), 18–26. https://doi.org/10.1016/J.JECP.2005.10.003

Smith, L., & Yu, C. (2008). Infants rapidly learn word-referent mappings via cross-situational statistics. *Cognition, 106*(3), 1558–1568. https://doi.org/10.1016/J.COGNITION.2007.06.010

Spencer, J. P. (2020). The Development of Working Memory. *Current Directions in Psychological Science, 29*(6), 545–553. https://doi.org/10.1177/0963721420959835/ASSET/IMAGES/LARGE/10.1177_0963721420959835-FIG2.JPEG

Spiegel, J. A., Goodrich, J. M., Morris, B. M., Osborne, C. M., & Lonigan, C. J. (2021). Relations between executive functions and academic outcomes in elementary school children: A meta-analysis. *Psychological Bulletin, 147*(4), 329–351. https://doi.org/10.1037/bul0000322

St Clair-Thompson, H. L., & Gathercole, S. E. (2006). Executive functions and achievements in school: Shifting, updating, inhibition, and working memory. *Quarterly Journal of Experimental Psychology, 59*(4), 745–759. https://doi.org/10.1080/17470210500162854

Steelandt, S., Thierry, B., Broihanne, M. H., & Dufour, V. (2012). The ability of children to delay gratification in an exchange task. *Cognition, 122*(3), 416–425. https://doi.org/10.1016/J.COGNITION.2011.11.009

Suanda, S. H., Mugwanya, N., & Namy, L. L. (2014). Cross-situational statistical word learning in young children. *Journal of Experimental Child Psychology, 126*, 395–411. https://doi.org/10.1016/J.JECP.2014.06.003

Tamm, L., Menon, V., & Reiss, A. L. (2002). Maturation of Brain Function Associated with Response Inhibition. *Journal of the American Academy of Child and Adolescent Psychiatry, 41*(10), 1231–1238. https://doi.org/10.1097/00004583-200210000-00013

Tekulve, J., & Schöner, G. (2022). A Neural Dynamic Network Drives an Intentional Agent That Autonomously Learns Beliefs in Continuous Time. *IEEE Transactions on Cognitive and Developmental Systems, 14*(1), 90–101. https://doi.org/10.1109/TCDS.2020.3013768

Thelen, E., Schöner, G., Scheier, C., & Smith, L. B. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. Retrieved from *The Behavioral and Brain Sciences, 24(1), 1–34* http://www.ncbi.nlm.nih.gov/pubmed/11515285.

Treisman, A., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12*(1), 97–136.

Verdine, B. N., Lucca, K. R., Golinkoff, R. M., Hirsh-Pasek, K., & Newcombe, N. S. (2016). The Shape of Things: The Origin of Young Children's Knowledge of the Names and Properties of Geometric Forms. *Journal of Cognition and Development, 17*(1), 142–161. https://doi.org/10.1080/15248372.2015.1016610

Vlach, H. A., & DeBrock, C. A. (2019). Statistics learned are statistics forgotten: Children's retention and retrieval of cross-situational word learning. *Journal of Experimental Psychology: Learning Memory and Cognition, 45*(4), 700–711. https://doi.org/10.1037/XLM0000611

Vlach, H. A., & Johnson, S. P. (2013). Memory constraints on infants' cross-situational statistical learning. *Cognition, 127*(3), 375–382. https://doi.org/10.1016/J.COGNITION.2013.02.015

Vygotsky, L. (1962). *Thought and Language*. Cambridge, MA: MIT Press.

Wei, Z., Wang, X.-J., & Wang, D.-H. (2012). From distributed resources to limited slots in multiple-item working memory: A spiking network model with normalization. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience, 32*(33), 11228–11240. https://doi.org/10.1523/JNEUROSCI.0735-12.2012

Wiebe, S. A., Espy, K. A., & Charak, D. (2008). Using confirmatory factor analysis to understand executive control in preschool children: I. *Latent structure. Developmental Psychology, 44*(2), 575–587. https://doi.org/10.1037/0012-1649.44.2.575

Wiebe, S. A., Sheffield, T., Nelson, J. M., Clark, C., & a C., Chevalier, N., & Espy, K. A.. (2011). The structure of executive function in 3-year-olds. *Journal of Experimental Child Psychology, 108*(3), 436–452.

Yerys, B. E., & Munakata, Y. (2006). When labels hurt but novelty helps: Children's perseveration and flexibility in a card-sorting task. *Child Development, 77*(6), 1589–1607.

Yu, C., & Smith, L. B. (2011). What you learn is what you see: Using eye movements to study infant cross-situational word learning. *Developmental Science, 14*(2), 165–180. https://doi.org/10.1111/J.1467-7687.2010.00958.X

Zelazo, P. D., Anderson, J. E., Richler, J., Wallner-Allen, K., Beaumont, J. L., & Weintraub, S. (2013). II. NIH Toolbox Cognition Battery (CB): measuring executive function and attentioN. *Monographs of the Society for Research in Child Development, 78*(4), 16–33. https://doi.org/10.1111/mono.12032

Zelazo, P. D., Muller, U., Frye, D., & Marcovitch, S. (2003). The development of executive function in early childhood. *Monographs of the Society for Research in Child Development, 2*(68), 1–137.