

Feature Review

# Rationalizing constraints on the capacity for cognitive control

Sebastian Musslick <sup>1,\*</sup> and Jonathan D. Cohen<sup>1,2</sup>

Humans are remarkably limited in: (i) how many control-dependent tasks they can execute simultaneously, and (ii) how intensely they can focus on a single task. These limitations are universal assumptions of most theories of cognition. Yet, a rationale for why humans are subject to these constraints remains elusive. This feature review draws on recent insights from psychology, neuroscience, and machine learning, to suggest that constraints on cognitive control may result from a rational adaptation to fundamental, computational dilemmas in neural architectures. The reviewed literature implies that limitations in multitasking may result from a trade-off between learning efficacy and processing efficiency and that limitations in the intensity of commitment to a single task may reflect a trade-off between cognitive stability and flexibility.

## The role of capacity constraints in human cognition

One of the most remarkable features of human cognition is the ability to rapidly adapt behavior in a changing world. This is often attributed to the capacity for cognitive control: the ability to flexibly direct behavior in pursuit of a goal (Box 1). Cognitive control is engaged by all of the higher mental faculties that distinguish humans from other species, including reasoning, problem solving, planning, and the use of symbolic language [1]. Yet, humans are strikingly limited in how many control-demanding **tasks** (see Glossary) they can perform simultaneously (e.g., reading a document while listening to a friend) or how intensely they can focus on a single task (e.g., parsing a mathematical equation in a noisy environment). The significance of these limitations is not only apparent in daily life. They are also a fundamental premise of general theories of human cognition (e.g. [2–7]). These theories posit that the exertion of cognitive control is associated with a cost, and that humans consider this cost when making decisions about how to allocate control [7–9]. The notion of a cost and concomitant constraints on control, can help integrate a wide range of empirical findings concerning the allocation of mental effort [10–13], the selection between cognitive heuristics [5], planning [14,15], or cognitive impairments in depression [16]. Yet, none of these theories provides an explanation for why control-dependent processing would be subject to these limitations in the first place.

Here, we review two fundamental, computational dilemmas that arise in neural systems and suggest that these provide a rational account of constraints on cognitive control. First, we review empirical and computational evidence suggesting a trade-off between the rapid acquisition of novel tasks (learning efficacy), that is promoted by sharing **representations** across tasks, on the one hand; and **multitasking** capability (processing efficiency) that is achieved by separating representations and dedicating them to individual tasks, on the other hand. The work reviewed suggests that limitations in the ability to simultaneously execute multiple tasks may reflect a preference to learn tasks more quickly. Immediate rewards associated with quickly acquiring a task may often outweigh greater but later rewards associated with the ability to execute that task in parallel with others. On this view, a purpose of cognitive control is to prevent the simultaneous execution of tasks that share representations, to avoid interference that could arise if those

## Highlights

To explain human behavior, most general theories of cognition assume, rather than explain, limitations in: (i) the number of control-dependent tasks that can be performed simultaneously (i.e., multitasked); and (ii) the amount of cognitive control that can be allocated to a single task.

Limitations in the capability to multitask can be explained by representation sharing between tasks. Computational modeling suggests that neural systems trade the benefits of shared representation for rapid learning and generalization (a mechanism increasingly exploited in machine learning) against constraints on multitasking performance.

Experimental studies posit a trade-off between cognitive stability and cognitive flexibility. Computational analyses of this trade-off suggest that adaptations to high demands for flexibility limit the amount of control that can be allocated to a single task.

<sup>1</sup>Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08544, USA

<sup>2</sup>Department of Psychology, Princeton University, Princeton, NJ 08544, USA

\*Correspondence: [sebastian@musslick.de](mailto:sebastian@musslick.de) (S. Musslick).



### Box 1. Cognitive control and the mitigation of conflict

Cognitive control can be defined as the collection of mechanisms responsible for flexibly adapting information processing in the service of task goals. While most theories of cognitive control comport with this definition, they differ in their assumptions about how cognitive control is implemented. Traditionally, psychological theory (and most, though not all symbolic architectures used to model cognitive function) has assumed that control relies on a common, centralized set of mechanisms responsible for coordinating task execution (e.g., either goal representations [149] or production rules [2]) that are assumed to have an intrinsic capacity limitation. By contrast, in connectionist systems, control is not assumed to rely on a centralized, dedicated, or necessarily capacity limited set of mechanisms. Rather, it is considered to be a more distributed and functional characteristic of the system, that reflects the influence that one set of units (often representing more abstract forms of information) can have over another set of units (representing more specific forms of information) in a processing pathway, biasing them in a manner that supports execution of a task (see Box 3) [51,83–85,150,151]. In a given context, these sets of units can be viewed, respectively, as supporting ‘control representations’ or ‘task representations’, and process(es) relying on the former as ‘control-dependent’. However, we emphasize that these designations are relative rather than absolute, that depend on the particular tasks being performed. Thus, a representation that serves as control in one context (e.g., of the dimension ‘color’ in the Stroop task) may be a task representation in another (e.g., when reporting whether two objects differ along the dimension of color or shape).

Another point of distinction is that, in connectionist architectures, a fundamental purpose of control is the avoidance of conflict that arises when two tasks compete for the same resource (i.e., task representations; see Figure 1 in main text and Box 3) [51,83,84], by limiting the engagement of shared representations to a single task (see Figure 2 and section ‘Trading the costs and benefits of shared representations’ in main text). This may explain why tasks that rely on ‘general-purpose’ representations (e.g., those required for the processing of language) are so closely associated with control: the very nature of being ‘general-purpose’ means they are accessible to (i.e., shared by) many tasks. The more general a representation, the more potential for conflict it poses among tasks and, thus, the greater the demand for control to mitigate this conflict by limiting processing to use by one task at a time (see Box 4). The avoidance of conflict aligns with another function of cognitive control: the flexible scheduling of task processes over time [40,41,152]. This coordination function plays a crucial role in more complex tasks, such as problem solving or planning, but a specific treatment of those is beyond the scope of this article.

tasks required conflicting representations to be active at the same time. It is thus the sharing of representations that makes tasks control-dependent, forcing people to allocate control to only one such task at a time and requiring them to flexibly switch between tasks in order to achieve more than one. This need to switch between tasks gives rise to a second dilemma, the trade-off between **cognitive stability** and **cognitive flexibility**: Greater allocation of control to one task (e.g., parsing a mathematical equation) results in greater activation of the neural representations needed to perform that task; but, due to persistence of this neural activity, this makes it more difficult to switch to another task (e.g., responding to a phone call) [17–19]. In the second part of this article, we review evidence for the hypothesis that constraints on the amount of control allocated to a single task result from a bias in this trade-off, toward the ability to flexibly switch between tasks. Finally, we discuss how the two dilemmas may account for limitations in other domains of human cognition, such as constraints on working memory and visual attention.

### Constraints on multitasking capacity

One of the key characteristics of cognitive control is a limitation in the number of control-demanding tasks that humans can execute simultaneously. This constraint is intuitively obvious (e.g., the inability to carry out a mathematical calculation while planning a grocery list). This has been defined as a core characteristic of cognitive control [20,21] and used to distinguish control-dependent from automatic processes: the former are assumed to rely on control for execution, and therefore must be performed in serial (i.e., one at a time), whereas the latter do not rely on control and can operate in parallel (i.e., be multitasked). The distinction between serial versus parallel execution is literally paradigmatic. **Dual-task interference** is universally used to operationalize (experimentally ‘diagnose’) control-demanding processes in the laboratory: a task is considered to be control-dependent if it cannot be executed in parallel with another control-dependent task without interference [20–22,148]. Thus, understanding the constraints associated with cognitive control amounts to understanding the inability to execute tasks in

### Glossary

**Cognitive flexibility:** the ability to quickly adapt information processing to changing task demands.

**Cognitive stability:** the ability to successfully pursue a task in the face of distraction.

**Dual-task interference:** decrements in task performance when an attempt is made to execute a task at the same time as another task, relative to when it is performed alone.

**Multitasking:** the simultaneous execution of two or more tasks. We distinguish this use of the term from broader uses, such as the switching between multiple tasks.

**Resource:** a cognitive commodity (e.g., a set of representations) that are required to execute a task. It is often assumed that resources are limited in some way (e.g., that only one or a restricted subset of representations of a given type can be engaged at a single time).

**Task:** a consistent mapping from the features of a stimulus along a given dimension (e.g., the color of a word) to a set of actions along a given response dimension (e.g., verbal responses).

**Task representation:** the full representation of all information (e.g., perceptual, contextual, and response) needed to execute a particular task.

parallel and requires identifying the cause of dual-task interference. A number of paradigms are amenable to the systematic study of dual-task interference, the most prominent of which is the psychological refractory period (PRP) paradigm (Box 2). Others include visual search [23], perceptual discrimination [24], or short-term memory search [25]. However, in the latter paradigms it can be difficult to infer from reaction times (RTs) and accuracy alone whether participants are processing tasks in serial or parallel. This issue has been interpreted as ‘model mimicry’: a parallel processing model can mimic the behavior of a serial processing model, for example, in visual search (see section ‘Visual attention and the binding problem’) [26]. To dissociate serial versus parallel processing in these paradigms, researchers developed mathematical tools relying on cumulative RT distributions [27,28].

The most common explanation for the processing constraints associated with control is that these reflect a limitation intrinsic to the mechanisms responsible for control itself (e.g., reliance on a single, centralized, and capacity-limited control mechanism [20,21,29–31], akin to the central processing unit (CPU, or core) of a traditional computer. However, this analogy is problematic, given the enormous capacity that the brain holds for parallel processing in other domains (e.g., the simultaneous integration of hundreds of visual features into an object, or simply walking

#### Box 2. Measuring dual-task interference in the laboratory

Perhaps the most thoroughly studied marker of dual-task interference is the psychological refractory period (PRP) [153]. In the PRP paradigm, participants are asked to respond as quickly as possible to two tasks within the same trial. Each trial begins with the presentation of a stimulus relevant to the first task (Task 1), followed by an experimentally manipulated delay [the stimulus onset asynchrony (SOA)] and then the stimulus for the second task (Task 2; Figure 1A). The critical observation is that response times to the second stimulus progressively increase as the SOA is reduced [143,153]. This increase in response time is referred to as the PRP (Figure 1B) and commonly interpreted as evidence of a structural bottleneck (see Figure 1A in main text) that delays execution of Task 2 while Task 1 is still being processed [31,32,143]. This bottleneck has been ascribed to the limited capacity of a central control mechanism responsible for executing both tasks. By contrast, connectionist (and some symbolic) models suggest that the PRP reflects strategic adjustments by control in response to processing interference that is induced by task-specific resources (e.g., representations; see Box 3) that are shared by both tasks [44,54]. To prevent ongoing or residual processing of Task 1 from interfering with Task 2, processing of the latter is strategically delayed. It is important to note that sharing need not be directly between the two tasks (i.e., due to structural overlap), but can also arise from indirect sharing (i.e., functional dependence) mediated by a third task, as shown in simulations using connectionist models described in Figure 2C in the main text [54].

A growing number of studies demonstrate that dual-task performance can improve with dual-task practice relative to single task practice [87,154–156]. Such effects can reduce [154,155] or even eliminate [157] the PRP. Proponents of the structural bottleneck hypothesis suggest that, with dual-task practice, two tasks can get compiled into tasks-specific productions, or even a ‘super-task’ processed more efficiently by a central resource [154,158] (for a review, see [156]). However, this fails to explain the observation that improvements with dual-task practice can transfer to other, related tasks [155,159]. The strategic bottleneck hypothesis suggests this results from improvements in the allocation and scheduling of resources [157,160]. By contrast, connectionist models suggest that improvements with dual-task practice result from the separation of resources (i.e., task representations; see Box 3) [54]. The latter conforms with observations that transfer is restricted to similar tasks [161] and that separation of neural representations is predictive of dual-task improvements with practice [87].

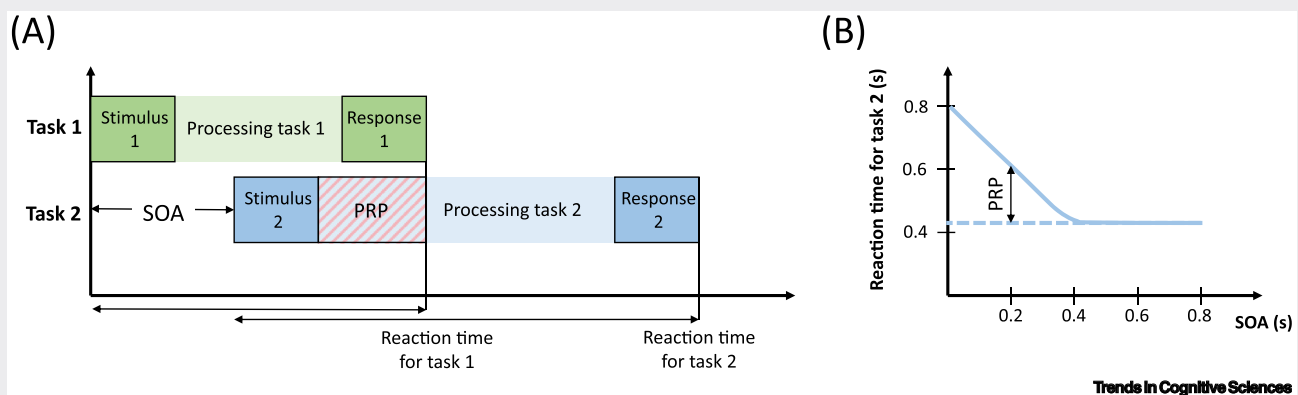


Figure 1. The dual-task interference with the psychological refractory period (PRP). (A) The PRP paradigm. (B) Reaction time (RT) of the second task as a function of stimulus onset asynchrony (SOA). The unbroken line indicates the presence of a psychological refractory period. The broken line serves as a reference for the RT of Task 2 with long SOA. Dual-task practice can eliminate the PRP (cf. broken line).

and talking at the same time). This analogy also seems outdated, since even the most basic of computers now almost always use more than one core. Alternative accounts align more closely with the idea that control reflects a response to, rather than the source of, constraints imposed on processing. In the section that follows, we review primary theories concerning constraints on human multitasking performance and their relationship to cognitive control.

### The costs of shared representation for multitasking performance

Historically, there have been three classes of theories that address the limitations in human multitasking and their relationship to cognitive control (see Table 1 for an alternative classification). All of them assume that these limitations reflect some form of shared **resource(s)** (Box 3), but differ

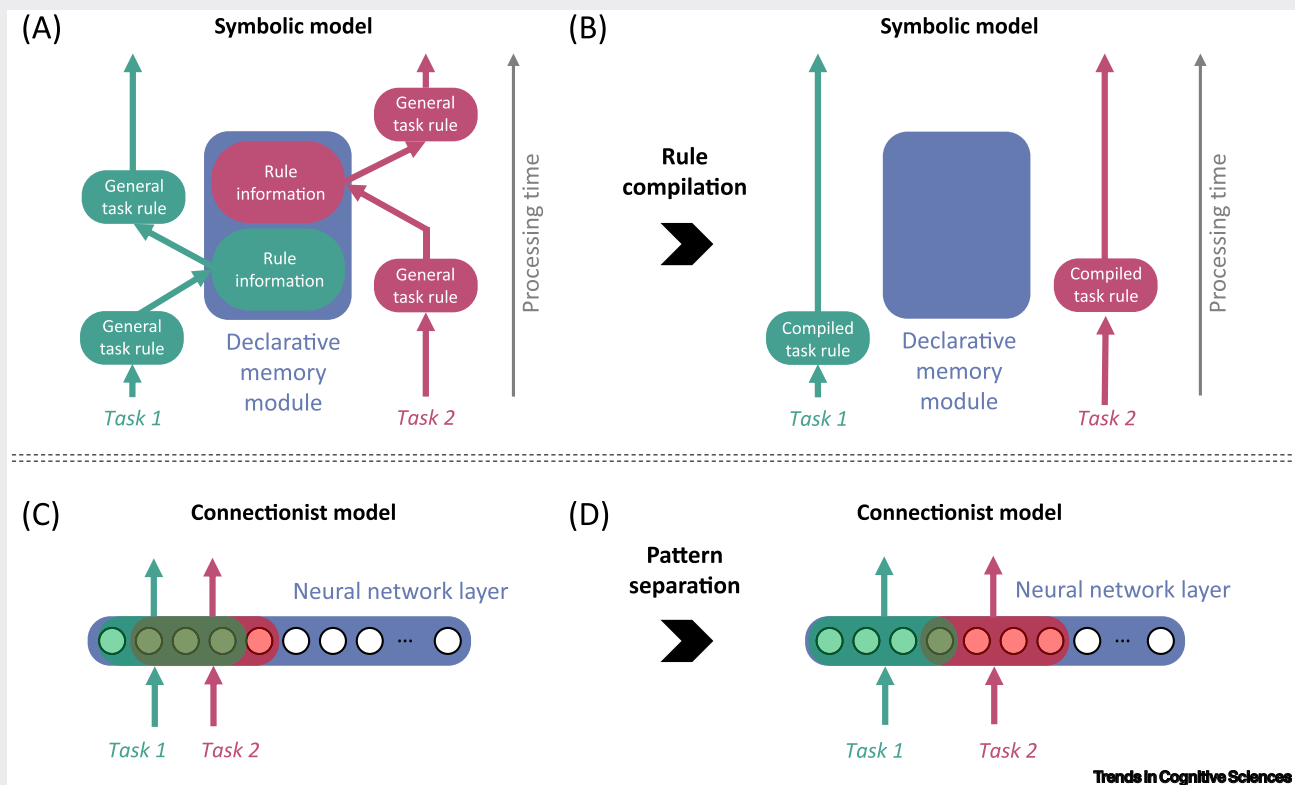
Table 1. Classification of multitasking theories based on their assumptions about resources. In contrast to task-specific resources, central resources are required to operate every task, and can be distinguished by the stage of processing they reside. Resources can be used either by one task (not divisible) or multiple tasks (divisible) at the same time. Check marks indicate assumptions that are stated in the respective article. Circles indicate cases where an assumption is not explicitly stated but is either (1) acknowledged as a possibility or (2) can be derived from other assumptions of the model/theory. Theories are grouped by whether they are expressed in verbal or quantitative form (cf. Box 3 for a comparison of symbolic and connectionist architectures). Note that the present list of references (Refs) is not meant to be complete, but rather representative of the range of theories in the literature.

Article	Number and types of resources				Divisibility of resources		Refs
	Central			Task-specific	Yes	No	
	perception	Response selection	Response initiation				
Verbal theories							
Welford (1952)		✓				✓	[143]
Broadbent (1957)	✓					✓	[32]
Moray (1967)		✓			✓		[144]
Smith (1967)		✓				✓	[145]
Posner and Boies (1971)		✓				✓	[146]
Allport, Antonis, and Reynolds (1972)				✓		✓	[38]
Kahneman (1973)		✓		✓	✓		[30]
Keele (1973)			✓			✓	[35]
Shiffrin and Schneider (1977)		✓				✓	[21]
Kinsbourne and Hicks (1978)				✓		✓	[147]
Navon and Gopher (1979)				✓	✓		[42]
Navon and Miller (1987)				✓	✓		[132]
Wickens (1991)		◦		✓	✓		[43]
DeJong (1993)		✓	✓			✓	[33]
Pashler (1994)		✓				✓	[31]
Mathematical models							
Logan and Gordon (2001)				✓	✓		[95]
Townsend and Wenger (2004)				✓	✓		[27]
Computational models based on symbolic architectures							
Meyer and Kieras (1997)				✓	✓	✓	[40]
Byrne and Anderson (2001)		✓		✓		✓	[39]
Salvucci and Taatgen (2008)		✓		✓		✓	[41]
Computational models based on connectionist architectures							
Feng et al. (2014)		◦		✓	✓		[50]
Musslick et al. (2016, 2020)		◦		✓	✓		[51,66]

**Box 3. Resource sharing in symbolic architectures and connectionist models**

The terms 'shared resource' and 'shared representation' describe similar concepts in different models of human multitasking. Models based on symbolic architectures [39–41], such as ACT-R [2] and EPIC [40,44], consist of different modules (e.g., for processing sensory information, retrieving semantic knowledge, etc.). A module can be considered a 'shared resource' if it is required by two or more tasks. It is generally assumed that if two tasks require the same module (e.g., declarative memory) at the same time, they interfere with one another (Figure IA). However, practice may result in the compilation of task processes (e.g., ones that rely on declarative memory) into specialized, task-dedicated processes that are independent of shared modules [41,162], leading to improvements in multitasking (Figure IB).

In a connectionist model, a module can be thought of as a set of closely interconnected processing units (physical substrates corresponding to individual neurons or populations of neurons) that are used to represent a particular type of information, that is represented as a pattern of activity over those units. Interactions among modules are generally assumed to be more common, direct, fine-grained, and continuous than is assumed in most symbolic architectures. However, as in symbolic models, a module can be considered a 'resource' in that it can support only one pattern of activity and thus represent only one item of information at a time. Thus, interference can arise if a resource is shared by two tasks that require it to represent different (incongruent) information at the same time (Figure IC). However, this raises the question of how the sets of processing units that constitute a resource are defined or identified, making it difficult to determine the extent to which two tasks share resources [56]. Operationally, this can be addressed by correlating the average pattern of activities across all units for pairs of tasks [51,62,66]. This operationalization of resource sharing exploits the graded and distributed nature of representations in connectionist models, by: (i) allowing it to be treated as a continuous rather than all or nothing factor, and (ii) allowing it to change as a consequence of learning through connection modifications. The reliance on shared representational resources initially during learning, that supports more rapid acquisition and generalization, may help explain reliance on control early during acquisition; whereas the progressive development of separate task-dedicated representations may help explain the gradual development of automaticity (i.e., diminished reliance on control) and greater capacity for multitasking that comes with extensive practice [54,66] (Figure ID and see Figure 2D in the main text).

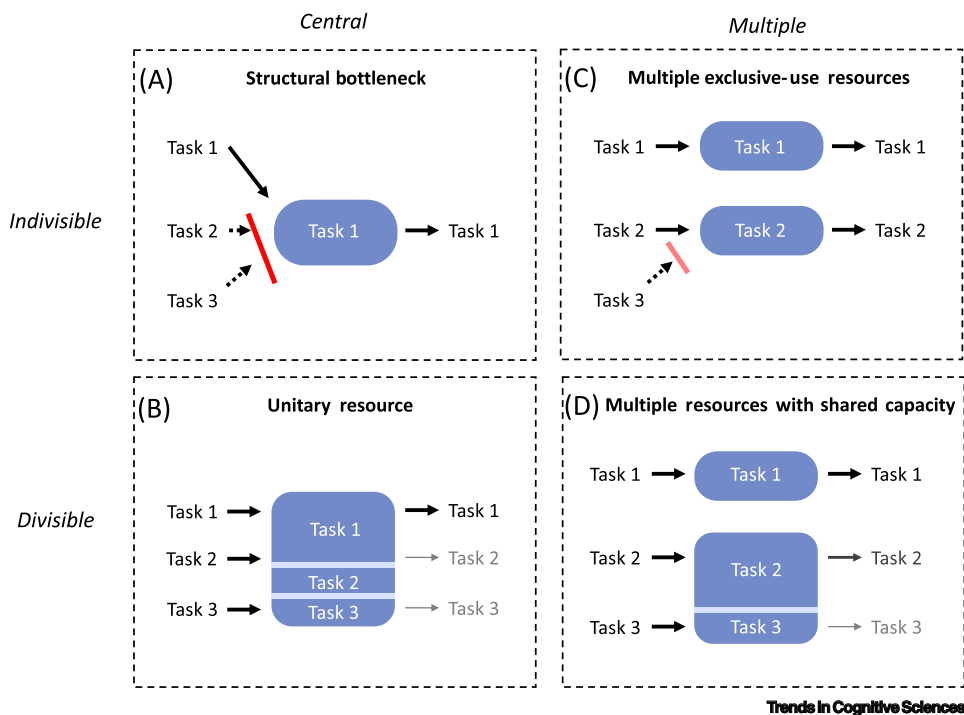


**Figure 1. Resource sharing and separation in symbolic and connectionist models.** (A) Depiction of a resource shared between Task 1 and Task 2 in a symbolic architecture. Both tasks rely on the serial execution of general task production rules the implantation of which requires the retrieval of information from declarative memory (shared resource, shown in purple). Processing of Task 2 is delayed as long as the shared resource is occupied with processing the production rule associated with Task 1. (B) Improvements in multitasking can be achieved by compiling general production rules into specialized, task-dedicated rules that no longer make use of the shared resource [41,162]. (C) In a connectionist model, Task 1 and Task 2 may use some of the same processing units constituting a module (e.g., the three units shown in dark green), leading both tasks to share a representation. As we discuss in the Concluding remarks, this may also occur in neural network layers responsible for encoding declarative memories (cf. [163]), paralleling (A). (D) Improvements in multitasking can be achieved by separating representations between tasks [54,66].

in where this occurs and the capacity for sharing. The first class, 'structural bottleneck theories', assume the existence of a central resource and further assume that this resource can be used for only one task at a time (Figure 1A). The bottleneck is often considered to be 'amodal' insofar as it is required by all tasks, irrespective of their modalities (sensory, motor, or otherwise). Structural bottleneck theories differ in where they locate this central resource, ranging from perceptual processing [32] to response selection [31,33,34] or response initiation [33,35].

The second class of theories, 'unitary resource theories', share the assumption of a central, limited resource. However, they posit that the central resource (sometimes labeled 'attention' [30]) can be divided between tasks (Figure 1B). Thus, unitary resource theories assume that tasks can be executed in parallel, with the caveat that the unitary resource must be divided between them, leading to a trade-off in performance between tasks [30,36,37]. In this sense, they represent a graded version of structural bottleneck theories.

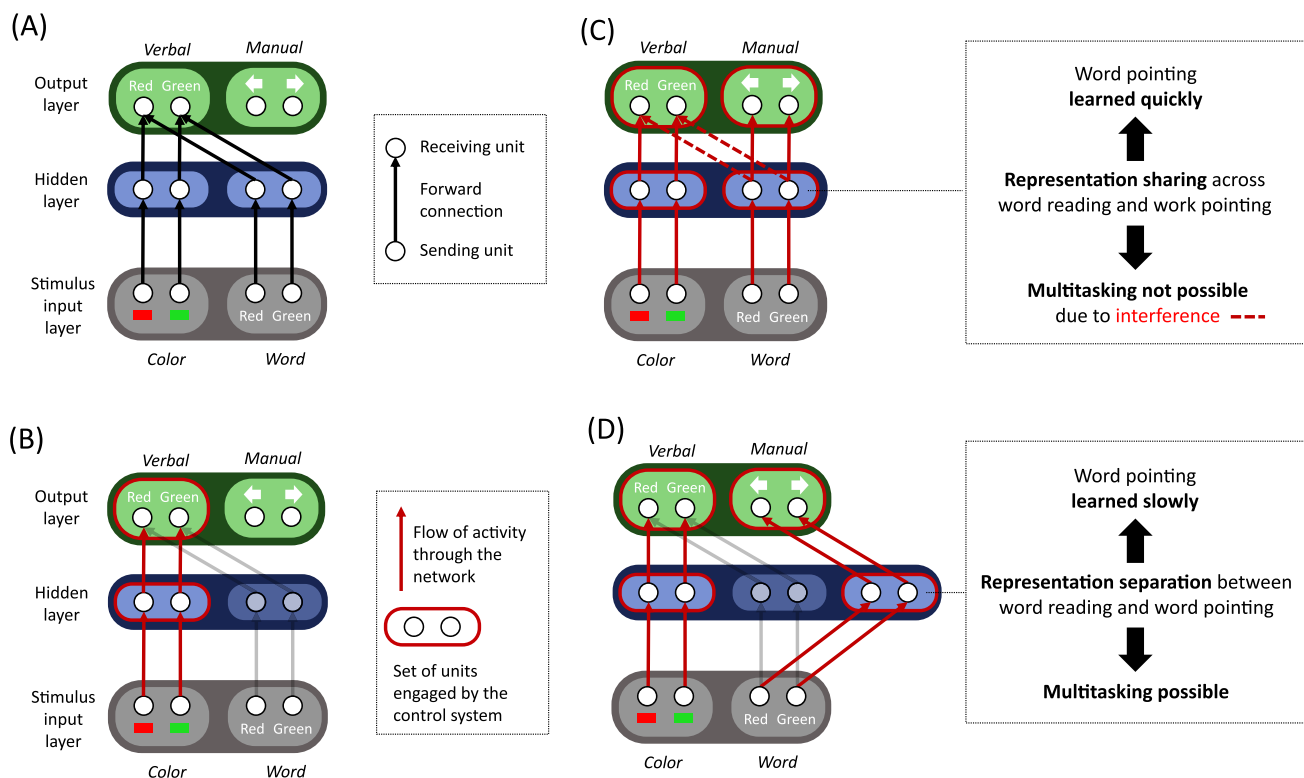
Finally, 'multiple resource theories' renounce the idea of a single, centralized, limiting resource. Instead, they presume a multitude of independent, specialized ('local') resources. A task may demand one or more of such resources in various combinations. Some propose that each



**Figure 1. Classification of resources in theories of human multitasking.** Theories differ according to whether they assume that tasks require the same, central resource, or local task-specific resources (central versus multiple) and the way in which those resources can be allocated (indivisible or divisible). (A) Structural bottleneck. A central resource constitutes a bottleneck in that it is required for execution of all tasks and can operate only one of those at a time; if the resource is engaged by one task, it causes a delay in the processing of others. (B) Unitary resource. Tasks rely on a unitary centralized resource, but it can be allocated to multiple tasks at the same time; task interference occurs if the demands of those tasks exceed the available capacity of the unitary resource. (C) Multiple exclusive-use resources. Tasks rely on local, task-specific resources, each of which can only be used for one task at a time; interference arises if two tasks make simultaneous use of the same resource. (D) Multiple resources with shared capacity. Local, task-specific resources can be shared; interference arises if the capacity of a local resource is exceeded by the number of tasks using it at the same time.



local resource can only be used by one task at a time [38–41] (Figure 1C), whereas others propose that resources can be shared between tasks (Figure 1D), similar in concept to a unitary resource but without assuming that such resources are required by all tasks [42,43]. Multiple-resource theories [41,42,44,45] became increasingly successful in explaining multitasking phenomena in laboratory tasks, such as higher dual-task interference for structurally overlapping tasks [46,47], and in real-world scenarios, such as the effects of phone dialing on speed control while driving [48,49]. In addition, these theories are supported by recent numerical and analytical work, suggesting that even modest amounts of resource sharing between tasks can be sufficient to drastically limit the multitasking capacity of a neural system [50–52] and that this effect scales with the number of processing steps (layers) in the network [53]. Thus, even small amounts of representation sharing are sufficient to induce constraints on multitasking that may invite misinterpretation as a central bottleneck [54]. As illustrated in Figure 2C, connectionist implementations of multiple-resource theory suggest that two tasks can interfere even if they rely on separate representations. They can be ‘functionally’ dependent if the representations on which they rely



**Figure 2. Trade-off between learning efficacy and processing efficiency in a neural network.** Tasks are implemented as mappings between stimulus features (units in the input layer) and responses (units in the output layer), through an internal representation of stimulus features (units in the hidden layer). (A) The network implements two tasks: mapping colors to verbal responses (color naming) and mapping words to verbal responses (word reading). All units are assumed to be inhibited at rest, diminishing incidental flow of activity through the network. (B) To execute color naming (pathway highlights in red), the control mechanism (not shown) engages representations for colors and verbal responses in the hidden and output layers, respectively. This favors color naming, diminishing interference from the word reading task (gray pathway). (C,D) The network may learn a new task, for example, indicate the word with a manual response (word pointing), by mapping the existing representations for words to manual responses (C) or, alternatively, by learning new representations dedicated to mapping from words to manual responses (D). The former requires less time to learn but results in a shared representation between word reading and word pointing and thus dependence on control; when prompted to multitask color naming with word pointing (red pathways), the control mechanism must engage representations for colors and words in the hidden layer, as well as verbal and manual responses in the output layer. In (C), this results in an implicit engagement of the word reading task (broken red pathway) because word reading and word pointing share a representation, leading to interference with color naming. Separated representations in (D) prevent such interference, but take longer to learn. Unpublished work demonstrates that human behavior is consistent with (C), suggesting that humans are biased toward representation sharing early in training [66].

Trends In Cognitive Sciences

can be recombined to form another (interfering) task [51,52]. However, despite growing empirical and quantitative support, multiple-resource theories lack a principled explanation for why a neural system, such as the human brain, would rely on shared resources between tasks at all, given the constraints on multitasking that this imposes.

All three classes of theory outlined above ground their assumptions about resource sharing in empirical evidence. If the simultaneous execution of two tasks reveals dual-task interference then they are likely to share a common resource (whether central or local) [43]. However, to avoid circularity, one needs more than behavioral criteria when deciding about the number and types of resources [55,56]. This theoretical issue is exaggerated in theories that propose multiple resources, as highlighted by Meyer and Kieras (1997, p. 11) [40]: 'One [...] [concern] is that the concept of multiple resources lacks sufficient principled constraints. In the absence of such constraints, there is a temptation to hypothesize new sets of resources whenever additional problematic data are collected. This could lead ultimately to an amorphous potpourri of theoretical concepts without parsimony or predictive power'. This criticism extends to central resource theories: why would a central shared resource exist in the first place and what are the processes invoked by this resource [56]? In neural architectures, two tasks can be described as sharing a resource if they rely on the same set of representations for processing (Box 3). From this perspective, explaining multitasking limitations requires specifying when and why representations might be shared. Computational investigations in the domains of semantic cognition and task acquisition, as well as machine learning, have begun to provide insights into this, most of which identify benefits of representation sharing for learning and generalization. Here, we present a mechanistic implementation of multiple-resource theory in terms of neural network architectures, that provides a quantitatively explicit and potentially normative interpretation of how and why resources may be shared.

#### The benefits of shared representations for learning

The sharing of representations between tasks is evident in many domains of cognition, ranging from visual processing [57] and numerical judgement [58] to language acquisition [59] and social cognition [60]. Yet, little is known about when and why people acquire shared representations across tasks. That is, what are the conditions under which neural systems develop shared representations and what are their benefits for behavior? Answers to these questions arise from the study of category learning in semantic cognition and are exploited in the form of machine learning paradigms for the acquisition of multiple related tasks.

Studies in semantic cognition suggest that shared representations support the learning of meaningful associations between object features. Learning a category can be a challenging task, given that features defining that category may never co-occur. For instance, one never observes that birds fly and lay eggs at the same time; Yet, both features define a bird [61]. To associate features observed in different contexts into a single category, the cognitive system must use a common representation across contexts [61,62]. The same principle applies to the learning of relationships between categories. Computational studies suggest that neural systems are likely to learn shared representations between semantic categories if they are statistically related [63,64]. This idea was recently formalized in a mathematical theory of semantic cognition, demonstrating that neural networks are biased toward learning shared representations between two categories (e.g., trees and flowers) if the set of features defining each of the categories overlap (e.g., trees and flowers have in common that they grow but are not motile) [65]. A similar observation can be made for task acquisition: neural networks are more likely to acquire shared representations between tasks if they overlap in terms of task-relevant stimulus features (e.g., visual features relevant for task performance) [66–68]. These and other findings [61,62,69] suggest that



shared representations are not just a 'byproduct' of learning; they allow networks to learn more rapidly and generalize better by capturing statistical regularities in the world and/or the actions required to act within it (i.e., its affordances). For instance, network architectures that promote the learning of shared representations across sensory modalities facilitate the acquisition of novel concepts [61]. Similarly, architectural biases toward the learning of shared representation between tasks can accelerate the sequential acquisition of these tasks [68,70]. This benefit of representation sharing may also underlie advantageous effects of learning task-sets over individual stimulus-response mappings. When learning a new task (e.g., mapping word stimuli to manual responses), participants' responses can be impacted by task-irrelevant information (e.g., the stimulus color) [71]. The effect of task-irrelevant information diminishes if participants are instructed to group novel stimulus-response mappings into response-relevant categories (task sets, e.g., whether the word describes a moving or nonmoving object), as opposed to memorizing them 'by heart' [72,73]. The benefit of grouping stimuli into categories may result from representation sharing at the level of task-sets: representations for abstract categories confer, by definition, insensitivity to category-irrelevant information.

Machine learning applications have begun to exploit the benefits of shared representations for learning. In that field, 'multitask learning' (as distinct from 'multitasking') refers to settings in which an agent is trained on a set of auxiliary tasks that share representations with the task to be learned, exploiting the fact that similarities among those tasks can lead to shared representations that can help improve generalization and, thereby, acquisition of the target task [74–76]. This has produced significant improvements in computer vision [77], natural language processing [78], and speech recognition [79].

These converging lines of work all indicate that sharing of representations can improve the efficacy of learning (i.e., speed of acquisition and generalization). However, as noted in the previous section, this introduces the potential for interference in multitasking and thus may come at the cost of efficiency of processing [68]. In the next section, we discuss how this trade-off can help rationalize constraints associated with cognitive control, working memory, and visual attention.

#### Trading the costs and benefits of shared representations

The trade-off between learning efficacy and processing efficiency suggests that the benefits and costs of control-dependent processing may be computationally intertwined (Figure 2). The benefits of cognitive control are that it allows novel tasks to be acquired rapidly by exploiting the advantages of shared representations while minimizing the risk of conflict that this introduces [21,80,81]. For instance, participants can quickly learn to map familiar stimuli onto new responses [76,82] (e.g., learning to map word stimuli to button presses). Biologically inspired models of cognitive control suggest that the brain may achieve this by recombining existing representations for task-relevant stimuli and responses [66,67,81,83]. This suggests that novel tasks may rely, at least initially, on representations shared with other tasks and thus may be subject to interference from those tasks (Figure 2C). However, this can be mitigated by the engagement of control, by limiting the engagement of representations shared by multiple tasks to one task at a time [83,84]. From this perspective, constraints on multitasking capability reflect the engagement of control as a rational adaption to resource sharing, rather than an intrinsic limit in control mechanisms themselves [51,66]. Furthermore, these constraints can be considered a cost incurred in exchange for the benefits gained from shared representations: the repurposing of existing representations for new tasks allow these to be acquired rapidly, but at the cost of constraints on multitasking capability and dependence on control [66,85]. This hypothesis is supported by a mathematical analysis of this relationship in linear networks [66,68], as well as artificial agents that optimize the trade-off between learning efficacy and multitasking capability [70,86]. This work suggests that representation sharing can pay off in learning single tasks more quickly, but

at the cost of having to execute them in serial. However, once acquired, repeated multitasking practice can lead to a separation of representations between tasks [87] (Box 3). The separation of representations between tasks may be a crucial ingredient of automaticity in that it frees tasks from interference, thereby reducing their reliance on control (*cf.* Figure 2D) [51,66]. The trade-off between shared versus separated representations may also inform constraints in cognitive functions related to cognitive control, such as working memory and visual attention.

### Working memory

The costs and benefits of sharing may apply at all levels of representation, including task-relevant representations that are maintained for extended periods of time (Box 4). This embraces representations that are held in short-term buffers, such as contextual information, or 'task goals' in working memory. Such abstract representations that influence processing at a broader level (e.g. at the 'task selection' level) are traditionally considered to support control-dependent processing [2,88]. The principle introduced earlier, that multitasking interference arises if two or more tasks require engagement of different representations (*cf.* Box 3), may equally apply to such abstract representations. Indeed, the limited capacity of short-term buffers, such as working memory, is often attributed to interference between the representations they maintain [89–91] (for a review, see [92]).

The role of shared representation for the active maintenance of information is nicely illustrated in a recent network model of visual working memory [89]. That model consists of two layers: (i) a 'sensory network', which is composed of independent subnetworks, each dedicated to represent a visual stimulus; and (ii) a 'random network', which is randomly and reciprocally connected to the sensory network. Representations for visual stimuli in the sensory network lead to corresponding activations in the random network, which then engage, through feedback connections, the same representations in the sensory network. This reciprocal connectivity ensures that representations for visual stimuli are maintained, despite removal of external input (the visual stimulus)

#### Box 4. Mutual exclusivity and persistence of task representations

Both symbolic and connectionist architectures assume that processing components cannot be engaged for different purposes at the same time (Box 3). This mutual exclusivity produces interference among tasks that share resources unless they are regulated by control. In connectionist architectures, this involves the engagement of representations that favor the processing of one task over others (Box 1). This implies that higher level representations favoring competing tasks should themselves be mutually exclusive. The latter can arise in connectionist architectures, as the system learns that simultaneously activating representations that engage tasks sharing lower level representations results in interference and poor performance. That is, the system may learn to implement mutual exclusivity at multiple levels of representation, from more concrete ones (e.g., of competing stimulus features [17,127,164,165]) to more abstract ones (e.g., corresponding to task goals [91,119–121,166]).

Mutual exclusivity associated with representational sharing, together with the graded nature of representations in connectionist models [107], can also interact with the persistence of representations to produce a trade-off between stability and flexibility of control [119,120]. On the one hand, strong mutual exclusivity and persistence favor robustness of the current representation (i.e., state of activity) and, thereby, the influence of control on the execution of any task that depends on it. However, for the same reasons, this makes it more difficult to switch to another representation (and corresponding task), owing to the greater persistence of the current representation and/or the greater extent to which it has suppressed competing representations due to mutual exclusivity. Conversely, weaker mutual exclusivity and/or persistence compromise robustness. On the other hand, this confers greater flexibility, by making it easier to switch between tasks. Thus, persistence coupled with mutual exclusivity of resources between tasks (a consequence of shared representations in neural architectures) induces a trade-off in the ability of control to confer cognitive stability (strong activation of a task-relevant representation supporting high performance) and cognitive flexibility (the ability to rapidly switch between different task representations). One solution to this trade-off may be the flexible 'gating' of task representations into the system, by dynamically reducing persistence once a task switch is required [150,167]. However, this requires additional mechanisms to optimize the persistence of individual representations over time [168] and appropriate learning mechanisms to implement it in an effective form, an important current direction of research in both natural and artificial systems [70,117].

to the sensory network. The random connectivity provides the network with the flexibility to store any arbitrary item. However, the downside of this random connectivity are shared representations in the random network between stimuli from different sensory subnetworks. That is, there need not be any constraint on the number of units in (i.e., the structural capacity of) the random network to induce capacity constraints; the latter may arise simply from the fact that, once two stimuli share a representation in the random network they may interfere with one another, thereby limiting working memory capacity. From this perspective, working memory can be viewed as an extreme example of a set of representations shared by any task that requires sustained activity of a stimulus and is in this respect subject to the same limitations, for the same reasons, as any other control-dependent process. Further on (*cf.* 'Concluding remarks and future directions'), we present a similar view of mechanisms responsible for episodic memory.

### *Visual attention and the binding problem*

An important consequence of shared representations in the visual system is the 'binding problem', that has been closely linked to the role of attention. The binding problem concerns the assignment of features to individual objects [93] and arises if different objects are represented by the same (shared) set of feature representations. For example, in a display containing a red car and gray house, if the two objects engage the same set of color representations (i.e., use the same population of neurons to represent color), then it may be impossible to determine which one activated red and which one gray. That is, the sharing of representations of features, sometimes referred to as 'compositional coding' [94], poses the risk of confusion if more than one object must be identified at the same time [55]. This is exactly analogous to the problem posed by the sharing of representations between tasks [95]. Figure 2 shows two ways in which the system can deal with this problem in the task domain. If tasks share representations then control must be used to limit processing to only one of them at a time (Figure 2C). However, the system may also, at the expense of additional learning, commit separated, task-dedicated representations to the mappings from the stimuli to the responses for each task (Figure 2D).

These two solutions correspond directly to ones that have been proposed to solve the binding problem in the visual system: use composition coding [94] and limit processing by allocating attention to only one object at a time so that only the features associated with that object are active [55,95]; or dedicate separate representations to the combinations of features for each object ('conjunctive coding' [96]), thus binding those features directly to the object. The latter is observed for highly familiar objects with combinations of features [97] (e.g., facial features of a familiar person). Individuals with autism spectrum disorder (ASD) have been hypothesized to exhibit relatively greater reliance on separated than shared representations [98]. Consistent with the aforementioned trade-off, individuals with ASD outperform individuals without autism in visual search tasks [99,100] but, at the same time, are less likely to abstract and synthesize information across experiences [101].

There is evidence that, despite the risk of interference, the visual system also makes use of compositional coding (i.e., shared representations) of features such as line segments and letters [102–104]. For instance, compositional neural codes of words based on neural responses to letters have been found to explain behavioral and neural correlates of visual search [104]. It is assumed that compositional coding supports spatial invariance (e.g., the ability to detect a letter with the same population of neurons, irrespective of its location). As noted earlier, this approach has been used in machine learning to improve efficacy of learning (e.g., to achieve object recognition) [105,106]. Findings from classic studies of visual attention also provide strong support for the use of compositional coding, evidenced by the binding problem that it poses, and the serialization of processing used as a solution [55]. In such studies, participants must detect an object composed of an arbitrary combination of features (e.g., a red T) in a field of

distractor objects that combine the same features in other ways (e.g., green Ts and red Ls). Errors in such tasks often involve confusion over the assignment of features to objects; and RTs typically scale linearly with the number of distractors, which is taken as evidence of serial allocation of visual attention to each object individually. These observations inspired ‘feature integration theory’, which proposed that the purpose of attention is to integrate percepts by serially binding objects and their features [55], though the mechanism for this process was not specified.

An elaboration of this view is that binding is not implemented directly by attentional mechanisms, but rather by the connections in the network that link objects with compositional representations of their features (i.e., ones that are shared across objects); and that the role of attention is simply to insure that only one object is represented at a time, to prevent the confusions that would arise if the features of different objects were simultaneously represented over the same set of representations. This directly parallels the role of control in averting conflict among tasks that share representations, and constraints in the number of representations that can be actively maintained in working memory. The appeal of this view is that it provides a unified explanation of the conditions under which the constraints of seriality arise and why these are associated with the engagement of control (whether to select what to attend visually, represent in working memory, or task to perform), all in terms of a common set of fundamental principles of representation and processing in neural systems [107]. All reflect a trade-off that favors the advantages of shared (i.e., compositional) representations, whether for efficacy of learning, flexibility of processing, and/or efficiency of representation, at the expense of the efficiency of execution afforded by the ability to process separated (e.g., conjunctive) representations in parallel.

### Constraints on control allocated to a single task

Models of cognitive control imply that higher amounts of control (e.g., increasing the activity of relevant task representations; cf. Box 1) allocated to a single task lead to higher cognitive stability [7, 10–12, 83]. Yet, participants generally refrain from allocating control in a cognitive task unless they are rewarded to do so [11, 108]. Thus, there appears to be a constraint on how much control participants are willing to allocate [8]. This constraint seems puzzling from a rational perspective: Why would a system refrain from allocating maximal control to a task to which it is already committed? One possible explanation is that there are opportunity costs associated with the dynamics of allocating control: the greater the allocation of control to one task, the harder it is to switch to others. Such dynamics can arise from competitive interactions between representations required to allocate control in just the way they do for any other set of representations (Box 4) and may help explain the stability–flexibility dilemma that has been described for control and constraints observed on the intensity of its allocation.

#### The stability–flexibility dilemma

Successful goal-directed action requires balancing antagonistic demands. On the one hand, humans need to maintain and protect task goals in the face of distraction (cognitive stability); on the other hand, they require quick and flexible reconfiguration to perform a different task when the environment changes (cognitive flexibility) [17–19, 109]. Cognitive stability can be quantified in the laboratory, by instructing participants to perform a single task in the presence of distractors and by measuring how much the distractors affect participants’ performance (Box 5). Cognitive flexibility can be assessed by measuring how fast participants can switch from one task to another. Critically, there appears to be a trade-off between the two quantities across participants. On the one hand, individuals with greater flexibility tend to be more distracted by task-irrelevant information [110, 111]; on the other hand, individuals with high resistance to distraction tend to be cognitively inflexible [112, 113]. In addition, reward and positive affect can bias individuals toward greater flexibility, at the expense of cognitive stability [114, 115]. A growing

Box 5. Measuring cognitive stability and flexibility in the laboratory

There are multiple ways to measure cognitive stability and flexibility. Here, we describe how both can be operationalized in tasks that demand cognitive control (also see Table 1 in [126]). Perhaps the most common performance metric used to characterize cognitive stability is response-interference. A classic example of this is the Stroop task [169], in which participants have to name the color in which a word is displayed. Participants are more error-prone and slower to respond if the task-relevant feature (the color) and the task-irrelevant feature (the word) are associated with different (incongruent) responses (e.g., say 'red' in response to the word 'GREEN' colored in red; Figure 1), compared with trials in which the responses associated with both features are congruent (e.g., the word 'RED' colored in red; Figure 1). Cognitive stability can be assessed as the difference in error rate and/or reaction time (RT) between incongruent and congruent trials: The larger the congruency effect for an individual, the lower their cognitive stability. Cognitive flexibility is commonly assessed in terms of an individual's ability to switch from one task to another (e.g., switching between naming the color and reading the word in the Stroop task). In task switching paradigms (Figure 1), participants are asked either to repeat the task they performed on the previous trial (repetition trials) or to perform a different task (switch trial). The switch cost for an individual can be quantified as the difference in error rate and/or RT on switch trials versus repeat trials. The switch cost is taken as an inverse measure of cognitive flexibility: the less the switch cost, the greater the cognitive flexibility. Finally, unpublished work suggests that the stability–flexibility trade-off can also be assessed in working memory tasks, in which cognitive stability is measured as the ability to actively maintain items in the presence of distractors and cognitive flexibility is measured as the ability to encode and maintain new items [170].

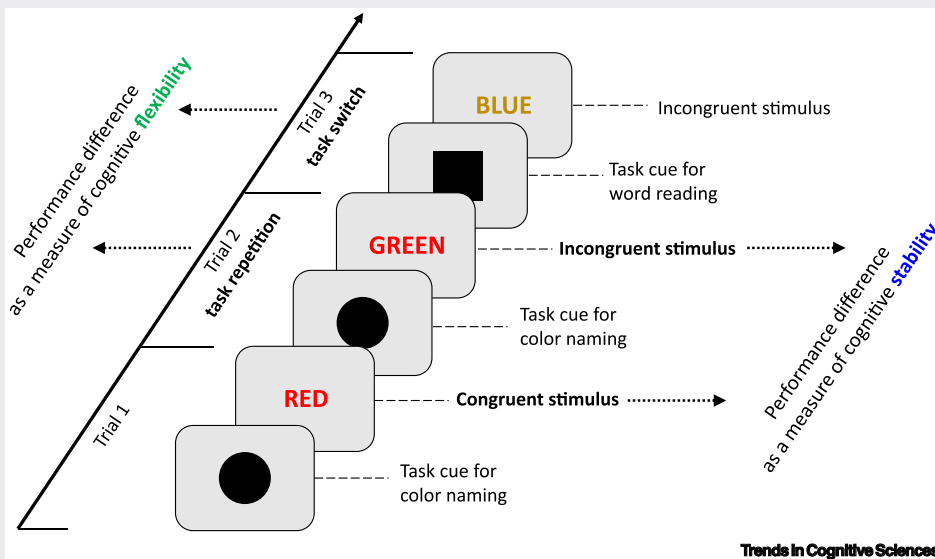


Figure 1. Cued task switching paradigm. Each trial consists of a task cue (e.g., either a circle or a square) that instructs participants which task to perform (e.g., either color naming or word reading), followed by the stimulus to which they must respond.

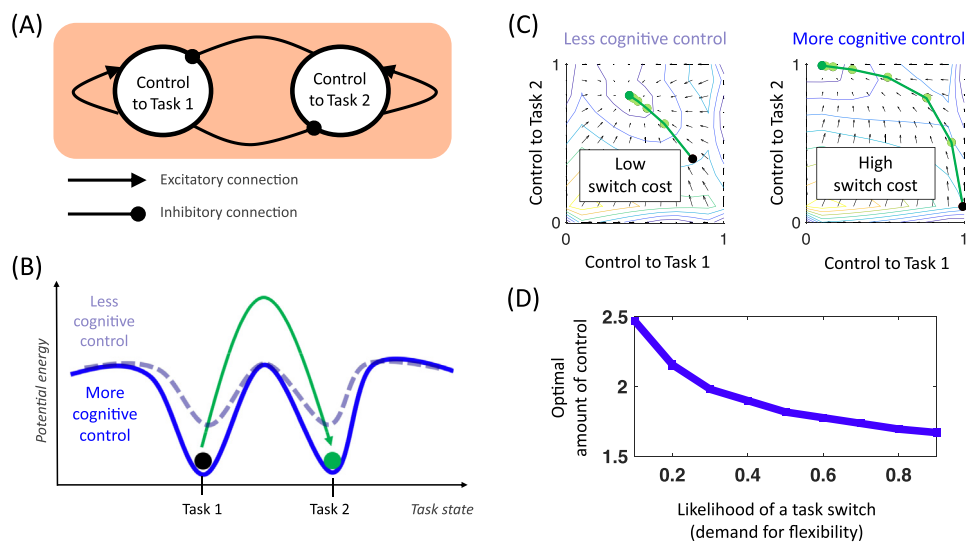
body of work suggests that the neuromodulatory transmitters dopamine and norepinephrine play a role in regulating this balance (for comprehensive reviews, see [116–118]).

Optimizing the balance between cognitive stability and flexibility

Computational analyses of the trade-off between cognitive stability and flexibility suggest that constraints on control may help optimize this trade-off. A number of biophysically inspired models of task control [91, 119–121] rely on the assumptions that: (i) representations for different tasks compete with one another (due to representation sharing); and that (ii) the representation of a task can persist in time, slowing representational reconfiguration for a subsequent task (Box 4). These models describe competing task configurations as different minima (or wells) in an energy landscape (Figure 3A,B). In these landscapes, the information that specifies which task to perform (e.g., task instructions) is assumed to be represented as stable patterns of neural activity that correspond to states with low energy, located at the bottom of an energy well in the

landscape (an ‘attractor’). When such task information is presented to the network, its pattern of activity evolves such that the system moves to that attractor (analogous to a ball rolling along a surface to the bottom of the nearest well). If the representation of the task information corresponds to a deep attractor, then even with small perturbations (e.g., due to noise) the system is most likely to settle back to the same state (akin to a ball bouncing around in a deep well). Thus, deep attractors make the system robust to noise. Conversely, shallow attractors make the system more susceptible to noise (i.e., make it easier for the ball to pop out of the well), but also make it easier to switch from one state to another.

Restrictions on the depth of attractors for task representations, implementing constraints on control intensity, can promote flexible task switching but come at the expense of robustness to distractors (Figure 3C) [120]. Simulation work suggests that higher constraints on control allocation (shallower attractors) yield a higher reward rate in environments with higher demand for (e.g., greater frequency of) task switches (Figure 3D). Furthermore, the behavior of participants in environments with a high rate of task switches can be best explained with higher constraints on control, compared with participants in environments with a low rate of task switches [121]. This is in line with a growing number of studies showing that participants shift their balance to favor cognitive flexibility over stability if task switches become more likely [122–125] (for a review, see [126]). Together, these computational and empirical results suggest that it can be useful to limit the amount of control allocated to a single task, given that this facilitates flexible switching between tasks.



**Figure 3. Modeling the stability–flexibility trade-off.** (A) Two-unit network used to illustrate the stability–flexibility trade-off [119–121]. Each unit represents a control signal for one of two tasks. The activity of each unit corresponds to the amount of control allocated to the corresponding task, that is determined by a recurrent excitatory input from itself and an inhibitory input from the other unit, as well as external input such as a task cue (not shown). (B) The network implements a dynamical system the state of activity of which (x-axis) is determined by its energy (y-axis; cf. [17,119,120,127]). The system has two stable states (attractors), one for executing each of the two tasks. The network’s parameters determine the depth of the attractors. Deep and shallow attractors correspond to networks with high and low amounts of control allocated to each task, respectively. Thus, deep attractors implement cognitive stability due to stronger activation of the control representation; whereas shallow attractors implement greater flexibility, making it easier to switch from one state to another (green arrow). (C) Simulated activation trajectories [120] for shallow (left) and deep (right) attractors are shown as a series of connected light green dots, evolving from the control attractor for Task 1 (black) to the control attractor for Task 2 (green). Contour lines and arrows indicate the energy and shape of the attractor landscape after a task switch from Task 1 to Task 2. With more control allocated to Task 1, the network requires more time steps to switch to Task 2. (D) Simulations show that the amount of control allocated to a single task yielding the highest reward rate decreases with the frequency of task switches [120].

Trends in Cognitive Sciences



### Concluding remarks and future directions

Constraints on our capacity for cognitive control pervade all forms of human cognition that involve control-dependent processes, ranging from lower-level processes such as perception, memory retrieval, and action selection, to higher level ones such as reasoning, problem solving, and language. Here, we considered constraints on: (i) the number of control-demanding tasks that humans can execute simultaneously, as well as (ii) the amount of control that they allocate to a single task, and reviewed recent work that suggests why the allocation of control might be subject to these constraints, in terms of a common set of principles concerning the representation and processing of information in neural network architectures (Boxes 3 and 4). The literature reviewed suggests two fundamental computational dilemmas that arise from these principles. The first is reflected in recent insights from studies of category learning and task acquisition, which suggest a trade-off between learning efficacy and processing efficiency. The former is promoted by the sharing of representations across tasks, whereas the latter is achieved by separating representations between tasks thereby permitting parallel processing (i.e., multitasking). Findings concerning working memory and visual attention suggest a similar trade-off between the representational efficiency of compositional coding using shared representations but at the expense of serial processing, and the efficiency of conjunctive coding that affords parallel processing but at the expense of dedicating separate representation to each item that must be represented. From this perspective, capacity constraints and the requirement for serial processing arise not from limitations intrinsic to the mechanisms responsible for control, but rather from the use of shared representations (whether for features of visual objects, information actively maintained in working memory, or the mappings required to perform a task) that demand the allocation of control in order to avoid interference or confusion that would arise from their use for different purposes at the same time. Furthermore, the competitive dynamics within a set of shared representations in a network can simultaneously account for the trade-off between cognitive stability and flexibility and constraints in the intensity of control allocated to any single task that arise because of this.

The principles of representation and processing reviewed in this article suggest a fundamental relationship between the behavioral and neural correlates of learning, multitasking, and task switching. A key prediction pertains to the use of shared representations early in task acquisition. The extent to which novel tasks are control-dependent (i.e., subject to dual-task interference) should depend on the extent to which they share representations. Furthermore, this should determine performance costs associated not only with dual-tasking [such as the psychological refractory period (PRP); Box 2] but also task switching (Box 5). To our knowledge, there has not yet been a direct empirical test of this prediction, although the effects of structural overlap between tasks in both dual-task and task switching paradigms suggest such a relationship [46,128–130]. This prediction extends from structural overlap on which most previous instances of multiple resource theory have focused [131], to the case of functional dependence addressed by crosstalk models [27,47,51,52,132]. Even if two tasks do not directly share representations (e.g., color naming and word pointing in Figure 2C), they may still interfere with one another by means of another process (e.g., word reading). In addition, sharing may involve either task-specific representations (as in Figure 2C), or more general-purpose representations (such as language [59]) and processes (such as cue retrieval from episodic memory [133]). The idea that representations in episodic memory constitute a generally shared resource, reliance on which requires the engagement of control, may provide a link to symbolic architectures, such as ACT-R (Box 3), in which it is a central feature of declarative memory that is considered as a ‘bottleneck’ in processing. The identification of neural substrates shared by multiple tasks, both directly and indirectly, may ultimately help identify the specific substrates of shared resources, including more general sources of dual-task interference [134,135]. Neural measurements with sufficient temporal resolution (e.g., electroencephalography and magnetoencephalography)

### Outstanding questions

How does the brain assess the risk of interference in the presence of shared representation, and which biophysical signals does it rely on to guide interference-free allocation of cognitive control? Moreover, how does it balance the benefits and costs of shared versus separate representation over its lifetime?

What are the neural mechanisms that underlie improvements in our ability to multitask? Symbolic models suggest that multitasking improvements result from more efficient use of task representations and sophisticated scheduling of tasks, whereas connectionist models suggest that multitasking improvements can be achieved by separating representations between tasks. Neurophysiological and neuroimaging studies of multitask training may be useful in testing these hypotheses.

Are there meta-control mechanisms for balancing the trade-off between cognitive stability and flexibility, and what are their computational and neural underpinnings? Humans have been shown to adapt to changing demands for cognitive flexibility. Yet, it is still unknown what information they use to regulate this balance, how they make relevant adjustments, and over what timescale they are able to do so.

Are there computational principles that motivate constraints on the duration of control allocation? This article focused on constraints in parallel processing (e.g., the number of tasks that can be performed simultaneously) and in the intensity of control allocation (i.e., to individual tasks). However, another universal observation is that control-dependent tasks are experienced as more effortful and therefore more readily subject to fatigue, that can be thought of as the loss of intensity over time. To what extent does this also reflect rational adaptations to fundamental computational dilemmas (such as the explore–exploit trade-off) as opposed to lower level, physical factors such as energetic or metabolic constraints?

may also permit the measurement of persistence characteristics of neural representations at such loci, that should predict the magnitude and timing of the PRP and adaptations in the stability–flexibility trade-off during task switching. Finally, another key prediction of the presented framework is that improvements in dual-task performance depend on the separation of representations between tasks with practice, a relationship that has been suggested by recent neuroimaging work [136].

This review has focused on two trade-offs: learning efficacy versus processing efficiency, and stability versus flexibility; however, there are of course others faced by cognitive systems [18,137,138]. For example, another one concerns the decision between selecting actions that yield known rewards (exploitation) and ones with unknown rewards but that may yield new information that leads to greater rewards in the future (exploration) [139]. It has been hypothesized that boredom may reflect an adaptive signal meant to manage the explore–exploit dilemma, by leading agents to explore new options when current ones are highly predictable [140–142]. A promising avenue for future research is to explore how other limitations on human cognition, such as the inability to exert control over extended periods of time, may result from such trade-offs (see [Outstanding questions](#)).

Constraints of cognitive control remain a crucial building block for general theories of cognition. Thus, an improved understanding of these constraints may help understand their consequences in other domains of cognition. The study of computational dilemmas, as reviewed here, presents a novel approach to this matter and may help rationalize these constraints in neural systems and relate them to the broad inventory of phenomena associated with cognitive function.

### Acknowledgments

This project was made possible through the support of a grant from the John Templeton Foundation. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation. We thank Timothy T. Rogers, Niels A. Taatgen, Bernhard Hommel, James T. Townsend, Gordon Logan, Ivan Grahek, Markus Spitzer, and Annika Geisemeyer for their valuable feedback on this manuscript.

### Declaration of interests

No interests are declared.

### References

- Cohen, J.D. (2017) Cognitive control: core constructs and current considerations. *Wiley Handb. Cognit. Control* 1–28
- Anderson, J.R. and Lebiere, C.J. (2014) *The atomic components of thought*, Psychology Press
- Gigerenzer, G. (2008) Why heuristics work. *Perspect. Psychol. Sci.* 3, 20–29
- Kurzban, R. et al. (2013) An opportunity cost model of subjective effort and task performance. *Behav. Brain Sci.* 36, 661–679
- Lieder, F. and Griffiths, T.L. (2015) When to use which heuristic: A rational solution to the strategy selection problem. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society*, Cognitive Science Society
- Simon, H. (1957) *Models of man; social and rational*, Wiley
- Shenhav, A. et al. (2013) The expected value of control: an integrative theory of anterior cingulate cortex function. *Neuron* 79, 217–240
- Botvinick, M. and Braver, T. (2015) Motivation and cognitive control: from behavior to neural mechanism. *Annu. Rev. Psychol.* 66, 83–113
- Kool, W. et al. (2010) Decision making and the avoidance of cognitive demand. *J. Exp. Psychol. Gen.* 139, 665–682
- Musslick, S. et al. (2015) A computational model of control allocation based on the expected value of control. In *Reinforcement Learning and Decision Making Conference*
- Manohar, S.G. et al. (2015) Reward pays the cost of noise reduction in motor and cognitive control. *Curr. Biol.* 25, 1707–1716
- Lieder, F. et al. (2018) Rational metareasoning and the plasticity of cognitive control. *PLoS Comput. Biol.* 14, e1006043
- Grahek, I. et al. (2020) A computational perspective on the roles of affect in cognitive control. *Int. J. Psychophysiol.* 151, 25–34
- Ho, M.K. et al. (2020) The efficiency of human cognition reflects planned information processing. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence*
- Kool, W. et al. (2017) Cost-benefit arbitration between multiple reinforcement-learning systems. *Psychol. Sci.* 28, 1321–1333
- Grahek, I. et al. (2019) Motivation and cognitive control in depression. *Neurosci. Biobehav. Rev.* 102, 371–381
- Durstewitz, D. and Seamans, J.K. (2008) The dual-state theory of prefrontal cortex dopamine function with relevance to catechol-o-methyltransferase genotypes and schizophrenia. *Biol. Psychiatry* 64, 739–749
- Durstewitz, D. (2013) Volition in Action: Intentions, Control Dilemmas, and the Dynamic Regulation of Cognitive Control. In *Action Science: Foundations of an Emerging Discipline*, pp. 409–434, The MIT Press
- Hommel, B. (2015) *Chapter Two - Between Persistence and Flexibility: The Yin and Yang of Action Control, Vol. 2 of Advances in Motivation Science*, Elsevier, pp. 33–67

20. Posner, M.I. and Snyder, C. (1975) *Attention and cognitive control. Information processing and cognition: The Loyola symposium*, pp. 55–85
21. Shiffrin, R.M. and Schneider, W. (1977) Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychol. Rev.* 84, 127–190
22. Lavie, N. *et al.* (2004) Load theory of selective attention and cognitive control. *J. Exp. Psychol. Gen.* 133, 339–354
23. Fifić, M. *et al.* (2008) Studying visual search using systems factorial methodology with target–distractor similarity as the factor. *Percept. Psychophys.* 70, 583–603
24. Egeth, H.E. (1966) Parallel versus serial processes in multidimensional stimulus discrimination. *Percept. Psychophys.* 1, 245–252
25. Sternberg, S. (1966) High-speed scanning in human memory. *Science* 153, 652–654
26. Townsend, J.T. (1972) Some results concerning the identifiability of parallel and serial processes. *Br. J. Math. Stat. Psychol.* 25, 168–199
27. Townsend, J.T. and Wenger, M.J. (2004) A theory of interactive parallel processing: new capacity measures and predictions for a response time inequality series. *Psychol. Rev.* 111, 1003–1035
28. Townsend, J.T. and Fifić, M. (2004) Parallel versus serial processing and individual differences in high-speed search in human memory. *Percept. Psychophys.* 66, 953–962
29. Anderson, J.R. *et al.* (2004) An integrated theory of the mind. *Psychol. Rev.* 111, 1036–1060
30. Kahneman, D. (1973) *Attention and effort*, Vol. 1063. Prentice-Hall, Englewood Cliffs, N.J.
31. Pashler, H. (1994) Dual-task interference in simple tasks: data and theory. *Psychol. Bull.* 116, 220–244
32. Broadbent, D.E. (1957) A mechanical model for human attention and immediate memory. *Psychol. Rev.* 64, 205–215
33. De Jong, R. (1993) Multiple bottlenecks in overlapping task performance. *J. Exp. Psychol. Hum. Percept. Perform.* 19, 965–980
34. Welford, A.T. (1967) Single-channel operation in the brain. *Acta Psychol.* 27, 5–22
35. Keele, S.W. (1973) *Attention and human performance*, Goodyear Publishing Company
36. Norman, D.A. and Bobrow, D.G. (1975) On data-limited and resource-limited processes. *Cogn. Psychol.* 7, 44–64
37. Tombu, M. and Jolicoeur, P. (2003) A central capacity sharing model of dual-task performance. *J. Exp. Psychol. Hum. Percept. Perform.* 29, 3–18
38. Allport, A. *et al.* (1972) On the division of attention: A disproof of the single channel hypothesis. *Q. J. Exp. Psychol.* 24, 225–235
39. Byrne, M.D. and Anderson, J.R. (2001) Serial modules in parallel: The psychological refractory period and perfect time-sharing. *Psychol. Rev.* 108, 847–869
40. Meyer, D.E. and Kieras, D.E. (1997) A computational theory of executive cognitive processes and multiple-task performance: Part I. Basic mechanisms. *Psychol. Rev.* 104, 3–65
41. Salvucci, D.D. and Taatgen, N.A. (2008) Threaded cognition: An integrated theory of concurrent multitasking. *Psychol. Rev.* 115, 101–130
42. Navon, D. and Gopher, D. (1979) On the economy of the human-processing system. *Psychol. Rev.* 86, 214–255
43. Wickens, C.D. (1991) Processing resources and attention. In *Multiple Task Performance*, pp. 3–34, Talor & Francis, Ltd., Bristol
44. Meyer, D.E. and Kieras, D.E. (1997) A computational theory of executive cognitive processes and multiple-task performance: Part II. Accounts of psychological refractory-period phenomena. *Psychol. Rev.* 104, 749–791
45. Wickens, C.D. *et al.* (1980) The structure of processing resources: Implications for task configuration and workload. In *Proceedings of the Human Factors Society Annual Meeting* (Vol. 24), pp. 253–256, SAGE Publications Sage CA, Los Angeles, CA
46. Göthe, K. *et al.* (2016) Eliminating dual-task costs by minimizing crosstalk between tasks: The role of modality and feature pairings. *Cognition* 150, 92–108
47. Koch, I. (2009) The role of crosstalk in dual-task performance: evidence from manipulating response-code overlap. *Psychol. Res. PRPF* 73, 417–424
48. Brumby, D.P. *et al.* (2009) Focus on driving: How cognitive constraints shape the adaptation of strategy when dialing while driving. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1629–1638
49. Salvucci, D.D. and Macuga, K.L. (2002) Predicting the effects of cellular-phone dialing on driver performance. *Cogn. Syst. Res.* 3, 95–102
50. Feng, S.F. *et al.* (2014) Multitasking versus multiplexing: Toward a normative account of limitations in the simultaneous execution of control-demanding behaviors. *Cogn. Affect. Behav. Neurosci.* 14, 129–146
51. Musslick, S. *et al.* (2016) Controlled vs. automatic processing: A graph-theoretic approach to the analysis of serial vs. parallel processing in neural network architectures. In *Proceedings of the 38th Annual Meeting of the Cognitive Science Society, Philadelphia, PA*, pp. 1547–1552
52. Petri, G. *et al.* (2021) Topological limits to the parallel processing capability of network architectures. *Nat. Phys.* 17, 646–651
53. Alon, N. *et al.* (2017) A graph-theoretic approach to multitasking: advances in neural information processing systems. In *Advances in Neural Information Processing Systems, Long Beach, CA*, pp. 2097–2106
54. Musslick, S. and Cohen, J.D. (2019) A mechanistic account of constraints on control-dependent processing: Shared representation, conflict and persistence. In *Proceedings of the 41st Annual Meeting of the Cognitive Science Society, Montreal, CA*, pp. 849–855
55. Treisman, A.M. and Gelade, G. (1980) A feature-integration theory of attention. *Cogn. Psychol.* 12, 97–136
56. Hommel, B. (2020) Dual-task performance: Theoretical analysis and an event-coding account. *J. Cogn.* 3, 1–13
57. Albers, A.M. *et al.* (2013) Shared representations for working memory and mental imagery in early visual cortex. *Curr. Biol.* 23, 1427–1431
58. Notebaert, W. *et al.* (2006) Shared spatial representations for numbers and space: the reversal of the snarc and the simon effects. *J. Exp. Psychol. Hum. Percept. Perform.* 32, 1197–1207
59. Weber, K. *et al.* (2016) fMRI syntactic and lexical repetition effects reveal the initial stages of learning a new language. *J. Neurosci.* 36, 6872–6880
60. Decety, J. and Sommerville, J.A. (2003) Shared representations between self and other: a social cognitive neuroscience view. *Trends Cogn. Sci.* 7, 527–533
61. Jackson, R.L. *et al.* (2021) Reverse-engineering the cortical architecture for controlled semantic cognition. *Nat. Hum. Behav.* 5, 774–786
62. Rogers, T.T. and McClelland, J.L. (2004) *Semantic cognition: A parallel distributed processing approach*, MIT press
63. Hinton, G.E. (1989) Learning distributed representations of concepts. In *Parallel distributed processing: Implications for psychology and neurobiology* (Vol. 1), pp. 46–61, Clarendon Press/Oxford University Press
64. Rumelhart, D.E. *et al.* (1993) *Learning and connectionist representations*. Attention and performance XIV: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience 2 pp. 3–30
65. Saxe, A.M. *et al.* (2019) A mathematical theory of semantic development in deep neural networks. *Proc. Natl. Acad. Sci.* 116, 11537–11546
66. Musslick, S. *et al.* (2020) On the rational boundedness of cognitive control: Shared versus separated representations. *PsyArXiv* <https://doi.org/10.31234/osf.io/jkhdh>
67. Yang, G.R. *et al.* (2019) Task representations in neural networks trained to perform many cognitive tasks. *Nat. Neurosci.* 22, 297–306
68. Musslick, S. *et al.* (2017) Multitasking capability versus learning efficiency in neural network architectures. In *Proceedings of the 39th Annual Meeting of the Cognitive Science Society, London, UK*, pp. 829–834
69. Frankland, S.M. and Greene, J.D. (2020) Concepts and compositionality: in search of the brain's language of thought. *Annu. Rev. Psychol.* 71, 273–303

70. Ravi, S. et al. (2020) Navigating the tradeoff between multi-task learning and learning to multitask in deep neural networks. *arXiv* 2007.10527
71. Dreisbach, G. and Haider, H. (2008) That's what task sets are for: shielding against irrelevant information. *Psychol. Res.* 72, 355–361
72. Dreisbach, G. and Haider, H. (2009) How task representations guide attention: further evidence for the shielding function of task sets. *J. Exp. Psychol. Learn. Mem. Cogn.* 35, 477–486
73. Dreisbach, G. (2012) Mechanisms of cognitive control: The functional role of task rules. *Curr. Dir. Psychol. Sci.* 21, 227–231
74. Baxter, J. (1995) Learning internal representations. In *Proceedings of the eighth annual conference on Computational learning theory*, pp. 311–320
75. Caruana, R. (1997) Multitask learning. *Mach. Learn.* 28, 41–75
76. Flesch, T. et al. (2018) Comparing continual task learning in minds and machines. *Proc. Natl. Acad. Sci.* 115, E10313–E10322
77. Lu, X. et al. (2014) Semi-supervised multitask learning for scene recognition. *IEEE Trans. Cybern.* 45, 1967–1976
78. Collobert, R. and Weston, J. (2008) A unified architecture for natural language processing: Deep neural networks with multi-task learning. In *Proceedings of the 25th International Conference on Machine Learning*, pp. 160–167
79. Deng, L. et al. (2013) New types of deep neural network learning for speech recognition and related applications: An overview. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 8599–8603, IEEE
80. Goschke, T. (2000) *Intentional reconfiguration and involuntary persistence in task-set switching*. Control of cognitive processes: Attention and performance XVIII pp. 331–355
81. Verguts, T. (2017) Binding by random bursts: A computational model of cognitive control. *J. Cogn. Neurosci.* 29, 1103–1118
82. Ruge, H. and Wolfensteller, U. (2010) Rapid formation of pragmatic rule representations in the human brain during instruction-based learning. *Cereb. Cortex* 20, 1656–1667
83. Cohen, J.D. et al. (1990) On the control of automatic processes: a parallel distributed processing account of the stroop effect. *Psychol. Rev.* 97, 332–361
84. Botvinick, M.M. et al. (2001) Conflict monitoring and cognitive control. *Psychol. Rev.* 108, 624–652
85. Badre, D. et al. (2021) The dimensionality of neural representations for control. *Curr. Opin. Behav. Sci.* 38, 20–28
86. Sagiv, Y. et al. (2018) Efficiency of learning vs. processing: Towards a normative theory of multitasking. In *Proceedings of the 40th Annual Meeting of the Cognitive Science Society, Madison, WI*, pp. 1004–1009, arXiv:2007.03124
87. Garner, K. and Dux, P.E. (2015) Training conquers multitasking costs by dividing task representations in the frontoparietal-subcortical system. *Proc. Natl. Acad. Sci.* 112, 14372–14377
88. Miller, E.K. and Cohen, J.D. (2001) An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.* 24, 167–202
89. Bouchacourt, F. and Buschman, T.J. (2019) A flexible model of working memory. *Neuron* 103, 147–160
90. Oberauer, K. and Kliegl, R. (2006) A formal model of capacity limits in working memory. *J. Mem. Lang.* 55, 601–626
91. Usher, M. and Cohen, J.D. (1999) Short term memory and selection processes in a frontal-lobe model. In *Connectionist models in cognitive neuroscience*, pp. 78–91, Springer
92. Oberauer, K. et al. (2016) What limits working memory capacity? *Psychol. Bull.* 142, 758–799
93. Treisman, A.M. (1996) The binding problem. *Curr. Opin. Neurobiol.* 6, 171–178
94. Biederman, I. (1987) Recognition-by-components: a theory of human image understanding. *Psychol. Rev.* 94, 115–147
95. Logan, G.D. and Gordon, R.D. (2001) Executive control of visual attention in dual-task situations. *Psychol. Rev.* 108, 393–434
96. Barlow, H.B. (1972) Single units and sensation: a neuron doctrine for perceptual psychology? *Perception* 1, 371–394
97. Liang, J.C. et al. (2020) Experience transforms conjunctive object representations: Neural evidence for unitization after visual expertise. *Cereb. Cortex* 30, 2721–2739
98. McClelland, J.L. (2000) The basis of hyperspecificity in autism: A preliminary suggestion based on properties of neural nets. *J. Autism Dev. Disord.* 30, 497–502
99. Joseph, R.M. et al. (2009) Why is visual search superior in autism spectrum disorder? *Dev. Sci.* 12, 1083–1096
100. Plaisted, K. et al. (1998) Enhanced visual search for a conjunctive target in autism: A research note. *J. Child Psychol. Psychiatry Allied Discip.* 39, 777–783
101. Plaisted, K.C. (2015) Reduced generalization in autism: An alternative to weak central coherence. In *Attention, space, and action: Studies in cognitive neuroscience* (Jacob, N.Y. et al., eds), pp. 149–169, Lawrence Erlbaum Associates
102. Desimone, R. (1991) Face-selective cells in the temporal cortex of monkeys. *J. Cogn. Neurosci.* 3, 1–8
103. Eickenberg, M. et al. (2017) Seeing it all: Convolutional network layers map the function of the human visual system. *NeuroImage* 152, 184–194
104. Agrawal, A. et al. (2020) A compositional neural code in high-level visual cortex can explain jumbled word reading. *eLife* 9, e54846
105. LeCun, Y. et al. (2015) Deep learning. *Nature* 521, 436–444
106. Schmidhuber, J. (2015) Deep learning in neural networks: An overview. *Neural Netw.* 61, 85–117
107. O'Reilly, R.C. (1998) Six principles for biologically based computational models of cortical cognition. *Trends Cogn. Sci.* 2, 455–462
108. Padmala, S. and Pessoa, L. (2011) Reward reduces conflict by enhancing attentional control and biasing visual cortical processing. *J. Cogn. Neurosci.* 23, 3419–3432
109. Fischer, R. and Hommel, B. (2012) Deep thinking increases task-set shielding and reduces shifting flexibility in dual-task performance. *Cognition* 123, 303–307
110. Dreisbach, G. et al. (2005) Dopamine and cognitive control: the influence of spontaneous eyeblink rate and dopamine gene polymorphisms on perseveration and distractibility. *Behav. Neurosci.* 119, 483–490
111. Tharp, I.J. and Pickering, A.D. (2011) Individual differences in cognitive-flexibility: The influence of spontaneous eyeblink rate, trait psychoticism and working memory on attentional set-shifting. *Brain Cogn.* 75, 119–125
112. Crofts, H. et al. (2001) Differential effects of 6-ohda lesions of the frontal cortex and caudate nucleus on the ability to acquire an attentional set. *Cereb. Cortex* 11, 1015–1026
113. Cools, R. et al. (2009) Striatal dopamine predicts outcome-specific reversal learning and its sensitivity to dopaminergic drug administration. *J. Neurosci.* 29, 1538–1543
114. Dreisbach, G. and Goschke, T. (2004) How positive affect modulates cognitive control: reduced perseveration at the cost of increased distractibility. *J. Exp. Psychol. Learn. Mem. Cogn.* 30, 343–353
115. van Wouwe, N.C. et al. (2011) Positive affect modulates flexibility and evaluative control. *J. Cogn. Neurosci.* 23, 524–539
116. Aston-Jones, G. and Cohen, J.D. (2005) An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annu. Rev. Neurosci.* 28, 403–450
117. Cools, R. (2016) The costs and benefits of brain dopamine for cognitive control. *Wiley Interdiscip. Rev. Cogn. Sci.* 7, 317–329
118. Cools, R. (2019) Chemistry of the adaptive mind: Lessons from dopamine. *Neuron* 104, 113–131
119. Ueltzhöffer, K. et al. (2015) Stochastic dynamics underlying cognitive stability and flexibility. *PLoS Comput. Biol.* 11, e1004331
120. Musslick, S. et al. (2018) Constraints associated with cognitive control and the stability-flexibility dilemma. In *Proceedings of the 40th Annual Meeting of the Cognitive Science Society, Madison, WI*, pp. 806–811
121. Musslick, S. et al. (2019) Stability-flexibility dilemma in cognitive control: A dynamical system perspective. In *Proceedings of the 41st Annual Meeting of the Cognitive Science Society, Montreal, CA*, pp. 2420–2426
122. Mayr, U. (2006) What matters in the cued task-switching paradigm: Tasks or cues? *Psychon. Bull. Rev.* 13, 794–799
123. Monsell, S. and Mizon, G.A. (2006) Can the task-cuing paradigm measure an endogenous task-set reconfiguration process? *J. Exp. Psychol. Hum. Percept. Perform.* 32, 493–516
124. Siqi-Liu, A. and Egner, T. (2020) Contextual adaptation of cognitive flexibility is driven by task- and item-level learning. *Cogn. Affect. Behav. Neurosci.* 20, 757–782

125. Liu, C. and Yeung, N. (2020) Dissociating expectancy-based and experience-based control in task switching. *J. Exp. Psychol. Hum. Percept. Perform.* 46, 131–154
126. Dreisbach, G. and Fröber, K. (2019) On how to be flexible (or not): Modulation of the stability-flexibility balance. *Curr. Dir. Psychol. Sci.* 28, 3–9
127. Panichello, M.F. et al. (2019) Error-correcting dynamics in visual working memory. *Nat. Commun.* 10, 3366
128. Shaffer, L. (1975) *Multiple attention in continuous verbal tasks. Attention and performance V* pp. 157–167
129. Stephan, D.N. and Koch, I. (2010) Central cross-talk in task switching: Evidence from manipulating input-output modality compatibility. *J. Exp. Psychol. Learn. Mem. Cogn.* 36, 1075–1081
130. Koch, I. et al. (2018) Cognitive structure, flexibility, and plasticity in human multitasking—an integrative review of dual-task and task-switching research. *Psychol. Bull.* 144, 557–583
131. Nijboer, M. et al. (2014) Single-task fMRI overlap predicts concurrent multitasking interference. *NeuroImage* 100, 60–74
132. Navon, D. and Miller, J. (1987) Role of outcome conflict in dual-task interference. *J. Exp. Psychol. Hum. Percept. Perform.* 13, 435–448
133. Webb, T. et al. (2020) Learning representations that support extrapolation. In *International Conference on Machine Learning*, pp. 10136–10146, PMLR
134. Kang, Y.H. et al. (2021) Multiple decisions about one object involve parallel sensory acquisition but time-multiplexed evidence incorporation. *eLife* 10, e63721
135. Panichello, M.F. and Buschman, T.J. (2021) Shared mechanisms underlie the control of working memory and attention. *Nature* 592, 601–605
136. Dux, P.E. et al. (2009) Training improves multitasking performance by increasing the speed of information processing in human prefrontal cortex. *Neuron* 63, 127–138
137. Del Giudice, M. and Crespi, B.J. (2018) Basic functional trade-offs in cognition: An integrative framework. *Cognition* 179, 56–70
138. Masis, J.A. et al. (2020) Rats strategically manage learning during perceptual decision making. *bioRxiv* Published online September 2, 2020. <https://doi.org/10.1101/2020.09.01.259911>
139. Cohen, J.D. et al. (2007) Should I stay or should I go? how the human brain manages the trade-off between exploitation and exploration. *Philos. Trans. R. Soc. B Biol. Sci.* 362, 933–942
140. Geana, A. et al. (2016) Boredom, information-seeking and exploration. In *Proceedings of the 38th Annual Meeting of the Cognitive Science Society, Austin, TX*, pp. 1751–1756
141. Schmidhuber, J. (1991) A possibility for implementing curiosity and boredom in model-building neural controllers. In *Proc. of the international conference on simulation of adaptive behavior: From animals to animats*, pp. 222–227
142. Agrawal, M. et al. (2021) The temporal dynamics of opportunity costs: A normative account of cognitive fatigue and boredom. *Psychol. Rev.* <https://doi.org/10.1037/rev0000309>
143. Welford, A.T. (1952) The psychological refractory period and the timing of high-speed performance—a review and a theory. *Br. J. Psychol.* 43, 2–19
144. Moray, N. (1967) Where is capacity limited? a survey and a model. *Acta Psychol.* 27, 84–92
145. Smith, M.C. (1967) Theories of the psychological refractory period. *Psychol. Bull.* 67, 202–213
146. Posner, M.I. and Boies, S.J. (1971) Components of attention. *Psychol. Rev.* 78, 391–408
147. Kinsbourne, M. and Hicks, R.E. (1978) *Functional cerebral space: A model for overflow, transfer and interference effects in human performance. Attention and Performance VII* pp. 345–362
148. Logan, G. (1978) Attention in character-classification tasks: Evidence for the automaticity of component stages. *J. Exp. Psychol. Gen.* 107, 32–63
149. Salvucci, D.D. et al. (2004) *Toward an ACT-R General Executive for Human Multitasking*, ICCM, pp. 267–272
150. O'Reilly, R.C. and Frank, M.J. (2006) Making working memory work: a computational model of learning in the prefrontal cortex and basal ganglia. *Neural Comput.* 18, 283–328
151. Eisenreich, B.R. et al. (2017) Control without controllers: toward a distributed neuroscience of executive control. *J. Cogn. Neurosci.* 29, 1684–1698
152. Fischer, R. and Plessow, F. (2015) Efficient multitasking: parallel versus serial processing of multiple tasks. *Front. Psychol.* 6, 1366
153. Telford, C.W. (1931) The refractory phase of voluntary and associative responses. *J. Exp. Psychol.* 14, 1–36
154. Ruthruff, E. et al. (2006) How does practice reduce dual-task interference: Integration, automatization, or just stage-shortening? *Psychol. Res.* 70, 125–142
155. Liepelt, R. et al. (2011) Improved intertask coordination after extensive dual-task practice. *Q. J. Exp. Psychol.* 64, 1251–1272
156. Strobach, T. (2020) The dual-task practice advantage: Empirical evidence and cognitive mechanisms. *Psychon. Bull. Rev.* 27, 3–14
157. Schumacher, E.H. et al. (2001) Virtually perfect time sharing in dual-task performance: Uncorking the central cognitive bottleneck. *Psychol. Sci.* 12, 101–108
158. Oberauer, K. and Kliegl, R. (2004) Simultaneous cognitive operations in working memory after dual-task practice. *J. Exp. Psychol. Hum. Percept. Perform.* 30, 689–707
159. Schubert, T. et al. (2017) Transferability of dual-task coordination skills after practice with changing component tasks. *Front. Psychol.* 8, 956
160. Kieras, D.E. et al. (2000) *Modern computational perspectives on executive mental processes and cognitive control: Where to from here. Control of cognitive processes: Attention and performance XVIII*, pp. 681–712
161. Strobach, T. and Torsten, S. (2017) Mechanisms of practice-related reductions of dual-task interference with simple tasks: data and theory. *Adv. Cogn. Psychol.* 13, 28–41
162. Taatgen, N.A. and Lee, F.J. (2003) Production compilation: A simple mechanism to model complex skill acquisition. *Hum. Factors* 45, 61–76
163. McClelland, J.L. et al. (1995) Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol. Rev.* 102, 419–457
164. O'Reilly, R.C. et al. (2016) The leabra cognitive architecture: How to play 20 principles with nature. *Oxf. Handb. Cognit. Sci.* 91, 91–116
165. Usher, M. and McClelland, J.L. (2001) The time course of perceptual choice: the leaky, competing accumulator model. *Psychol. Rev.* 108, 550–592
166. Kalanithroff, E. et al. (2018) Task conflict and proactive control: A computational theory of the stroop task. *Psychol. Rev.* 125, 59–82
167. Braver, T.S. et al. (1999) Cognition and control in schizophrenia: a computational model of dopamine and prefrontal function. *Biol. Psychiatry* 46, 312–328
168. Herd, S.A. et al. (2014) A neural network model of individual differences in task switching abilities. *Neuropsychologia* 62, 375–389
169. Stroop, J.R. (1935) Studies of interference in serial verbal reactions. *J. Exp. Psychol.* 18, 643–662
170. Papadopettraki, D. et al. (2019) Quantifying the cost of cognitive stability and flexibility. *bioRxiv* <https://doi.org/10.1101/743120>