Advances in Cognitive Systems 6 (2018) 85–100 Submitted 6/2018; published 9/2018

Towards an Attention-Driven Model of Task Switching

Will Bridewell WILL.BRIDEWELL@NRL.NAVY.MIL
Christina Wasylyshyn CHRISTINA.WASYLYSHYN@NRL.NAVY.MIL
Paul F. Bello PAUL.BELLO@NRL.NAVY.MIL
U.S. Naval Research Laboratory, 4555 Overlook Ave SW, Washington, DC 20375 USA

Abstract

We describe a computational model of task switching developed within the ARCADIA framework. This model emphasizes the attentional component of task sets and the locality of processing for task selection, stimulus–response rules, and motion planning. We contrast the structure and processing of the model to one developed in ACT-R and illustrate its dynamics on one trial from a task switching experiment. We conclude by discussing the broader implications of the model for cognitive modeling and cognitive systems research.

1. Introduction

We engage in tasks everyday like answering an email, making dinner, or participating in a psychology experiment. Sometimes the tasks are well practiced, other times they are newly conceived. Within cognitive psychology, the prevailing thought on how we arrange our minds to pursue these tasks is that we rely on structured representations called *task sets*. These aggregate representations shape our perceptions and organize our behaviors so that we can engage in temporally extended, goal-directed activities. A considerable amount of literature in psychology and neuroscience has attempted to uncover the structure and dynamics of task sets, but a clear image is still out of sight.

If we are to build artifacts that engage in diverse, goal-directed behavior, then knowing how task-related information is structured in the human mind could provide much needed insight into how we might design those artifacts. Currently, there is general agreement that task sets include a *stimulus set* that identifies the objects or properties relevant to the task and a *response set* that includes all the task-relevant actions. These sets are organized under a task identifier that is sometimes called a goal or simply a task. From this point of commonality, models of task sets diverge and are shaped by implementation-specific concerns.

Grange and Houghton (2014) provide an excellent overview of existing models, grouping them as mathematical, computational, or neural network approaches. All the models that they review were developed to account for human behavioral data in the paradigm of task switching. However, because switching among tasks requires people to replace one task set with another, each of these approaches takes a stand on how these are organized. Of the three groups, we are concerned with computational models because they emphasize the reproduction of human behavior without adopting the strictures of neural architecture. Moreover, our efforts to represent task sets within the

ARCADIA framework (Bridewell & Bello, 2016b) fall naturally into this category. Specifically, we concentrate our attention on the model implemented by Altmann and Gray (2008) in the ACT-R cognitive architecture, using it as a point of comparison with the model that we have developed.

Like Altmann and Gray, along with others who investigate task switching, we are interested in questions of cognitive control. How do we engage in flexible, goal-directed behavior in a dynamic world that contains an essentially infinite variety of obstacles, distractions, and other interferences? Models in ACT-R address this question starting from a theory of memory and, as a result, the main emphasis is on what the individual elements in memory are and how their activation levels are modified throughout task execution. In contrast, we start from a theory of attention. This perspective lets us reflect on how our task guides our attention to relevant objects in the environment and to important properties of those objects that we have perceived or inferred. In this paper, we emphasize the key point that conflict resolution in the ACT-R model takes place at a global, task level, whereas in ARCADIA it can also occur locally, in parallel, and in many cases without explicit attention or cognitive control.

2. The Psychology of Task Switching

To discuss task sets, we need to introduce the task-switching paradigm. This experimental paradigm is not the only one where task sets are discussed, but it is arguably the most well studied, and it is the target of both our research and the work reported by Altmann and Gray. Task switching is primarily used to investigate cognitive control processes, particularly those associated with efficiently and rapidly switching between different tasks. Accordingly, models of task switching behavior often assume the role of executive control functioning, but some do not (e.g., see Logan & Bundesen, 2003). All models, however, provide some description of how people can adapt to a dynamic environment either through changes in their actions and/or deployment of attention.

Task switching is studied with experiments that can vary on several dimensions. In terms of control, the stimuli can differ in whether and how an external cue is presented, denoting which task to perform. In terms of complexity, they can differ in the number and types of tasks and the modalities in which tasks are presented. In terms of time, they can differ in how long participants have both to prepare for an upcoming trial and to recover between response completion and the presentation of the next cue or task (see Monsell, 2003, for a review). In conventional, externally cued, task-switching paradigms, participants perform two reaction-time tasks within the same block of trials. Often, *switch costs* are the dependent measurement and reflect latency differences between trials on which participants repeat the same task and trials on which they switch between tasks. The tasks usually involve the basic identification or categorization of digits, objects, or words within a single sensory modality. In this paper, a digits task is modeled, such that participants make judgments on either a single digit's parity (even or odd) or its magnitude (greater or less than five), and an explicit cue signals which task to perform on each trial.

Figure 1 illustrates a typical trial of the experiment modeled in Section 4.5. First, a cue is presented for a certain amount of time to denote what the upcoming task will be. During this interval, participants determine which attributes of the target stimulus to distinguish, and this often involves switching between category rules. In this example, a participant would see the cue "M" for

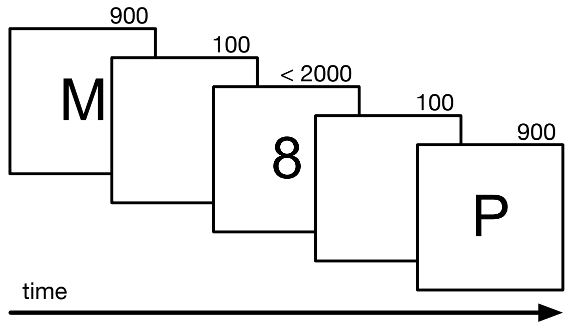

Figure 1. Time course of experimental display. Cues appear for 900 ms, delays between cue and stimulus and between trials are 100 ms, and target stimuli appear until a button press, up to 2000 ms.

"magnitude" and know that a greater or less than five judgment will be made on the upcoming trial. In this case, odd and low responses share one button and even and high responses share another button. This means that the number eight maps to the same response across the tasks whereas nine maps to different responses. After a brief presentation of a blank screen, a target is presented until a response is made. In this experiment, the stimulus can be any integer from one to four and from six to nine. A decision for a particular task is expressed in an action and often involves switching between response rules: for example, seeing the number "6" and deciding that a response press denoting "even" should be made. After the subject responds or a time limit is reached, another blank screen appears, and the next trial begins.

How these two components, switching category rules and switching response rules, are accomplished in the course of a given trial defines a task set, or the effective intention to perform a task. However, as we mentioned in the introduction, the representation of a task set in the literature has been hazy. In response, we suggest that viewing the task set through the lens of attention can provide some clarity in addition to any provided by the lens of memory.

3. Task Switching in ACT-R

In this section, we describe key elements of Altmann and Gray's (2008) ACT-R model of task switching. Our emphasis is on how they represent the components of a task set, the dynamics of task switching and task execution, and to a lesser extent, the productions that encode control knowledge. To be clear, we do not address the activation and retrieval processes in ACT-R or the specifications and values of the 14 model parameters. Additionally, Altmann and Gray evaluated their model with *repeat runs* when task switching does not occur, a measure that we also do not consider. For those details, we direct readers to their original paper.

For Altmann and Gray, a task set is composed entirely of elements in semantic memory. These include the tasks, stimulus categories, and responses. The parity task is symbolized by "evenodd" and the magnitude task, which involves determining whether a single digit falls above or below five on a number line, is symbolized by "highlow." The stimulus categories are "even," "odd," "high," and "low." Finally, the responses are "left" and "right" corresponding to the positions of buttons to be pressed. For the purposes of spreading activation, the tasks are connected to the stimulus categories, which are connected to the responses. All of these elements may be activated by percepts, such that the percepts associated with the cues activate the tasks and those from the task stimuli (i.e., the numerals 1–4 and 6–9) activate the stimulus categories and possibly the responses. In addition, episodic memory stores *task codes* that are connected to the tasks in semantic memory and are used to track which one is currently being executed.

Execution in the model is controlled by productions and follows a routine pattern. The productions direct ACT-R to create a task code based on a cue, to increase its activation, and to manage the retrieval of elements from semantic memory during task execution. This leads to an execution trace of the following sort. The cue appears, a task code is created, and its activation level is increased. After the stimulus appears, the model continues to increase task code activation until it has reached a threshold value. Then the task code is retrieved ("evenodd"), the stimulus category is identified ("6"), the response is selected ("left"), and the response is executed. Each cycle of execution is specified to take 50 ms with the exception of response execution, which is specified to take 345 ms. Assuming the task code is activated strongly enough to enable immediate retrieval upon stimulus presentation, task execution takes 200 ms for stimulus processing for a minimum of a 545 ms response time.

Although, at a high level, Altmann and Gray's model agrees with the generally accepted view that a task set comprises links between stimulus categories and responses, the specific representation is at odds with other findings. For instance, this model relies on a limited notion of episodic memory that keeps track of only the task codes ("evenodd" or "highlow") and relies on activation levels when recalling the current task. This characterization fails to cohere with results from Altmann's (2013) work where he concedes that based on latency data, episodic memory is not involved in explicitly cued task-switching. Although Altmann does argue that error rates reflect an episodic memory component, the simpler explanation is that participants failed to switch tasks after cue presentation. As a result, it is likely that there is some encoding of the current task, but asserting a role for episodic memory remains unwarranted.

In addition, Altmann and Gray's reliance on episodic memory raises the issue of local versus global conflict resolution. According to their account, errors in task execution result from retrieving the wrong task code from episodic memory. However, Brown, Reynolds, and Braver (2007) developed a neural network model and argued for at least two conflict resolution mechanisms in task switching. Their claim aligns with Regev and Meiran's (2016) empirical work showing that conflict resolution occurs not only at the task-set level but also when processing stimulus–response rules. Specifically, their work found separate indicators of suppression associated with selection of responses and with selection of tasks. At the least, these findings suggest localized conflict resolution within the task set. Compounding these results, we see that alternative motions are planned in parallel and introduce a further point of conflict resolution (Gallivan et al., 2015). Such localized activity is not accounted for by Altmann and Gray, and it is unclear with their current representation and modeling commitments how they could shift from a global view of conflict resolution to a localized view. As evidence accumulates, it appears that we should move to a modeling paradigm that supports local, and even parallel, conflict resolution for task switching.

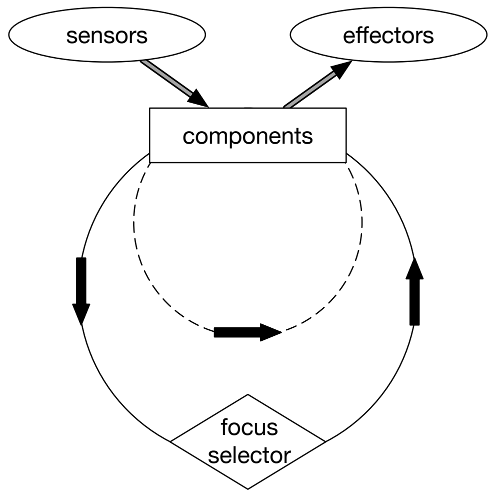

Figure 2. Pictorial representation of ARCADIA's 25 ms cycle. Each component processes a subset of information read from sensors during the current cycle and produced by all the components during the previous cycle. The focus selector uses an attentional strategy to pick the focus of attention, an element of information that is given priority by the components that can operate on its encoded data. Effectors enable ARCADIA to interact with its external environment.

4. An Attention-Driven Model of Task Switching

In this section, we introduce a model called Attentional Strategies and Task Reconfiguration in Action (ASTRA), which is developed using the ARCADIA framework. After providing a brief overview of ARCADIA, we discuss the representation of task sets and describe components that load, store, execute, and swap the task sets. We then provide a specific task-switching scenario and discuss components that process task-specific information. Finally, we walk through an example of cue and stimulus processing in ASTRA.

4.1 The ARCADIA Framework

ARCADIA is at once both a theory of human cognitive architecture and a computational framework for building attentionally guided, cognitive systems (Bridewell & Bello, 2016b). Whereas the next sections discuss portions of ARCADIA in terms of their relationship to psychological and neuroscientific theorizing, this abbreviated introduction treats it as a general framework, which is depicted in Figure 2. At its core, ARCADIA is a collection of *components*, each of which computes a function that maps input to the component into output. Information is exchanged among components using an *interlingua* that bundles information represented in heterogeneous data formats into discrete elements. For example, the element in Table 1 encodes a request to push a button. The components operate in parallel during a simulated 25 ms cycle. On each cycle, the elements produced during the previous one are visible to every component. In addition, components have access to an external environment through sensors, which are currently limited to carrying visual data, and can issue commands to effectors, which can alter that environment.

Table 1. An interlingua element for the action request to push the button assigned to odd/low numbers. This request is for the number parity task and was generated by the system component that processes stimulus–response links.

```
{:id g0112
 :name "push-button"
 :arguments {:button-id :b-odd-low
             :task :parity
             :effector :pointer}
 :world nil
 :source "sr-link-processor"
 :type "action"}
```

To carry out integrated processing in ARCADIA, a *focus of attention* directs components to apply their resources to a specific interlingua element. On each 25 ms cycle, one element is chosen as the focus, after which its status is announced to every component. Components that have routines for processing information in the focus of attention will do so, while others may continue to operate independently of what is attended. The choice of the focus of attention is determined by an *attentional strategy*: a task-specific priority list. A default strategy, which we have used in several cognitive models, is presented in Table 2, and ARCADIA's main cycle is shown in Table 3. The effects of attending to a particular element depend on the components within a model, and a typical strategy guides the components toward the fulfillment of a task without eliminating reactivity to the environment. For example, in the model described in the next section, attending to a visual representation of an object will enable various components to report features of that object. On the subsequent cycle, those features will get bound together into the object representation unless a component produces a request to press a button or to subvocalize a task name.

4.2 Task Sets in ASTRA

As with any model in ARCADIA, information processing in ASTRA is distributed and attention provides a control signal that enables adaptive, temporally extended, task-oriented behavior. The representation of tasks in this model draws from the psychological and neuroscientific literature on task switching and is intended to provide a general framework for specifying tasks within all ARCADIA models. Similar to the Altmann and Gray view, tasks in ASTRA comprise a set of stimulus–response (SR) rules and a name that serves as a semantic tag for deliberate selection (e.g., "This is the *odd–even* task."). Unlike Altmann and Gray's model, in which control knowledge for task execution exists as a set of productions unconnected to the task representation, the task sets in ASTRA include an attentional strategy that prioritizes responses related to that task. Specifically, the strategy enables the resolution of conflicts when SR rules for different tasks are activated during the same trial. These parts of a task are bound together in a sense but, when a task is active, individual parts are stored and processed across different components. Therefore, task representation is a cluster of other representations, each in its own format and mechanisms for information processing.

Table 2. The default attentional strategy for ARCADIA, which searches the output of components for requests (1) to adopt an attentional strategy, (2) to carry out an action, (3) to attend to an internal object representation, or (4) to shift covert attention to a potential object in the world.

```
(defn select-focus-default [expected]
  ;; input: the current set of interlingua elements produced by the components.
  ;; output: the interlingua element that will be the next focus of attention.
  (or (rand-element expected :name "adopt-attentional-strategy")
      (rand-element expected :type "action")
      (rand-element expected :name "object" :type "instance" :world nil)
      (rand-element expected :name "fixation")))
```

By taking a closer look at SR rules and attentional strategies, we can point out key commitments behind this view of a task set. In ASTRA, the SR rules connect a stimulus through its properties to a request for action as shown in Table 4. The *stimulus* part is a predicate that returns true when one or more elements produced by components match its definition. The *response* part is an action-request element. This request is only executed by the system if it receives attention. Therefore, the task's attentional strategy plays a key role in determining whether a particular SR rule leads to an action, not only by focusing attention on the cue or stimulus, but also by prioritizing action requests associated with the current task.

The core of ASTRA is a set of components developed to support task switching, but we did make a key change in how ARCADIA selects the focus of attention. In previously reported models, a static, monolithic, attentional strategy was sufficient, because each model accounted for a single psychological task. Enabling flexible, dynamic shifts among tasks requires a focus-selection mechanism that supports swapping strategies. To this end, we extended the existing mechanism so that it can be given an alterable, task-related strategy whose preferences are prioritized over the default strategy. In this model, the task-related strategies prioritize actions produced by the task's SR rules followed by requests to subvocalize words. The default strategy operates whenever no elements match the task strategy. In that case, as shown in Table 2, the priorities are (1) requests to change attentional strategies, (2) action requests, (3) newly created object representations, and (4) requests to shift covert visual attention. If no preferred elements exist, one is chosen uniformly at random.

4.3 The Dynamics of Task Switching

We intend for the task-processing components introduced in ASTRA to serve as a general-purpose group for use in any ARCADIA model that switches among multiple tasks or carries out tasks with complex, hierarchical structure. As designed, this group includes five components that emphasize task sets as distributed, structured representations and processes, the dynamics of which are illustrated in Figure 3. These components include *task-memory*, *task-wm*, and *sr-link-processor*, which act in part as memory systems, and *task-refresher* and *attention-manager*, which request updates to memory contents. Before describing the specifics of the task shown in Figure 3, we discuss the relationships among the task-switching components used in the example.

Table 3. Clojure code for executing a single cycle in ARCADIA.

```
(defn ARCADIA-Cycle
  ;; input: the current state of the simulation environment
  ;; output: the next state of the simulation environment
  [env]

  ;; broadcast the focus of attention to all the components so that they can
  ;; execute their functions over the current interlingua elements.
  (broadcast-focus (access/get-focus) (access/get-content))

  ;; retrieve the resulting interlingua elements from all the components.
  (doseq [c (get-components)] (component/deliver-result c))

  ;; select the next focus of attention using the current attentional strategy.
  (focus-selector/select-focus!)

  ;; step through the environment, executing any actions that are slated to
  ;; be carried out on this cycle.
  (simulator/step env (environment-actions (access/get-content)))
  env)
```

At the moment, ARCADIA lacks components that turn instructions into task representations, so we assume that each task is initially defined and stored in a long-term *task-memory* component. On most cycles, this component is inactive, but when the focus of attention is on a request to subvocalize a word that corresponds to the name of a known task, *task-memory* issues a command to load the associated task into working memory. If this memorization request receives attention, then *task-wm* replaces its current task with the newly requested one. This activity is similar to task-code creation and retrieval in Altmann and Gray's model, but instead of invoking episodic memory, we have a specific limited-capacity, working memory location for tracking the current task. This implementation is in line with evidence that, typically, the monitoring and motivational factors for a single, active goal are represented in the medial frontal cortex (Charron & Koechlin, 2010). Like other short-term memory components in our models, *task-wm* reports its current element (a task structure) on each cycle. When ARCADIA starts, *task-wm* is empty and ASTRA relies on the default attentional strategy. To ensure that the modeled task is active, we use a component called *task-initializer* (not illustrated) to load the top-level task into working memory during the first few cycles of operation.

When a task is memorized, not only is it loaded into task-specific working memory but also its parts are activated in three components of the model. The first is *task-refresher*, which upon seeing a request to load a task issues a request to update the current attentional strategy. If this request receives attention, then *attention-manager* ensures that ARCADIA's focus selection mechanism updates its priority structure to use the new strategy. While *task-refresher* brings the system's attentional priorities in line with the specified task, *sr-link-processor* stores the associated SR rules in its list of active rules. This implementation is in line with evidence that SR rules are stored in the premotor dorsal cortex and pre-premotor dorsal cortex (Badre & D'Esposito, 2009). In ASTRA,

Table 4. An example of how to construct a stimulus–response rule in the ASTRA model, which consist of (i) a predicate that matches interlingua elements with a name of "object-property" and an argument list that includes the values ":parity" for the attribute :property and ":odd" for :value, and (ii) a function that produces an action request interlingua element for pushing the button assigned to "odd" for the "parity" task.

```
;; descriptor constructs a predicate that can be used to filter interlingua
;; elements for only those that match all the specified values.
(sr-link [(descriptor :name "object-property", :property :parity, :value :odd)]
         ;; push-button constructs a function that produces an
         ;; interlingua element for an action request.
         (push-button :b-odd-low :parity))
```

rules, once loaded, are never inactivated. Although this limitation should be addressed in later versions, we note that in a complex task environment such as one that requires task switching, multiple sets of potentially contradictory SR rules can be active simultaneously (Vandierendonck, 2016). As a result, the model must be robust to responses being generated in parallel.

To present trials of this sort to ARCADIA, we give it images at a simulated rate of 40 Hz. Interpreting this imagery relies on a group of components for visual processing described elsewhere (Bridewell & Bello, 2016a). In short, the current image is read by a sensor component and operated on by other low-level visual components. After a few cycles, an object representation is stored in visual short-term memory. The important points about this approach are that ARCADIA operates on image data and that creating stored visual representations of objects is attention-dependent and takes multiple cycles. As a result, these vision-related components account for part of the time required for cue and stimulus interpretation in a manner that maps to what occurs in the course of human visual processing.

4.4 Task-Specific Components

Interpreting the cue and target characters is handled by four task-specific components. First, when an object is the focus of attention, *character-reporter* identifies any single letters or numbers that appear in the associated image segment. This character is associated with the object. Second, when that character is a numeral, *number-identifier* reports the number that it represents, which is also associated with the object. The third and fourth components, *number-parity* and *number-magnitude*, report whether that number is odd or even and greater or less than a given threshold. Although the components are task-specific, we note that they each provide capabilities in a manner general enough for inclusion in other models. However, they differ from the visual processing and the new task processing components, which we consider practical for use in all ARCADIA models.

In addition to these components, there are specific SR rules for cue interpretation and both the magnitude and parity tasks. Cue interpretation uses two rules, one that connects identifying the character 'P' to a request to subvocalize "parity" and another that connects 'M' to "magnitude". The task rules connect relevant properties to requests to press a button. For instance, if the parity property "odd" is detected, then issue an action request associated with the parity task to press the "odd-low" button. A sketch of these rules appears in the memory row of Figure 3.

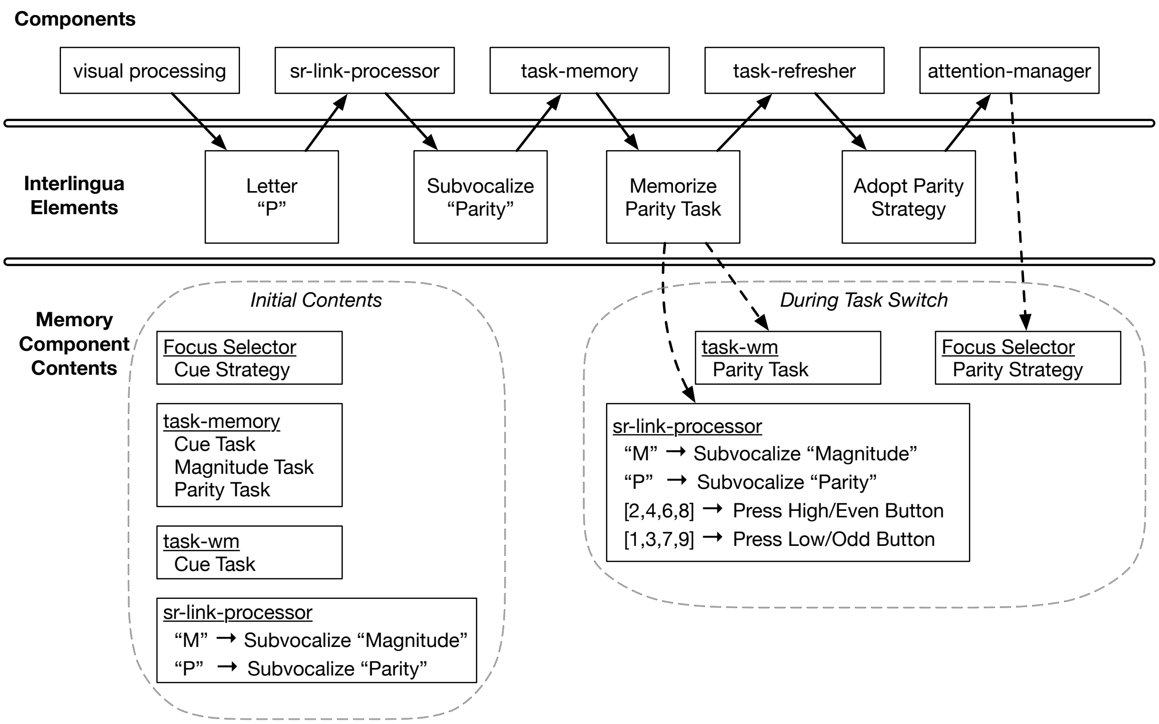

Figure 3. Example time course of a cue-based task switch, starting at the point where the cue-interpretation task is currently active. The top row displays components that contribute to the task switch. The middle row shows a gloss of the interlingua elements produced during the switch. The bottom row shows the contents of the memory components and focus selector (a part of the ARCADIA architecture) before the task switch and as the task switch occurs. Note that the representation of the active task is distributed across multiple memory-related parts of the model.

Finally, cue interpretation and both task sets have associated attentional strategies that are prioritized over ARCADIA's default strategy when they are loaded. As shown in Table 5, the strategy for each task prioritizes (1) an action produced by an SR rule for that task and (2) a specific action request to subvocalize a word. The assumption behind this implementation is that the provenance of an action request is encoded with it. For example, the parity task, the odd or even property, the SR rule, and the response are connected through the task set. Consequently, both the activation of an SR rule ("push 'b' when the number is even") and the corresponding response are traceable as following from the parity task. The consensus holds that task-set representations are distributed and that parts can be shared, but nevertheless parts of task sets can be exchanged or reprioritized while leaving other parts unaltered (Rangelov et al., 2013).

4.5 An Example of Task Switching in ASTRA

To illustrate how these parts of the model come together and point toward the implications of the perspective that they reflect, we walk through one cue and target presentation as depicted in Figure 1. Our goal with this model is not to match any specific human data, but instead to characterize the

Table 5. The attentional strategy for the parity task, which searches the output of components for requests (1) to carry out an action instigated by a parity-related, stimulus–response rule and (2) to subvocalize some word or phrase. If no elements match these preferences during a cycle, then ARCADIA relies on the default strategy.

```
(defn select-focus-parity [expected]
  ;; input: the current set of interlingua elements produced by the components.
  ;; output: the interlingua element that will be the next focus of attention.
  (or (rand-element expected :type “action” :task :parity)
      (rand-element expected :type “action” :name “subvocalize”)))
```

stages of information processing involved, their ordering, and to what extent they may occur in parallel. As a result, this model may be seen to place a lower bound on the time required to interpret the stimuli in this particular task-switching paradigm. Additionally, the example helps address the question of what a task set is by showing how its different aspects are dynamically invoked.

Beginning with visual processing, Figure 3 shows the time course of cue processing with respect to the model. Within ARCADIA, each of these stages is assumed to take 25 ms with the exception of visual processing which is simplified in this picture and takes 100 ms, or four cycles. Visual processing ends with *character-reporter* identifying the associated letter (“M” or “P”). In the next cycle, *sr-link-processor* responds to the character stimulus and issues a request to subvocalize the task name associated with the letter (“magnitude” or “parity”). That request receives attention and the subvocalization process begins simultaneously with the request to memorize or refresh the task-set associated with the name. In response to this request, *task-refresher* prepares a request to adopt the attentional strategy to fit the associated task. At the same time, *task-wm* updates its stored task set and *sr-link-processor* ensures that the associated rules are active. Finally, *attention-manager* switches its attentional strategy so that the model is guided to task-relevant representations. As modeled, this entire process takes eight cycles or 200 ms of simulated time.

Figure 4 provides a similar view for processing the target stimulus. After visual processing, three stages of property identification are interleaved with the binding of that information to the object representation. These properties include the numeral (*character-reporter*), the number represented by the numeral (*number-reporter*), and the task-relevant characteristics of the number (*number-magnitude* and *number-parity*). The associated SR rules connect the parity and magnitude properties to task-specific action requests that are then sent to the motion planner. The planner holds the action requests in abeyance for 150 ms before executing them in the environment. This delay is based on the values reported by Haith, Pakpoor, and Krakauer (2016), and is an underestimate of the actual time that people take since there appears to be an extra 80 ms delay between motion preparation and action initiation. As a result, the entire process (as modeled) takes a simulated 400 ms: 225 ms of initial processing, 150 ms of motion planning, and 25 ms for movement initiation.

Finally, we provide a closer look at localized conflict resolution in ASTRA. After the magnitude and parity tasks have been loaded into task working memory at least once, their rules will be included in the *sr-link-processor*. As a result, both potential action requests associated with the tasks will be produced and both may receive attention. In cases where both requests appear during the

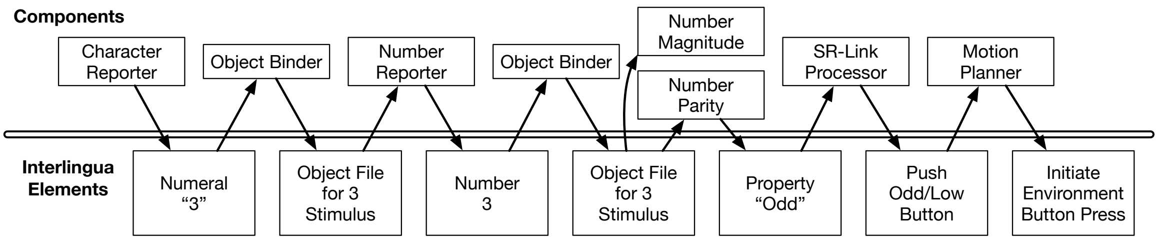

Figure 4. Example time course of target processing, starting at the point just after visual processing. The top row displays components that contribute to target interpretation. The bottom row shows a gloss of the interlingua elements produced during task execution.

same ARCADIA cycle, then the attentional strategy can make a local determination to prioritize the action associated with the current task. In other cases, for instance when one target property takes fewer cycles to process than others, the motion planner has a limited ability to resolve any conflicts. Specifically, when there are no other planned actions and an action request is the focus of attention, the motion planner stores that request and starts a timer. As that timer counts down, any other action requests that arrive will be held alongside the initial one. When the timer reaches zero, the request associated with the active task is enacted regardless of when, during that delay, it was made. If no task-specific action is requested, then the first to arrive is executed. Even though there is more research to be carried out on this topic, these two situations illustrate the localized approach and how different mechanisms for conflict resolution are realized within ARCADIA models.

5. Discussion

We began this paper by asking the question, "How do humans engage in flexible, goal-directed behavior in a dynamic world that contains an essentially infinite variety of obstacles, distractions, and other interferences?" The answer that we sought to give was one invested in the perspective of computational cognitive science, a perspective that addresses the representations and processes that support mental activity. To this end, we reviewed a model of task switching that was centered around a theory of memory and used episodic memory to drive its behavior. Mounting psychological, neuroscientific, and computational evidence for the presence of distributed, localized conflict resolution operating in human cognition led us to develop a novel task representation within the ARCADIA framework. In this section, we state the relevance of this approach to cognitive models and, more broadly, to cognitive systems research. Afterwards, we summarize our answer to our driving question.

5.1 Relevance to Cognitive Modeling

The value of the present research to the field of cognitive modeling can be expressed by the ways in which it differs from the ACT-R model discussed in Section 3. First, ASTRA incorporates control knowledge into the task set in the form of an attentional strategy. To our knowledge, this is a unique perspective among models of task switching. This choice is supported by recent work from Longman, Lavric, and Monsell (2016), who showed that spatial attention is reconfigured with the task set and, specifically, that errors in attending can be mitigated by letting study participants advance trials at their own speed. Second, ASTRA allows for localized conflict resolution. In the current implementation, such localized resolution occurs only in the motion planning component, where action requests are held in abeyance for a period of time and those associated with the currently active task are preferred. However, we plan to include conflict resolution in the component that processes SR rules. We also expect to add conflict resolution when switching task sets and attentional strategies. These changes will help account for recent reports on the influence of localized conflicts (e.g., Longman et al., 2016; Regev & Meiran, 2016).

In addition, there is a natural extension of the attentional approach of ASTRA to models of contingent capture (Folk & Remington, 2006), where the task set influences the visual salience of goal-relevant objects. Experiments for investigating this phenomenon are typically not carried out in the task-switching literature, but they are relevant to visual search where the task of finding a particular object could influence early visual processing. For example, green items may be more likely to receive attention than others to someone who is looking for lettuce in a refrigerator. This potential extension highlights the usefulness of task set representations in ARCADIA beyond basic models of task switching. Furthermore, research in this area will help identify how mental states such as goals and intentions can influence perception. The central idea being that attention as a mediator of priorities in perception, cognition, and action serves as the integrator of mental processing across modalities (Watzl, 2017).

5.2 Relevance to Cognitive Systems

This idea brings us to the broader value that ASTRA has for cognitive systems research. From a practical perspective, cognitive architectures including Soar (Laird, 2012) and ACT-R (Altmann & Gray, 2008) as well as early expert systems (Lindsay et al., 1980; Buchanan & Shortliffe, 1984) have illustrated the value of productions in the encoding and application of knowledge. The task representation described here extends ARCADIA so that it can use productions in the form of its SR rules and enables research on how stimuli and responses interact with attention. For example, in ASTRA, the attentional strategy prioritizes responses that are associated with the currently adopted goal. Building on this capacity, we intend to extend the reported approach to enhance perceptual representations that match the conditions in the SR rules with the intention that this will draw attention to task relevant stimuli. This tactic is related to modeling contingent capture and can be generalized to cover enhancement of non-visual stimuli and interlingua elements as well. Adding productions to ARCADIA also eases the integration of those productions with perceptual processing. In contrast to Soar, which requires special modules for perceptual input, and ACT-R, which assumes that percepts are already encoded symbolically, researchers can add components for computer vision and audition to ARCADIA that have the same relationship to the framework as every other component. In other words, components that process sensory information have access to the same interlingua elements as the others in a model, and they produce interlingua elements that are visible to and usable by all the components. Consequently, integration of perceptual information is no different from the integration of any other sort of information in the system.

From a theoretical perspective, there has been a longstanding question regarding how to incorporate working memory into a cognitive system. For the most part, system developers rely on

Baddeley's model of working memory for inspiration (e.g., Baddeley et al., 2001; Baddeley, 2012). A central aspect of that model is its multicomponent nature, which many cognitive architectures and systems implement, including ARCADIA. However, as shown in ASTRA, the ARCADIA view of working memory is more distributed than typically seen. In addition to task working memory and the phonological loop that supports subvocalization that were discussed in this paper, ARCADIA models often include visual short-term memory, episodic memory, and a general working memory component for storing arbitrary interlingua elements. The contents of these memory components are themselves spread across multiple representations and components. In this paper, we saw that different parts of the active task are represented in task working memory, the SR link processor, and the focus selector. Moreover, when action responses are initiated, the motion planner stores the instantiated responses of the SR rules for multiple cycles. This approach lets researchers control the granularity of their representations and processes are for any particular working memory implementation. In ASTRA, this meant that we could quickly develop a model that closely satisfied the constraints of the current literature.

With that said, how do humans engage in flexible, goal-directed behavior in a dynamic world that contains an essentially infinite variety of obstacles, distractions, and other interferences? Our tentative answer is that they control the deployment of attention to goal-relevant stimuli and responses. By adopting a task, a person not only identifies a collection of behaviors that would be useful in achieving a goal but also sets priorities on patterns of perception and action. These priorities help focus activity while being able to take advantage of opportunities that arise in the environment. To illustrate, suppose you were walking to a soda machine to get a drink. If a friend walked up and offered you a soda, you might take it to save the trip and expense. Importantly, these priorities are mixed with ones that guide perception in untasked situations, which lets people respond to hazards that are salient enough to break through their focused activity. For instance, if you were on the same walk to the soda machine and someone stepped out in front of you, it would be best to notice them and respond instead of walking into them, regardless of the overarching task. Whether this answer is ultimately correct or specific enough requires further research, which is faciliated by an attention-centric framework such as ARCADIA.

In conclusion, we have made progress toward a model of task switching that will add a global capacity to represent and switch among tasks to the ARCADIA framework. Moreover, we are already using the functionality in separate models of causal reasoning and intentional action that are under development. To date, ASTRA has been influenced by recent and landmark results in cognitive psychology and neuroscience. We plan to continue this trend, with a particular emphasis on localized information processing that is bound together by the direction of attention. And from a cognitive systems perspective, we will use the work here to further test ARCADIA's utility as a framework for integrating theories of cognition, perception, and action.

Acknowledgements

We thank Andrew Lovett for discussions that led to improvements in the paper. The authors also acknowledge generous support from the Office of Naval Research under Award N00014-18-WX-0-0166. The views expressed in this paper are solely the authors' and should not be taken to reflect any official policy or position of the United States Government or the Department of Defense.

References

Altmann, E. M. (2013). Fine-grain episodic memory processes in cognitive control. *Zeitschrift für Psychologie*, *221*, 23–32.

Altmann, E. M., & Gray, W. D. (2008). An integrated model of cognitive control in task switching. *Psychological Review*, *115*, 602–639.

Baddeley, A. (2012). Working memory: Theories, models, and controversies. *Annual Review of Psychology*, *63*, 1–29.

Baddeley, A., Chincotta, D., & Adlam, A. (2001). Working memory and the control of action: Evidence from task switching. *Journal of Experimental Psychology: General*, *130*, 641–657.

Badre, D., & D'Esposito, M. (2009). Is the rostro-caudal axis of the frontal lobe hierarchical? *Nature Reviews Neuroscience*, *10*, 659–669.

Bridewell, W., & Bello, P. F. (2016a). Inattentional blindness in a coupled perceptual–cognitive system. *Proceedings of the Thirty-Eighth Annual Meeting of the Cognitive Science Society* (pp. 2573–2578). Philadelphia, PA: Cognitive Science Society.

Bridewell, W., & Bello, P. F. (2016b). A theory of attention for cognitive systems. *Fourth Annual Conference on Advances in Cognitive Systems* (pp. 1–16). Evanston, IL: Cognitive Systems Foundation.

Brown, J. W., Reynolds, J. R., & Braver, T. S. (2007). A computational model of fractionated conflict-control mechanisms in task-switching. *Cognitive Psychology*, *55*, 37–85.

Buchanan, B. G., & Shortliffe, E. H. (1984). *Rule-based expert systems: The MYCIN experiments of the Stanford heuristic programming project*. Reading, MA: Addison Wesley.

Charron, S., & Koechlin, E. (2010). Divided representation of concurrent goals in the human frontal lobes. *Science*, *328*, 360–363.

Folk, C. L., & Remington, R. (2006). Top-down modulation of preattentive processing: Testing the recovery account of contingent capture. *Visual Cognition*, *14*, 445–465.

Gallivan, J. P., Barton, K. S., Chapman, C. S., Wolpert, D. M., & Flanagan, J. R. (2015). Action plan co-optimization reveals the parallel encoding of competing reach movements. *Nature Communications*, *6*, 1–9.

Grange, J. A., & Houghton, G. (2014). Models of cognitive control in task switching. In J. A. Grange & G. Houghton (Eds.), *Task switching and cognitive control*, 160–199. New York: Oxford University Press.

Haith, A. M., Pakpoor, J., & Krakauer, J. W. (2016). Independence of movement preparation and movement initiation. *Journal of Neuroscience*, *36*, 3007–3015.

Laird, J. E. (2012). *The Soar cognitive architecture*. Cambridge, MA: MIT Press.

Lindsay, R. K., Buchanan, B. G., Feigenbaum, E. A., & Lederberg, J. (1980). *Applications of artificial intelligence for organic chemistry: The DENDRAL project*. New York: McGraw-Hill.

Logan, G. D., & Bundesen, C. (2003). Clever homunculus: Is there an endogenous act of control in the explicit task-cuing procedure? *Journal of Experimental Psychology: Human Perception and Performance*, *29*, 575–599.

Longman, C., Lavric, A., & Monsell, S. (2016). Self-paced preparation for a task switch eliminates attentional inertia but not the performance switch cost. *Journal of Experimental Psychology: Learning Memory and Cognition*, *43*, 862–873.

Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, *7*, 134–140.

Rangelov, D., Töllner, T., Müller, H. J., & Zehetleitner, M. (2013). What are task-sets: A single, integrated representation or a collection of multiple control representations? *Frontiers in Human Neuroscience*, *7*, 1–11.

Regev, S., & Meiran, N. (2016). Cue-type manipulation dissociates two types of task set inhibition: Backward inhibition and competitor rule suppression. *Psychological Research*, *80*, 625–639.

Vandierendonck, A. (2016). A working memory system with distributed executive control. *Perspectives on Psychological Science*, *11*, 74–100.

Watzl, S. (2017). *Structuring mind: The nature of attention and how it shapes consciousness*. Oxford, UK: Oxford University Press.