

Neuroscience of Technology

Dietrich Stout

Abstract

Although there is a burgeoning neuroscience of tool use, there is nothing that might be properly called a neuroscience of technology. This review aims to sketch the outlines of such a subject area and its relevance to the study of cultural evolution. Technology is itself an ill-defined term and is often taken to correspond loosely to human action that (a) involves the use or modification of objects, (b) displays a complexly organized multilevel structure, and (c) is socially reproduced. These characteristics may be better understood with reference to neuroscience research on perceptual-motor control, object manipulation, motor resonance, imitation learning, and goal-directed action. Such consideration suggests a number of biases which may affect the cultural evolution of technologies.

Introduction

In the extended analogy developed by Mesoudi et al. (2006), neuroscience is described as the “molecular genetics” of cultural evolution. That is, neuroscience is meant to provide a mechanistic understanding of the way in which cultural traits are instantiated as neural processes that can be replicated across individuals and expressed as (more or less) isomorphic behaviors. Mesoudi et al. recognize that strict analogies with the molecular mechanisms of DNA replication and expression are not likely to be very helpful; rather, the intended analogy appears to be at the level of the explanatory work that needs to be done in each case. Whether or not this genetic analogy proves fruitful with respect to research in social neuroscience, analogies between biological and cultural evolution proposed at higher levels of analysis (e.g., the selection and drift of cultural traits) will stand or fall on their own merit.

As it has matured, neuroscience research has moved away from attempts to identify simple one-to-one structure–function mappings and recognized the need for analyses in terms of dynamic and variable neuronal networks that are soft-assembled in response to context-specific task demands. Neural systems, like genetic and immune systems, are massively *degenerate*, meaning that

different structural elements can produce the same functional output (Edelman and Gally 2001; Price and Friston 2002). Conversely, many structural elements also appear to be *pluripotent* (e.g., Anderson 2010), meaning that single structures are capable of supporting multiple, different functions. Such one-to-many and many-to-one structure–function mappings are present across all levels of analysis (Edelman and Gally 2001), invalidating attempts to derive function from structure in a purely bottom-up way. Processes at the neural level cannot be properly understood without reference to higher-order functional and contextual constraints any more than the genetic information “coded” by DNA sequences can be understood apart from the larger processes of cellular metabolism, somatic development (Mayr 1994), and organism reproduction in which it is embedded. In neither case is it possible to assign causal primacy to a “replicator” identified at one fundamental level of analysis. Instead, our objective should be to identify multilevel constraints acting on the reproduction of behavior across individuals including, but not limited to, species-typical learning mechanisms and biases.

Neuroscience can contribute to this enterprise through an iterative research program in which structural and physiological correlates of behavior are used to inform the fractionation of psychological processes, and the fractionation of psychological processes motivates increasingly refined neuroscientific investigation. A neuroscience of technology would seem to be a good place to begin, considering the central place that technology (e.g., Boyd et al. 2011 and this volume) and technological artifacts (e.g., Shennan 2011 and this volume) have had in studies of cultural evolution, as well as the volume of neuroscience research devoted to understanding the perception, execution, and imitation of goal-directed interactions with objects. First, however, it is necessary to step back and consider what exactly we mean by “technology” in this context.

What Is Technology?

Technology is a fuzzy category. There can be little doubt that central examples like laptop computers and atlatls belong, but more peripheral examples raise questions. Is music a technology? What about a martial art or sign language? Is tool use by nonhumans “technology,” and how should we categorize complex foraging techniques which do not involve tools (e.g., Byrne and Byrne 1993)? Attempting an exclusive and exhaustive definition of technology is likely to be neither possible nor profitable. It would appear to be more important to identify key dimensions of variation in the “family resemblance” that links exemplars, so that these may become the subject of further study. In other words, we should concentrate less on this question (“What is technology?”) and instead openly explore the issue: “What is interesting about technology?” For current purposes, some of the more interesting things about technology are that it (a) often involves the use and modification of objects, (b) is characterized by

complexly organized goal structures, and (c) is heavily reliant on social mechanisms for its reproduction.

Objects

From an anthropological perspective, technology's tendency to involve objects is of practical interest because it increases the chance of behaviors leaving physical traces for future study. It is also of theoretical interest because it leads to creation of a durable medium for human action and interaction, broadly referred to as "material culture." From a somewhat narrower cognitive and neuroscientific perspective (e.g., Arbib et al. 2009), the involvement of objects is interesting because it complicates the perceptual-motor and cognitive control of action by introducing a wider array of potential affordances and effectivities to be discovered and coordinated and by requiring mechanisms of perceptual monitoring and internal modeling in the absence of direct somatosensory feedback from the end effector. Furthermore, the potentially greater temporal persistence and causal diversity of object-mediated actions may support the production of more complexly organized and temporally protracted action goals and sequences.

Complex Organization

Complex perceptual-motor and cognitive organization is a basic characteristic of technology, whether or not this complex organization is directly occasioned by the use of tools. For example, weaving a basket is classically "technological" even if no tools are used, whereas sweeping the floor is a more liminal example, even though a tool is being used purposefully to alter the physical environment. This largely reflects the intuition that basket weaving is a more complex and organized activity, but what exactly is meant by "organization" and "complexity" in this context? As recently discussed by Deacon (2012), an information theoretical approach would define organization and complexity as opposite extremes of a scale measuring the redundancy of a system. For example, a pattern of random static on a television screen is maximally *complex* because it is not constrained by any redundant patterning: there is no way to summarize the image on the screen without specifying the state of each individual pixel. It follows that complexity is also increased by increasing the total number of elements (e.g., pixels) or the number of different possible states of each element (e.g., colors). In contrast, many photographs contain *redundant* (i.e., predictable) patterns such as edges, fields, or gradients which allow "lossless" image compression (e.g., the LZW algorithm used to generate GIF, TIFF, and PDF files). These images are more *organized*, but less complex, than an equally sized field of static. A blank white screen is maximally organized and minimally complex because there is only one color option for each pixel. As it turns out, then, maximal complexity and maximal organization are both quite dull.

What is interesting (and characteristic of technology) is *complex organization*. This apparent oxymoron is achieved by opposing complexity to organization at different levels of abstraction; in other words, by exploiting the generative potential of hierarchical¹ systems (Lashley 1951; Mesoudi and O'Brien 2008c; Simon 1962). A classic example from the movement sciences is provided by Bernstein (1996), who described the arm motions of blacksmiths striking a chisel with a hammer. Surprisingly, Bernstein found that the movement trajectories of individual joints in the arms of these expert craftsmen were relatively unpredictable (i.e., “complex”) across swings. Nevertheless, these complex movements produced a redundant (i.e., “organized”) action outcome across swings in the form of a highly consistent trajectory of the hammer head. Thus, the repeated hammering action is well organized in terms of its consistently reproduced goal but remains complex in terms of its actual kinematic means. Moving upward in scale, one might similarly consider the assembly of redundant action “types” (e.g., hammering, heating, quenching) into complexly contingent action sequences that are again redundant on the still higher level of the standardized artifacts produced (e.g., Japanese swords, Martin 2000). This logic is quite familiar to stone knappers, who must produce standard products “based on raw material which is never standard, and with gestures of percussion that are never perfectly delivered” (Pelegrin 1990:117). In such cases, redundancy in technological outcomes actually *requires* variation in means. Note that it is the increasing abstraction of goals at higher hierarchical levels which supports this generative interplay of complexity and organization by allowing heterogeneous subordinate elements to constitute uniform superordinate goals. This is closely analogous to the way in which standard grammatical units (e.g., noun phrases) can be constituted from an infinite variety of different words. Complexly organized goal hierarchies are not only characteristic of technology, they are also critical in supporting both the adaptive flexibility and the social reproduction of technological behaviors (Byrne and Russon 1998; Wolpert et al. 2003).

Social Reproduction

The heavy reliance of technology on social mechanisms of reproduction is, of course, another one of its interesting features, and the primary focus of the current discussion. Loosely speaking, such reliance differentiates technologies from more biologically determined “instincts,” such as the dam building of beavers or the nest building of birds. We must be careful, however, to avoid

¹ In this chapter, “hierarchy” refers specifically to compositional containment hierarchies in which superordinate behavioral elements are constituted by subordinate elements; causation may be bidirectional (bottom-up and top-down). There is no implication with respect to the organization, hierarchical or otherwise, of neural or cognitive systems implementing these behavioral hierarchies.

simply resurrecting the sterile nature–nurture dichotomy in a new context. Social context plays a role in the development of all human behaviors, and no behavior is completely unconditioned by biologically inherited characteristics. Even a prototypically “innate” human skill like bipedal walking develops through an interaction of physical maturation and socially mediated opportunities for practice (Adolph et al. 2003). The crucial distinction is that the constraints organizing human gait derive almost entirely from the interaction of evolved biomechanical and neural traits with physical substrates rather than from social influences. Social scaffolding serves to motivate the developing system rather than to structure the behavior. Such social motivation may also be critical to technological reproduction (cf. Lave and Wenger 1991), but it is the social transmission of behavioral targets (i.e., multilevel action “goals” as discussed above) that is distinctive.

For example, the remarkable quality and consistency of Japanese swords is achieved by following an elaborate production recipe organized around a vast number of subgoals (Martin 2000), ranging from the kinematics of properly executed hammer strokes, to subtle perceptual cues indicating desired material properties and transformations, to abstract conceptual representations of various production stages or sequences. The consistent outcomes achieved by master sword makers do not simply “fall out” of interactions between some global goal (i.e., “make a sword”) and preexisting anatomical and environmental constraints in the way that the regularities of human gait do. Rather they reflect huge amounts of work (cf. Deacon 2012), both individual and social, invested in generating and regenerating the particular array of behavioral constraints that allow practitioners, starting from variable initial conditions and using procedures that are never identical, nevertheless to converge reliably on an astronomically unlikely outcome. As we have seen, this requires a hierarchical structure in which behavioral complexity can be simultaneously preserved and constrained on different levels of abstraction. The interesting question for the study of technology is: How do culturally constructed objects, situations, and social interactions come to constrain individual behavior in precisely this manner?

Low-Level Constraints: Tools and Actions

Tools

Human behavior is constrained not only by inherited somatic and neural structures, but also by an inherited cultural niche filled with predesigned tools. For example, the design of a bicycle affords only a very narrow range of effective actions, and no one needs to show an infant which end of a toy hammer to hold or which end is used to strike an object (Lockman 2006). The constraints imposed by designed objects generate behavioral attractors that can be reproduced across individuals and generations without necessarily involving

the sharing of internal action representations or conceptual knowledge. This property of artifacts is exemplified in patients who suffer from ideational or ideomotor apraxias; they have difficulty describing or pantomiming tool actions due to impaired internal representations but often produce appropriate grasps and manipulations when allowed the sensorimotor feedback from actually handling tools (Johnson-Frey 2004).

Seminal studies recording single neurons in the parietal cortex of macaque monkeys (Maravita and Iriki 2004) have shown that the use of simple tools is associated with a modification of the “body schema” to quite literally incorporate the hand-held tool as an extension of the body. Lesion data suggest that a similar mechanism is involved in simple human tool use (Berti and Frassinetti 2000). An extended period (2–6 weeks) of highly structured training (reinforcing successive elements of a behavioral chain) is required to produce such simple tool use in macaques (Iriki et al. 1996; Peeters et al. 2009) and may reflect the experience-dependent formation of new afferent connections from temporoparietal and ventrolateral prefrontal cortex (PFC) to neurons in the intraparietal sulcus (Hihara et al. 2006). Thus, macaque tool use would seem to rely on the adaptive flexibility of bodily representations in an occipitoparietal “dorsal stream” of vision-for-action (Milner and Goodale 2008). However, it is not clear that similar mechanisms can explain the much more diverse, pervasive, complex, and rapidly learned manual tool use of humans. In particular, it seems doubtful that body schema alterations alone are sufficient to explain the use of tools to alter the basic functional properties of the hand (e.g., knives, hammers, potholders) as is commonly seen in humans (Arbib et al. 2009) and perhaps also our closest living relative, the chimpanzee (Mulcahy and Call 2006; Povinelli et al. 2010).

Whereas a causal understanding of tool properties as distinct from the hand may not be necessary to explain macaque use of “simple tools” like rakes or pliers (Maravita and Iriki 2004; Peeters et al. 2009; Umiltà et al. 2008), such understanding is clearly implicated in the human use of “complex tools” which convert hand movements into qualitatively different mechanical actions (Frey 2007). It has been proposed that the human capacity for complex tool use arises from a novel integration in the left inferior parietal lobule of semantic representations of tool function from a ventral, occipitotemporal stream of vision-for-perception with the sensorimotor transformations for action in the dorsal stream, thus allowing for functionally appropriate tool prehension and use (Frey 2007). More recently it has been reported (Peeters et al. 2009) that an anterior region of human parietal cortex (anterior supramarginal gyrus, aSMG) displays a specific response to the observation of simple tool use. This region is located posterior to phAIP, the putative human homolog of the motor part of monkey anterior intraparietal sulcus (AIP) (Frey et al. 2005; Orban et al. 2006), a region involved in visuomotor grip transformations for object manipulation (Fagg and Arbib 1998). Human aSMG has been associated with the planning, pantomiming, and execution of actions with tools (Lewis 2006) and,

in contrast to phAIP, is activated by the observation of simple tool use but not of unassisted hand actions (Peeters et al. 2009). It thus appears likely that, in humans, aSMG plays a specific role in coding the casual properties of simple tools as distinct from hands, and it may represent an important convergence point for dorsal and ventral streams. Indeed, training in the use of novel tools results in increased activation of a converging network of ventral and dorsal stream structures, including an anterior portion of left intraparietal sulcus (i.e., in the rough vicinity of phAIP and aSMG) (Weisberg et al. 2007).

Apart from these details of functional neuroanatomy, it is clear that tools constrain action in at least two ways. First, simple tools typically present a small number (perhaps just one) of efficient options for grasping that constrain the way in which they will typically be incorporated into the body schema. These constrained affordances, which themselves reflect the cultural evolution of artifact design, will be reliably and repeatedly discovered across individuals even in the absence of more “active” social transmission. Second, the performance characteristics of “complex” tools, likely represented in ventral stream regions concerned with nonbiological motion (posterior middle temporal gyrus) and object form (fusiform gyrus), will constrain the range of actions for which they are typically used. Given a somewhat longer time of exploration, and perhaps some socially structured motivation, these performance characteristics should also be more or less reliably rediscovered across individuals with minimal other social input, much as the dynamics of bipedal walking are.

The idea that tools constrain human action may, at first blush, seem reminiscent of anthropological arguments that attribute causal agency to artifacts independent of human users (Gosden 2005) or to suggest that artifacts themselves are “active replicators” evolving in the same way that living organisms evolve. However, artifactual constraints on behavior are relational properties which only emerge in the context of goal-oriented action by living agents and, even in this context, typically result in the reproduction of simple behaviors (e.g., particular grips) rather than reproduction of artifacts themselves. This dynamic seems better captured by the concept of ecological inheritance developed in niche construction theory (Odling-Smee et al. 1996) than by strict analogies between artifacts and organisms. Artifacts can indeed embody reproducible information about their own form, use, and construction (e.g., Caldwell and Millen 2009), *but only under the goal-oriented interpretation of a living agent* (cf. Deacon 2012). Although the simplifying assumption that artifact taxa evolve “as if” they were biological taxa has been empirically productive (O’Brien et al. 2001; Shennan 2011), if we wish to understand the actual mechanisms involved or to identify and address cases in which this simplifying assumption might not work, then we must consider the nature and transmission of constraints on interpretation by such agents.

Internal Models

Although an organism's physical environment constitutes a vast array of constraints organizing behavior (or, in positive terms, an array of "affordances" for action), it is also the case that many actions unfold too quickly to be guided by online sensory feedback and error correction. It is thought that this limitation is overcome through the use of *internal models* which predict movements and outcomes in advance of sensory feedback (Wolpert et al. 2003). More specifically, *forward models* predict the sensory consequences of motor acts (i.e., model action outcomes) whereas *inverse models* predict the motor commands necessary to produce a given action (i.e., model bodily states and transformations). It is perhaps more intuitive to refer to these two types of internal model as *predictors* and *controllers*, respectively. Briefly, predictors developed through prior experience can be used to select appropriate controllers in advance of actual sensory feedback, with *post hoc* comparison to actual outcomes allowing for error correction through the elimination of inaccurate predictors. For example, one reaches to pick up a full tea kettle expecting to require a certain amount of muscular force; upon finding the kettle empty there is a rapid reevaluation and correction. The concept of internal models developed out of computational modeling studies of motor control, and the question of their actual neural instantiation is a highly complex and controversial one. Perhaps the most consistently implicated structure is the inferior parietal cortex, which appears to play a key role in the integration of sensory and motor information, for example, during object manipulation (Arbib et al. 2009), the use of subvocal articulation to support speech perception (Price 2010), and the central cancellation of the sensory consequences of self-tickling (Blakemore et al. 1998). Such integration is obviously critical to imitation, in which the sensory consequences of others' actions must be matched to appropriate motor commands for self-execution (Wolpert et al. 2003), and numerous studies have confirmed inferior parietal cortex involvement in imitation (Buxbaum et al. 2005; Chaminade et al. 2005).

Motor Resonance

Internal models thus provide a useful framework for understanding the imitation of simple actions. The neuroscience of imitation and social cognition has been massively impacted by the description and study of *mirror neurons* in the inferior frontal and parietal cortex of macaque monkeys (Rizzolatti and Craighero 2004). These are neurons that respond both to observed actions and the self-performance of a similar action. Neurons with similar properties are thought to exist in humans (although the invasive recording techniques used in monkeys cannot be applied to humans to confirm this directly) and to reflect the direct mapping of motor representations of one's own actions to sensory representations of the actions of others. In the language of internal models,

predictors are compared with sensory perception of others' actions rather than feedback from one's own actions and then used to select the appropriate controllers needed to generate isomorphic movements (Wolpert et al. 2003). These activated controllers can then be executed to produce the actual movement (imitation) or suppressed to model the movement without overt action. It is thought that this automatic activation of motor controllers, or *motor resonance*, provides a mechanism of action understanding through internal simulation, and that this basic mechanism of action understanding is the foundation for more sophisticated forms of social cognition, including the understanding of intentions (Wolpert et al. 2003) and theory of mind (Gallese et al. 2009).

Imitation

There are several important issues for this account of action understanding and imitation, particularly with respect to the social reproduction of technology. Most fundamentally, this account requires a mechanism whereby predictors incorporating rich somatosensory or kinesthetic feedback from one's own body can be matched with the purely visual and auditory input generated by the actions of others: input which is not even presented in the same spatial perspective. This is the so-called *correspondence problem* and, although special purpose mechanisms have been proposed, a prevailing view is that it is solved by general purpose mechanisms of associative learning (Brass and Heyes 2005). Thus, internal models are linked to the observed behavior of others through stimulus generalization from one's own visible movements (i.e., recognizing that one's own hand posture is "the same" as that of another individual) and/or associated contextual cues, supported by simple mechanisms such as Hebbian learning (i.e., synaptic plasticity: "cells that fire together, wire together"). An elegant neural network model of this process is presented by Laland and Bateson (2001). However, this simple solution to the correspondence problem raises another issue: if imitation is enabled by the activation of associations with already existing internal models, how is it possible to imitate novel actions? In other words, how is imitation learning achieved?

This is a critical problem for the social reproduction of behavior generally and of technology specifically. The solution likely lies in the hierarchical structure of behavior. Because action goals are abstractions over constituent means (note that the means/goal distinction is a relative one: "goals" become "means" at a different level of analysis), it is possible to assemble new behaviors from familiar constituents (Buccino et al. 2004). Furthermore, because the relative abstraction of goals renders them robust to variation in lower-level means, it is possible to "copy" behaviors without matching these details perfectly. Indeed, it would never be possible to imitate precisely the kinematics of other individuals with differently sized and shaped bodies (de Vignemont and Haggard 2008). Conversely, it is quite possible to imitate a gesture with an entirely different effector (e.g., right vs. left hand). In cases where such variation in

lower-level action details does not make a difference to intended outcomes (i.e., is unconstrained), the purely observational learning of novel actions may be possible through statistical parsing of multilevel behavioral regularities (Buchsbaum et al. 2011; Byrne 1999), thus supporting fast imitation learning. However, when action details do matter and these action elements are not already in the behavioral repertoire, it will be necessary to engage in active behavioral exploration (i.e., practice or play) to (re)produce effective internal models through an iterative process matching self-actions to ever closer approximations of observed kinematic and/or environmental outcomes. This hierarchical model addresses the issue of novelty but raises a final issue for resonance accounts of imitation: Exactly what kind or level of information is being shared during imitation learning?

Many of the possible kinds of representations that might be shared during action observation and imitation have been reviewed by de Vignemont and Haggard (2008), who distinguish various types (sensory vs. motor, semantic vs. pragmatic) and degrees of abstraction (from specific motor commands and sensory predictions to abstract *prior intentions* specifying the global goal of an action, as in “to drink from a glass of water”). They conclude that resonance or mirror mechanisms most likely involve pragmatic, motor *intentions in action*, defined as “dynamic sequences of specific movements.” These intentions in action are more specific and less interpretive than prior intentions (e.g., “grasp object and bring to mouth” as opposed to “eat”), but are not specified to the level of particular motor commands. This is consistent with work in monkeys which shows that mirror neurons are selectively responsive to action types (e.g., grasp) despite substantial variation in the motor details (e.g., precision vs. power grip) (Rizzolatti and Craighero 2004), and that the response of inferior parietal mirror neurons coding for particular action types is modulated by the final goal of the motor sequence in which the actions are embedded (Fogassi et al. 2005). In humans, putatively homologous regions of inferior frontal and inferior parietal cortex display a similarly selective response to simple action goals (i.e., grasp a particular object) across variation in specific kinematic means (Grafton 2009). Thus it appears that classic motor resonance mechanisms operating in anterior inferior parietal cortex and posterior inferior frontal cortex achieve a best-fit matching of observed actions to “mid-level” internal models (representing goal-directed sequences of elementary actions, Wolpert et al. 2003) that are already in the motor repertoire of the observer.

This mechanism for observational action understanding (Rizzolatti and Craighero 2004), working together with cortical regions that support motor planning (dorsal premotor), spatial awareness (superior parietal), biological motion perception (superior temporal sulcus), object representation (inferior temporal), and working memory (middle frontal gyrus), may support imitation of simple goal-oriented actions (Buccino et al. 2004; Menz et al. 2009; Molenberghs et al. 2009), such as reaching to grasp or strike an object. Interestingly, Hecht et al. (2012) recently reported a pattern of increasing

connectivity of “core” mirror-system and related temporal-parietal regions across (nonimitating) macaques, (infrequently imitating) chimpanzees, and humans.

Imitative matching of shared internal models may constitute a basic unit for the social transmission of behavior. As a matter of speculation, it seems possible that particular internal models widely shared in a population for one reason (e.g., manipulating chopsticks) could bias the perception and reproduction of unrelated technical gestures. However, this mechanism by itself is insufficient to explain the transmission of complexly organized technologies, which typically require fidelity both at lower (e.g., embodied skills) and higher (e.g., sequences of goals) levels of action organization. This is likely to implicate additional mechanisms of more abstract goal representation in the PFC, as well as more concrete sensory processing of observed movements. For example, it is increasingly well documented that the “elementary” percussive gesture of stone knapping requires a highly coordinated and precise strike (Bril et al. 2000, 2010). In other words, kinematic variation is highly constrained by desired outcomes. However, the important parameters (e.g., kinetic energy) are not perceptually available to naïve observers nor captured in existing internal models for more generic percussive acts. Thus, the observer must begin by (incorrectly) imitating the observed gesture, checking the outcome against the predicted (desired) outcome, and then embarking on a lengthy process of goal-oriented behavioral exploration or deliberate practice (Ericsson et al. 1993) to (re)discover the relevant task constraints and develop corresponding internal models. Because there are a huge number of variables to be explored, this skill acquisition process may be quite lengthy. This process may be accelerated somewhat by continued observation of expert performance, which provides a sensory model to be matched through processes of stimulus generalization and associative learning discussed above, or through intentional teaching by a mentor, which might involve ostensive cues and/or modified performance (demonstration) to highlight relevant variables, structured coaching (Vygotsky 1978), or explicit semantic information about the task. In this context, social motivation for practice (implicating additional social cognitive and affective mechanisms) may also be critical for technological reproduction (Lave and Wenger 1991; Roux 1990; Stout 2002, 2010) although this is beyond the scope of the current review.

This account of best-fit action matching followed by individual kinematic rediscovery also requires the presence of more abstract goal representations. The initial (mis)matched internal model provides a starting point for generating a range of behaviors; however, subsequent selection on this behavioral variation can only occur with reference to a desired outcome. In line with a widely held distinction between imitation (of means) and emulation (of goals), neuroscience studies of imitation have used tasks in which the goal is simply to produce a particular movement or body posture (Buccino et al. 2004; Chaminade et al. 2005; see also, the do-as-I-do imitation task in nonhuman primates in

Custance et al. 1995). This avoids theoretical complications surrounding the imitation–emulation dichotomy by equating goals and means. However, technological action is characteristically organized (i.e., constrained) with respect to higher-level goals, most typically involving transformations of objects or substrates. How are these higher-order constraints represented and reproduced across individuals?

High-Level Constraints: Objects and Goals

The word “object” can refer to either a tangible physical entity or the goal of an action. This etymological association highlights the simple fact that action goals typically have to do with modifying, manipulating, or using physical objects. This is certainly true of “technological” actions. As noted in the introduction, the systematic transformation of durable objects provides an important medium for the elaboration of temporally extended, complexly organized action sequences. Unfortunately, and in sharp contrast to the growing body of research on (simple) tool use (e.g., Lewis 2006), there has been almost no neuroscientific investigation of such object transformation.

Action Outcomes

Our understanding of object-related goal representations thus comes primarily from studies of simple observation and manipulation. Such studies consistently implicate prefrontal and parietal cortex, with more abstract goal representations generally being associated with greater distance from the primary sensorimotor cortex surrounding the central sulcus (i.e., more anterior in frontal cortex, more posterior in parietal cortex). In monkeys, Nelissen et al. (2005) found that more posterior inferior frontal regions responded to variation in the specific context of observed actions (agents, effectors and actions) whereas the more anterior area 45B (a putative homolog of human area 45, i.e., anterior Broca’s area) was responsive to objects. More directly relating to object transformation, Hamilton and Grafton (2008) studied the representation of observed action outcomes in humans using a simple sliding-top box which could be opened in different ways (variation in outcome) and with different motions (variation in kinematic parameters). They found a selective response to outcomes in inferior frontal cortex (area 44, posterior to area 45) and inferior parietal cortex, both in the right hemisphere. This contrasts with a previous study by Grafton and Hamilton (2007) in which variation in the target object for a simple reach-to-grasp was associated with selective responses in inferior frontal cortex and a relatively anterior portion of left inferior parietal cortex (i.e., phAIP, discussed above), both in the left hemisphere. Across experiments, variation in low-level kinematics was associated with response in visual association cortex (implicating sensory matching rather than motor resonance).

These findings suggest the presence of a diversity of goal representations in human parietofrontal cortex (Grafton 2009), with higher levels (e.g., material outcomes) in particular being represented in the right hemisphere.

Involvement of the right hemisphere may seem surprising, considering the extensive evidence that simple tool use is left lateralized (Lewis 2006). However, action outcomes, particularly those which involve object transformation, may unfold on a longer temporal scale and involve larger-scale visuospatial processing, both of which may be preferentially associated with the right hemisphere (Stout et al. 2008). In fact, right hemisphere involvement has been consistently reported in the small number of imaging studies that have actually studied object transformations. Chaminade et al. (2002), who conducted a PET study of subjects imitating simple construction actions using Lego blocks, reported activation of the right dorsolateral PFC in cases where only partial information (goal only or means only) was available to guide action planning. Frey and Gerry (2006) had subjects learn by observation how to assemble different objects from Tinkertoys and found that the right anterior intraparietal sulcus was the only region in which activation correlated with successful imitation of the demonstrated (arbitrary) sequences of assembly. They concluded that this region is important in forming representations of the temporal ordering of component actions. This is consistent with patterns of impairment to complex action sequencing observed following right hemisphere lesions (Hartmann et al. 2005).

Finally, in the only imaging studies to date of actual technological production, a series of PET and fMRI investigations of stone tool making (Stout and Chaminade 2007; Stout et al. 2008, 2011) have consistently reported right hemisphere activation. Stout et al. (2008) found increased activation of right inferior parietal and inferior frontal cortex during skilled handaxe production as compared to simple Oldowan flake production, as result mirrored by fMRI data from the observation of tool production (Stout et al. 2011). This does not appear to reflect the presence of low-level differences in manipulative complexity across the two tasks (Faisal et al. 2010) and is thus attributable to the more complexly organized goal structure of handaxe production.

Goal-Organized Imitation

Though more study of object transformation is clearly needed, the emerging picture suggests that a right lateralized parietofrontal network is involved in representing goals at the level of discrete action outcomes and sequential object transformations. This would be in addition to a better-known left hemisphere motor resonance system involved in representing object-directed intentions in action, as discussed in the previous section. This distinction parallels the differential roles of left versus right hemispheres in rapid (e.g., phonology, syntax) versus slow (e.g., prosody, context) linguistic processing and may reflect a more general hemispheric division of labor between rapid, small-scale

action control on the left and large-scale, longer duration integrative functions on the right (Stout and Chaminade 2012). In any case, attention to action outcomes and object transformations is likely to be at least as important as motor resonance in the social reproduction of behavior. As argued above, selection among potential internal models for simple actions can only occur with reference to higher-order goals. It is likely that simple object-transformation targets often provide this constraint on action variation.

For example, Richard Byrne has proposed an influential “string parsing” model of imitation, in which “recurring patterns in the visible stream of behavior are detected and used to build a statistical sketch of the underlying hierarchical structure” (Byrne 1999:63). In this cognitively simple way, observers pick out essential actions or stages (i.e., constraints indicated by redundant patterning) from complex observed sequences, allowing for “program-level” copying (Byrne and Russon 1998) of large-scale behavioral organization. Byrne suggests that these redundancies might be bodily movements or effects on objects, but that the latter are likely to be much more easily observable. As we have seen, more or less separable neural systems exist for representing each of these levels of action organization. Importantly, Byrne’s model specifically does not require causal or intentional interpretation of the observed actions; such understandings may be developed later if at all. All that is required are opportunities and motivation for repeated observation. In addition to broader issues of social context alluded to above, this implies the need for at least one additional level of abstraction in goal representation: that of the overall goal of the sequence, the desirability of which motivates attention to and copying of subgoals. In fact, developmental (Bekkering and Prinz 2002; Flynn and Whiten 2008b) and experimental transmission studies (Mesoudi and Whiten 2004) indicate that there is a bias toward higher-fidelity copying at higher levels of hierarchical organization.

Such higher-order goal representations are most likely supported by PFC. Neurophysiologically, PFC is well suited to maintain stable superordinate goal representations over extended subordinate action sequences because prefrontal neurons are able to sustain firing over extended periods of time and across events (Barbey et al. 2009). Connectionally, PFC represents a high-level convergence zone for the brain’s sensorimotor systems and is thought itself to be organized in a multilevel fashion, with increasingly abstract representations being instantiated in increasingly anterior regions. The precise meaning of “abstraction” in this context remains controversial, with some nonexclusive alternatives being domain generality (integrating across cognitive domains), relational integration (relations between stimuli, relations between relations, etc.), temporal abstraction (maintaining goals over time), and policy abstraction (representing goals as abstractions over subgoals) (Badre and D’Esposito 2009). This raises the question of exactly what constitutes higher-order goals of technological action, and what needs to be represented.

Classically, technology is associated with achieving economic goals and, ultimately, with increasing the quantity and efficiency of energy capture from the environment (White 1943). In everyday life, however, such ultimate aims may be very far from the mind of individuals. More proximate goals are likely quite diverse, from the pursuit of a valued item or commodity (money, a useful tool) to an assertion of personal identity and pursuit of social status. In the case of technological skill acquisition, the latter is perhaps most often the motivating goal (Lave and Wenger 1991); it is even likely that the “value” attached to many items is itself socially motivated. These are questions largely beyond the scope of this chapter, but it is important to note that they do matter and are quite relevant to understanding the neural systems and mechanisms which may be involved. This is particularly true if one accepts the argument that there is a very important “top-down” or goal-directed element in imitation and technological transmission. For example, a fascinating intergenerational study of Zinacantec weavers (Greenfield 2003) found that the introduction of a market economy in the region changed the goal of weaving from the fulfillment of traditional social relations to commercial profit. This introduced a new subgoal of producing innovative designs, which was in turn realized by an increase in the complexity of conceptualization and manipulation of designs (from unitary blocks of color to a thread-by-thread basis). Thus, changes in high-level goals can have important “trickle-down” effects on all levels of action organization and technological transmission.

Structured Event Complexes

Structured event complex (SEC) theory (Barbey et al. 2009) provides a model of PFC function that might be useful in addressing the fundamental diversity of potential technological goals. An SEC is defined as “a goal-oriented set of events that is structured in sequence and represents event features (including agents, objects, actions, mental states and background settings), social norms of behavior, ethical and moral rules, and temporal event boundaries” (Barbey et al. 2009:1292). This information is “stored” throughout the brain in the form of embodied sensory, motor, and visceral associations or *feature maps* that are integrated into more abstract cross-modal representations in associative *convergence zones*, such as PFC. Barbey et al. (2009) review neuroscience evidence linking major dimensions of variation in SECs to functional gradients and regions in PFC. Thus, it is expected that more complex SECs (e.g., more relations, greater policy abstraction) will be associated with more anterior PFC, multiple event integration with right PFC, mechanistic plans and actions with dorsolateral PFC, and social norms and scripts with ventromedial PFC. This range, coupled with widely distributed representations in posterior cortex, suggests that virtually the entire brain may become involved in different SECs and in different activities we consider “technological.”

According to SEC theory, once the complex web of associations that constitutes an SEC is formed, the entire SEC can be activated in a bottom-up fashion by any of its elements. For example, upon walking into a restaurant, a wide array of expectations and intentions become active. Similarly, the need to drive a nail into a board can stimulate an entire SEC having to do with finding and using an appropriate tool. Thus, SECs are substantially similar to the scripts and schemas of cognitive psychology (e.g., Abelson 1981), the cultural models of cognitive anthropologists (Holland and Quinn 1987), and the concepts of *constellations of knowledge* and *umbrella plans* specifically used to describe technological proficiency (Keller and Keller 1996). Their acquisition and refinement may thus be related to cognitive development in children (e.g., Bruner 1990) and the acquisition of technical expertise in adults (Keller and Keller 1996).

The activation of related representations in SECs also resembles the priming of semantic concepts (e.g., “dachshund” primes “poodle” and “leash,” Patterson et al. 2007), but involves a wider diversity of representation types (e.g., internal models). Like webs of semantic association, different SECs would be overlapping, with fuzzy boundaries and ambiguous membership (e.g., should instant messaging etiquette be similar to a phone call or email) leading to substantial possibilities for creativity (i.e., a generative system). Indeed, a core function of SECs is thought to be the support of counterfactual thinking (e.g., predicting future outcomes, planning contingencies, and imagining alternatives). SEC theory is thus a promising direction for exploring the little known neuroscience of creativity, and for better understanding the source(s) of technological innovation.

Conclusion

The neuroscience of technology should be of interest to many different people for many different reasons. In this chapter, implications for our understanding of the cultural evolution of technologies have been highlighted. Of specific interest are mechanisms that might (a) bias or constrain the reproduction of technological practices and/or (b) generate new technological variants (complexity). With respect to the former, analysis suggests that there should be biases operating at multiple levels of organization. Starting at the bottom, there are the constraints imposed by designed tools themselves that strongly tend to encourage particular patterns of prehension which, in turn, affects the way simple tools are incorporated into the body schema, and thus their most likely patterns of usage. The existence of such constraints, even in archaeological tools of unknown function, could be directly discovered through experiments with modern subjects. At a somewhat higher level, the performance characteristics of tools also constrain usage. However, these constraints will tend to be less determining relative to the vast range of potential activities that could be

imagined: top-down goals and context become more important in narrowing the search space, presenting greater problems for experimental archaeologists.

Research on motor resonance suggests that simple goal-oriented gestures, such as reaching to grasp, place, or strike an object, may be the basic units of imitation and social transmission. The best-fit matching of observed gestures to functionally similar gestures already in the observer's motor repertoire likely presents a common occasion for biased transmission, where a more generic gesture is substituted for a specialized one. Stabilizing selection against such substitutions is likely an important dimension of technological apprenticeship and skill acquisition (cf. results bias; see also Shennan, this volume). This implies a top-down influence of desired goals constraining low-level action. It is likely that the most easily observable and transmissible goals providing such detailed constraints on action sequences are often visible transformations of objects. This creates a possibility for transmission biases against transformations that are subtle or otherwise difficult to observe. This will be the case, in particular, where the observer's causal/intentional understanding of the process is not well developed. For example, novice stone knappers often make mistakes with respect to subtle techniques, like platform preparation, and may not be able to tell that the nice large flake they just removed actually does not get them any closer to their goal. One way to counter such transmission bias in poorly understood systems is to slavishly "overimitate" all aspects of production, which can promote technological stasis (Martin 2000).

Indeed, the biases/constraints just discussed can be viewed in two ways. First, all other things being equal, they will tend to bias the kinds of technology that are successfully reproduced (i.e., against those with highly specialized gestures and cryptic elements). Second, however, they provide a key to understanding the types of errors that are most likely. Assuming that some such errors are actually beneficial, this might be seen as somewhat akin to understanding the mechanisms of mutation in biological evolution. In this way, these constraints can also act as an accidental mechanism generating novelty.

A major disanalogy with conventional conceptions of biological evolution is that technological innovation can also be intentional and goal oriented. The neuroscience of such creative action is not well understood, perhaps because creativity is itself hard to define and operationalize experimentally. It also seems likely that the generation and preservation of technological innovations will be more profitably studied as a social phenomenon than a psychological one. Nevertheless, a common conception of innovation by individuals is that it represents the recombination of existing ideas into a new framework. Insofar as this is the case, SEC theory, which addresses precisely the mechanisms of association between elements of goal-directed action, may provide a good framework for investigating the issue.

From “Cultural Evolution: Society, Technology, Language, and Religion,”
edited by Peter J. Richerson and Morten H. Christiansen. 2013. Strüngmann Forum Reports, vol. 12,
J. Lupp, series editor. Cambridge, MA: MIT Press. ISBN 978-0-262-01975-0.