Chapter 7
Development of the Quantified Human

Morley O. Stone, Jack Blackhurst, Jennifer Gresham, and Werner J.A. Dahm

Abstract A *Sense-Assess-Augment (SAA)* framework – originally outlined by Galster and Johnson (Sense-assess-augment: a taxonomy for human effectiveness. Technical report. United States Air Force Research Laboratory, Wright-Patterson Air Force Base, 2013) and based loosely on the adaptive system framework of Feigh et al. (Hum Fact 54(6):1008–1024, 2012) – is presented for approaching augmentation of human performance. While the SAA framework has broad application across all three elements of human-computer interaction, including the machine, the human-machine interface, and the human operator, here we focus on its role for human performance augmentation. SAA begins with the human, sensing their physical, physiological, and psychological state. *Sensing* is the most mature piece of the SAA paradigm, because it leverages the considerable commercial investments in wearable sensors for athletics, healthcare, and human productivity. As a result, sensors exist or are in development that can measure a wide range of physiological parameters, such as brain activity, eye movement, skin temperature, and increasingly biological performance markers, such as blood glucose levels and molecules like orexin that indicate the onset of fatigue. *Assessment* involves aggregation of data from multiple sensors, algorithmic processing of the data, and correlation of the results to behaviors and actions of interest. The challenge is to empirically make sense of the data in relation to baselines that vary between and within individuals, and the needs of a task at hand that is shared by both human and machine and

M.O. Stone (✉) • J. Blackhurst
Air Force Research Laboratory, Wright-Patterson AFB, OH 45433, USA
e-mail: morley.stone@us.af.mil

J. Gresham
Florida Institute for Human and Machine Cognition, Ocala, FL, USA

W.J.A. Dahm
Security & Defense Systems Initiative (SDSI), Arizona State University, Tempe, AZ 85287-5604, USA

P. Artemiadis (ed.), *Neuro-Robotics: From Brain Machine Interfaces to Rehabilitation Robotics*, Trends in Augmentation of Human Performance 2,
DOI 10.1007/978-94-017-8932-5_7,

that may occur both in real time and across the human lifetime. Finally, based on the assessment, appropriate *augmentation* is delivered, which can take many forms, including redistribution of tasks from man to machine, changes in the operating environment, influences from external hardware, or even the growing use of “electroceuticals” – the use of electric stimulation to augment performance. The SAA framework provides a way of approaching human performance augmentation that is consistent with and leverages the emerging understanding of how humans can interact effectively with autonomous systems in an entirely new socio-technical dynamic.

7.1 Introduction

In 1961, President Kennedy issued a call to Congress and the Nation to put a man on the moon by the end of the decade. So started a decade of incredible advancements commensurate with the “space race”—catapulting computer science and engineering into a central developmental role. The previous decade, John McCarthy had coined the term “artificial intelligence” paving the way for the promise of robots, intelligent machines and other kinds of autonomous systems that would drastically reduce workload and the eliminate drudgery. Optimism abounded as nascent scientific and engineering fields were driven to the design of systems capable of transforming society as we knew it.

The pace of fulfillment on that promise has been astounding, though it has not unfolded as expected. In a mere five decades, computers and robots have become a fixture in everyday life. From robotic vacuum cleaners to wearable computers to the wholesale revamping of how we do our work, the push for smaller, faster, and more intelligent machines has been a success beyond expectation. What’s missing is the concurrent reduction in human workload. The effect has been dubbed the “autonomy paradox”, where the very systems designed to reduce the need for human operators require more manpower to support them. For example, unmanned aircraft carry cameras aloft without a crew, yet require multiple shifts of operators, maintainers and intelligence analysts on the ground to extract useful data. Current estimates for one particular system, the MQ-1/9, show that with nearly 20,000 h of full-motion video being collected each month, analysis personnel outnumber aircrew by 8 to 1 [1].

In recent decades, advancements in the human and biological sciences have been as rapid and transformative as the information and computer sciences. The Human Genome Project has provided a wealth of information that is already changing the way we think about medical care delivery and human performance. And the new White House BRAIN Initiative [2] brings with it the potential for breakthroughs in neuroscience, biology and computational science—much as President Kennedy’s space challenge did for computer science and engineering.

One could view the ongoing interaction and intersection of these transformative areas as the basis for increased speculation around the art-of-the-possible with respect to human performance augmentation (HPA). Below, we present a framework for HPA that borrows heavily from ongoing revolutions spanning mobile health (mHealth) to quantified self (QS) movements.

In this chapter, we discuss a Sense-Assess-Augment (SAA) framework for reviving the role of the human in the development of autonomous, interdependent systems. This will require research and development at each of the touch points: the machine, the human-machine interface, and the human operator. Artificial intelligence and autonomy is where much of the work on the machine is ongoing and much progress is being made. There is tremendous need and opportunity to improve the human side of the equation, and although we will discuss all three pieces of human-computer interaction (HCI), we shall devote most of our focus to human augmentation.

In order to understand the role of the human in future autonomous systems, it is important to draw a distinction between automation and autonomy. Automated processes are now common, particularly in areas like manufacturing, and involve execution and control of a narrowly defined set of tasks without human intervention. Automation is used when the parameters are well-defined and the environment is highly structured. In contrast, autonomous systems can perform tasks in an unstructured environment. Such a system is marked by two attributes: self-sufficiency—the ability to take care of itself—and self-directedness—the ability to act without outside control. Most technological developments today, including "unmanned" air systems, would still be more classified as automation rather than autonomy. For systems to become more autonomous, that is, more self-governing, they will require some type of basal reasoning capability. Reasoning is critical in order to deal with the main sources of brittleness in our systems today: dynamic, complex requirements and environments.

Reasoning, however, does not mean the machine acts alone. The question is not "what can machines do without us?" but "what can machines do with us?" Consider that it was not until the 1990s that the "I" in HCI switched from "interface" to "interaction [3]." It is only recently that adaptive computing based on the executive, affective, and conative state of the user has risen in prominence.

Unfortunately, as Aryeh Finegold noted in 1984, "One of the big problems is the tendency for the machine to dominate the human [4]." Sadly, despite our progress in fields such as artificial intelligence and autonomy over the last several decades, this is still true. This is due in large part to what we call the "leftover" principle of interface design, where the goal is to automate as much work as possible as the human adapts to whatever is leftover. This produces a rigid interaction lacking both in transparency and a bi-directional understanding of intent. The point is that the design parameters for an interdependent (not dependent) human-machine system look very different than a machine designed to maximize autonomy. Johnson et al. proposed [5] that interdependent systems should possess mutual awareness (context), consideration (adjustability) and the capability to support (reciprocation). This means we must design systems such that a machine not only provides support for others' dependence on it but can also deal with its own dependence on others.

7.2 Historical Context and Overview of the Sense-Assess-Augment Model

In 1948, Claude Shannon published his seminal work on information theory [6], describing an ideal communications system where all information sources have a source rate, and the channel through which the source's data travels has a capacity, both of which can be measured. Information can be transmitted only if the source rate does not exceed the channel's maximum capacity, now known as the Shannon limit. This had important implications with the advent of radar, one of the first exquisite pairings of man and machine, where in order to get the system to work at optimal levels, the operator had to be trained in how to discriminate the appropriate signal from an incredibly noisy background. By 1954, signal detection theory was being formally applied to the study of perception and recognition [7].

Shannon's work heralded the start of what psychologists now refer to as the cognitive revolution, spurring the idea that information processing could be used to describe the human as a system consisting of interacting subsystems, each of which operated with various capabilities and capacities. This was a time when the study of human performance was becoming increasingly interdisciplinary [8], with significant influences from the fields of psychology, linguistics, and computer science. As Proctor and Vu [9] put it: "Given the close relation of the information-processing approach to computers and artificial intelligence, and given the view that both humans and machines can be conceived of as being types of symbol manipulators, it seems only natural that the information-processing approach has provided a primary basis for understanding and analyzing human-computer interaction (HCI)." This not only allows for a common lexicon between those studying both humans and machines, but breaks the human-computer interaction into machine and human subsystems which can be analyzed either separately or together.

World War II also heralded a time of immense improvement in aircraft design, accompanied by a feverish rollout of new aircraft models. In that race to production, human factors issues were often overlooked. As an example, there were several known aviation mishaps, where pilots confused landing gear knobs with flaps. This is what spurred the innovative work of Lieutenant Colonel Paul Fitts. In 1954, he published what became known as Fitts' Law [10], a quantitative model relating the speed-accuracy trade-off associated with pointing, whereby targets that are smaller and/or further away require more time to acquire.

$$T = a + b \log_2(1 + \mathrm{D}/\mathrm{W}) \tag{7.1}$$

With T equal to the average time required to complete the movement and D (distance) over W (width of target) as a proxy for accuracy; *a* and *b* are device dependent constants.

Fitts' law showed a linear relationship between task difficulty and movement time that has proved to be remarkably robust. Although there have been minor modifications since then, the mathematical relationship applies under a variety of conditions, with different limbs, and holds true even without overt motor movements [11]. More fundamentally, it advanced Shannon's work by providing the first empirical determination of the information capacity of the human motor system [12]. Providing a mathematically valid description of human performance was not only revolutionary for its time, it continues to be relevant and advantageous today.

Despite the cognitive revolution, much of the human sciences have remained rooted in empirical versus theory-driven studies. Indeed, as critical as the development and application of information theory has been, it has largely remained more descriptive than explanatory [13]. MacKenzie [14] stated that "despite being robust and highly replicable, Fitts' law remains an analogy waiting for a theory." In many ways, the real fruits of the cognitive revolution have yet to be picked. What distinguishes an engineering discipline is an objective feedback control mechanism. What's needed now is to "close the loop," where the physical and mental states of the operator are fed back into the machine, making the human a more seamless part of the overall system.

James Watson, in his explanation of the goals of behaviorism, says, "Its theoretical goal is the prediction and control of behavior [15]." Without that, as Donald Kennedy so aptly put it, neuro-imaging is akin to post-modern phrenology [16]. As Proctor and Vu stated in their review of Human Factors research progress, "One implication of an emphasis on paradigm shifts is that past research is of little relevance because it is from 'old' paradigms. This view is reinforced within human factors because the field deals with new, increasingly sophisticated technologies" [17].

What we propose in this chapter is a framework that reconciles the behaviorists' demand for objective data with the cognitive desire to understand mental processes directly. If one hopes to design human performance with the same precision as a circuit (or in concert with a circuit), a more quantitative, data-driven approach to human augmentation is needed.

With this in mind, we present the sense-assess-augment (SAA) framework [18], which is based loosely on the adaptive system framework originally proposed by Feigh et al. [19]. It begins with the human, sensing their physical, physiological, and psychological state. **Sensing** is the most mature piece of the paradigm, thanks to considerable commercial investment in athletics, healthcare, and productivity. Sensors exist or are in development that can measure a huge range of parameters, such as brain activity, eye movement, skin temperature, and biological performance markers, such as blood glucose levels or molecules like orexin that indicate the onset of fatigue. **Assessment** involves the interpretation of data from multiple, individual sensors and merging it into actionable information. The challenge is to empirically make sense of the data in relation to individual baselines and the needs of the task at hand. Ideally, this is a task shared by both human and machine and happens both in real time and across a lifetime. Finally, based on the assessment, the appropriate **augmentation** is delivered. Augmentation can take many forms, including the redistribution of tasks from man to machine, just in time or chronic uptake of drugs, external hardware, environmental changes, or even genetic engineering. We will discuss each of these pieces in more detail below.

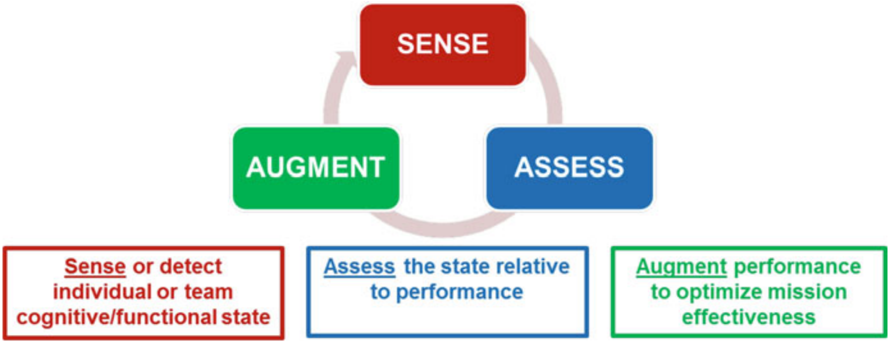

Fig. 7.1 Sense-assess-augment framework (From Galster and Johnson [18])

Each piece of the framework is critical to the design and deployment of human augmentation. Sensing without assessment is frustrating. It is, in fact, one of the most common complaints of consumers trying to make sense of the athletic, health, and productivity data they are collecting. Awash in a flood of data, many ask: what does the data mean and how do I alter my performance accordingly?

Augmentation without the sensing and assessment components is not only potentially dangerous, but breeds distrust among the public and policy-makers. For example, the Air Force pilots responsible for the friendly-fire deaths of Canadian troops in Afghanistan in 2003 implicated "go pills" as the cause of the accident. Although the official investigation found no contribution of the drug to the outcome, the public and media [20] were not persuaded. Physiological monitoring and assessment might have provided objective proof whether the cause was poor judgment by the pilot, a side effect of a widely used drug, or a combination of the two that stemmed from individual susceptibility.

Absent the framework described above, the sensing and augmentation communities have largely worked independently, and the assessment piece has lacked a research leader to make significant progress to bridge them. As we will discuss in detail, if there is one lesson learned from the decades of attempting to deliver the many promises of human performance augmentation, it is the necessity and interdependence of the three steps (Fig. 7.1).

7.2.1 Sense

In their article "Beyond Asimov: The Three Laws of Responsible Robotics [21]," Woods and Murphy proposed alternatives to Asimov's classic laws of robotics, stating "The capability to respond appropriately—responsiveness—may be more important to human-robot interaction than the capability of autonomy." An unfortunate case in point is the 2010 drone attack that killed 23 Afghan civilians. The primary cause of the accident cited by Air Force and Army officials was information overload [22]. In addition to keeping track of video from the drone, operators were also engaged in “dozens of instant-message and radio exchanges with intelligence analysts and troops on the ground.” They failed to mentally account for the children that were part of the civilian assembly.

This is but one illustration that stems from a lack of shared perception between human and machine. There is not only a need, but now the opportunity, to push beyond simple measurements of human experimental feedback, such as filling out surveys or asking people, “Was your workload diminished or not?” Despite unprecedented technological advances, our ability to assess an individual’s or team’s physical, psychological, and physiological readiness is startlingly unsophisticated. We are blind, for example, to any number of problems that plague human operators:

- When boredom or data overload lead to prolonged lapses of attention
- When emotional resilience hits its breaking point
- When exhaustion or hunger degrade cognitive abilities

The emerging field of neuroergonomics aims to remedy this by decoding the functioning of a healthy brain at work [23]. The work is highly interdisciplinary, drawing from human factors, ergonomics, neuroscience, and machine learning to develop adaptive interfaces that sense and respond to changes in an individual’s executive function, an umbrella term that refers to cognitive processes such as planning, working memory, task switching, initiative, and others. These studies are important because, as founder Parasuraman [24] explains, more traditional cognitive science and neuroscience work “often fails to capture the complexity and dynamics of behavior as it occurs naturally in everyday settings. In other cases, the tasks used in laboratory studies may have little or no relation to those confronting people in everyday life.”

Another important milestone in personal sensing came in 2007, when two editors at Wired Magazine noticed that trends in life logging, personal genomics, location tracking, and biometrics were starting to converge. Gary Wolf, one of the founders of what became known as Quantified Self, stated “These new tools were being developed for many different reasons, but all of them had something in common: they added a computational dimension to ordinary existence.” Today nearly anyone can record a half dozen physiological data streams in his quest to become fitter or healthier, including a log of alpha rhythms to diagnose sleep quality. For an elite athlete or corporate executive, the sky is the limit in terms of quantified physiological parameters. This made the development of unobtrusive, wearable, and robust sensors a commercial industry, enabling performance tracking at the individual level at a cost that would have been unfathomable just a decade ago.

The combination of neuroergonomics and individual tracking allows us to finally escape the tyranny of the “average user” which has dominated HCI philosophies. As discussed earlier, many protocols originate from Fitts and others, who examined the most complicated pairing of man and machine at that time—the airplane cockpit. The idea of an average user worked for pilots who simply had to distinguish between one knob or another on a panel and the time-accuracy trade-offs did not vary significantly across the population of users, given the right training. It was also fine for distinguishing the utility of a keyboard versus a mouse. The same cannot be said for today's information saturated, multi-tasking knowledge worker. There's huge variability in executive function between individuals, as well as differences that alter performance hour to hour, and from day to day. Thus, the complexity and the number of parameters that must now be optimized together fundamentally changes how we need to approach HCI.

Topol describes how individual tracking is already leading to massive changes in the approach to healthcare in his book *The Creative Destruction of Medicine* [25]. An example of particular relevance to HCI and the SAA model is blood glucose monitoring. Until a few years ago, the only way for diabetics to monitor their glucose levels in their day to day life was with finger sticks, using a device to lance one of the fingers to produce a drop of blood which must then be smeared onto a test-stick and read by a small device. This procedure is usually performed four times a day, is inconvenient, somewhat painful, but more importantly, still runs the risk of missing large spikes or drops due to food intake, exercise, or incorrect insulin dosages. Today, continuous glucose monitoring is possible with a sensor that samples glucose levels from the interstitial fluid just beneath the skin using a small, indwelling 27 gauge needle. The device still has its downsides, such as cost and the need to calibrate readings with finger sticks every 12 h. However, the sensor is robust enough that wearers can exercise and shower as normal. Topol describes additional sensors in development, noting "contact lenses can be embedded with particles that change color as the blood sugar rises or falls or the glucose level can be assessed through tears. Another imaginative solution has been dubbed a "digital tattoo" in which nanoparticles are injected to the blood that bind glucose, and emit a fluorescent signal that is quantified by a reader on a smart phone."

The challenge for HCI is to become equally imaginative in what to sense and how to sense it. The artificial intelligence and HCI communities have continued to focus on how the human can better access and utilize computer technology, without mention of how sensing of the human condition and capabilities might also augment the machine. For example, Sandberg writes, "What is new is the growing interest in creating intimate links between the external systems and the human user through better interaction. The software becomes less an external tool and more of a mediating 'exoself.' This can be achieved through mediation, embedding the human within an augmenting 'shell' such as wearable computers or virtual reality, or through smart environments in which objects are given extended capabilities [26]."

We now have the sensors and digital infrastructure to "remotely and continuously monitor each heart beat, moment-to-moment blood pressure readings, the rate and depth of breathing, body temperature, oxygen concentration in the blood, glucose, brain waves, activity, mood—all the things that make us tick [27]." And in response, we can imagine a machine that uses this information to asses the cognitive and affective state of its user and dynamically alter its level of automation and complexity in response. This is not a new idea—the field of human/brain-computer interface has sought such an adaptive interface since man became so dependent on his machine counterpart. But most of instruments used to examine mental workload today, such as electroencephalography (EEG), electrocorticography (ECoG), functional near-infrared spectroscopy (fNIRS), and functional magnetic resonance imagery (fMRI), were designed for laboratory use where issues of wearability, comfort, portability, and robustness are not an issue. In their review, Pickup et al. note [28], "The notion [of mental workload] has found widespread acceptance as of value in assessing the impact of new tasks, in comparing the effects of different or job interface designs and in understanding the consequences of different levels of automation." This highlights that much of the prior HCI work focused on initial design considerations rather than true adaptability.

Beyond simple user experience however, these instruments miss the more common and frequent sources of performance decrement, such as lack of sleep, low blood glucose, emotional distress, sickness, etc. Nor does it account for the growing source of information through mobile and social media. A recent survey [29] found that 75 % of workers access social media on the job from their personal mobile devices at least once a day (and 60 % access it multiple times a day). Without the ability to pinpoint the source of increased mental workload in real time, the proper augmentation strategy may not be implemented.

Biomarkers are essential to this endeavor. In addition to the readouts from EEG for example, peripheral measures largely associated with the autonomic nervous system have proven to be salient as well [30]. Biomarkers can mean different things: blood oxygen levels, eye movements, perspiration levels, posture, or any number of molecular metabolites.

Molecular monitoring has been aided significantly by the development of flexible, dissolvable electronics. Advances in electronics and microfluidics have led to the development of miniaturized "lab-on-a-chip" devices and unobtrusive wearable psychophysiologic sensors that can support the rapidly emerging need to instrument the user and monitor physical and mental states. This monitoring, when fed back into the machine system, can provide a "check engine light" for the operator as well as drive adaptive autonomy based on the real-time needs of the operator to improve overall sociotechnical team mission performance. Recent scientific studies have elucidated several molecular targets of opportunity. For example, the neuropeptide orexin A (hypocretin) has been implicated in arousal/alertness. Deficiency of orexin A results in narcolepsy, while other studies [31] suggest orexin is the central switch between sleep/wake states. Previously, monitoring this peptide in patients required a sample of cerebrospinal fluid—an impractical obstacle to widespread adoption. However, recent advances in biofunctionalized sensors have increased sensitivity for orexin detection over 3 orders of magnitude (pM levels) allowing for peptide detection in saliva—a more preferable biomatrix for sampling.

Another molecular target of opportunity is neuropeptide Y (NPY). This 36 amino acid peptide has been implicated in learning and memory and is produced by the hypothalamus. In one study [32], animals whose behavior was extremely disrupted by induced stress displayed significant down regulation of NPY in the brain, compared with animals whose behavior was minimally or partially disrupted and with unexposed controls. One-hour post-exposure treatment with NPY significantly reduced prevalence rates of extreme disruption and reduced trauma-cue freezing responses, compared with controls. Although most studies on NPY have been performed with rodents, there is accumulating data [33] from the genetic to the physiological to implicate NPY as a potential 'resilience-to-stress' factor in humans as well.

Diabetics are not the only ones who need to be concerned with blood glucose levels. Previous studies have not only shown decreased cognitive performance with low blood glucose, but that increasing blood glucose can partially compensate for decreases in procedural memory due to sleep deprivation [34], a condition that is increasingly common among workers across industries.

As mentioned, one of the biggest challenges is developing sensors that do not themselves impinge on human performance. Current "wet electrode" EEG monitoring, for example, is cumbersome enough to preclude its use except in the most extreme necessities where lapses in performance could mean loss of life (e.g. Flight traffic controllers). Arguably, the future of human performance monitoring may benefit most from advances in materials science, such as recent work [35] utilizing flexible, dissolvable, and unobtrusive electronics. Transient electronics, made of biocompatible metals and encased in silk, are meant to be implanted into the body, do their work for days, weeks, or even months, and then safely dissolve and resorb in the body.

In addition to measuring the human directly, we must also sense the environment to discover the right correlates to understand degraded executive function in context. Lighting, noise levels, and temperature can all impact cognitive function, and perhaps just as importantly, offer some the easiest of potential solutions.

7.2.2 Assess

Man's ability to understand is often outstripped by his ability to measure. Assessment of the context of the psychophysiologic and performance data represents a key underdeveloped area in many systems in need of future research. Knowledge of context and changes in context allow human team members to disambiguate underconstrained data that can have different meanings in different settings. Machine reasoning to understand sensor data related to environmental, system, task planning, and user physical and cognitive state will allow the system to share some level of perception with the human operator in the proper context. Fundamentally, assessment addresses three questions: who should augment, under what conditions, and how can we quantify the effects?

To better understand the challenge of assessment, consider the landmark work of Yerkes and Dodson [36], who in 1908 proposed a relationship between adverse reinforcement and discrimination learning in rats. What became known as the Yerkes-Dodson Law, popularized decades later in a review by Hebb [37], resembles an inverted, U-shaped curve, as shown below. The Hebbian version proposes that at low arousal, people are lethargic and perform badly. As arousal increases,

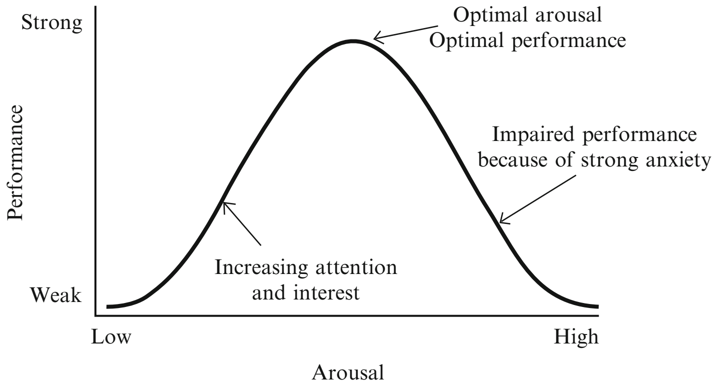

Fig. 7.2 Hebbian version of the Yerkes-Dodson law

performance also increases, but only to a point, after which increasing arousal actually decreases performance. Arousal in this context has often been equated with stress (Fig. 7.2).

Thus one might assume that, given the variety and robustness of sensors today, it should be straightforward to assess from physiological data when a user is experiencing less than optimal arousal in an operational setting and thus enhance the human or adjust the computer interface accordingly to maximize performance. However, there have been many criticisms of the Yerkes-Dodson Law, much of it relating to the misinterpretation [38] of the original work. For example, many modern references use terms such as arousal, stress, and performance, terms that were never used in the original paper and remain vague and un-quantified today. Nor was the original work, which was performed using rats, intended to extend to the relationship between stress and performance in humans. Even those experiments performed with rats produced notable exceptions to the expected curvilinear response. For example, as Easterbrook [39] describes in his paper on cue-utilization theory, “On some tasks, reduction in the range of cue utilization under high stress conditions improves performance. In these tasks, irrelevant cues are excluded and strong emotionality is motivating. In other tasks, proficiency demands the use of a wider range of cues, and strong emotionality is disorganizing. There seems to be an optimal range of cue utilization for each task.” Thus, Easterbrook goes on to explain, tasks can be considered complex if it involves attention to multiple cues and simple if it involves focused attention to a single cue. This may constitute Easterbrook’s definition of difficulty, but it is by no means widely accepted.

The problem extends throughout the human performance community, as well as medicine. In a now famous article, Ioannidis [40] suggested that much of what medical researchers conclude in their studies is misleading, exaggerated, or flat-out wrong. His conclusions are in keeping with the issues of the Yerkes-Dodson Law as well: (1) the smaller the studies, the less likely the research findings are to be true; (2) the smaller the effect, the less likely the research findings are to be true;

(3) the greater the financial and other interests, the less likely the research findings are to be true; (4) the hotter a scientific field (with more scientific teams involved), the less likely the research findings are to be true.

The appeal of the Yerkes-Dodson Law lies in its appeal to our intuition. We have all encountered cases where arousal, in the form of a cup of coffee or an impending deadline, allowed us to focus and perform better than we might have otherwise. Likewise, we have all experienced stress, in form of a cold or an overflowing email inbox, that appeared to degrade performance. What's missing from many of the studies today is the ability to determine the context of stress or arousal, and the patterns that link that context to individual performance. This is critical for determining whether augmentation is needed and the predicted improvement in performance based on the augmentation selected. Such an approach requires, at least initially, the fusion of much more sensory data from both the individual and the environment, than most research currently includes. As Tapscott and Williams warn [41], "When the devices we use to capture and process data are sparsely distributed and intermittently connected, we get an incomplete, and often outdated, snapshot of the real world."

The most common approach to pattern recognition is based on models, such as Markov models or neural networks, which provide some general knowledge of the system they are observing. However, both these approaches require large sets of training data in order to produce accurate results. For example, a study might monitor EEG channels combined with heart rate data as a participant is put through scenarios that are believed to represent low and high mental workload tasks. The training data establishes classification criteria for the two states. As the individual is then tested with real tasks, one sees attribution of low and high mental workload, typically with accuracy in the 80–85 % range. An issue manifests when baselining takes so long that the test subject enters a high stress, disengaged mental state before the experimental portion even begins. Thus, data which are supposed to indicate stress are already one or two standard deviations above baseline and thus little variation is seen in assessed response.

All of this data, however, is taken in a laboratory setting with very controlled parameters and tasks. If the characteristics of the data being analyzed deviate significantly from the training model, then previously learned data sets must be relearned along with the new data set. This means that without retraining, a model that relies on select EEG channels to produce impressive accuracy rates for a vigilance task, for example, often does not work well when applied to a different task in a different setting. This becomes even more problematic when you consider that the average worker engages in several tasks as part of his work, each of which may rely on a distinct assessment or augmentation. One task may require intense vigilance while another may require a mix of creativity, abstract thinking, and the ability to forecast. Today, a study that focuses on assessment of vigilance will be of little consequence to an assessment of creativity.

The answer then is not to collect less data, at least initially. The goal is to collect as much data as possible to discover the relevant performance patterns for each individual. This will likely require a data-driven algorithm that requires no a priori knowledge of the underlying system and can operate without a closed data set. The algorithm would be capable of learning and would include some or all of the following features:

1. Would not need to be tuned based on expected features in a data stream.
2. Able to learn and recognize patterns in an unlabeled data stream.
3. Works online—the process of learning and recognition would occur simultaneously without an offline training phase.
4. Quickly converges to recognize data patterns after only a few occurrences.
5. Finds patterns in nonlinear, nondeterministic, and non-Markovian systems.
6. Interpretable structure and produces an interpretable model.
7. Hierarchical pattern detection for the determination of context.

Such a system is particularly important when trying to merge sensor data across multiple time scales. Many parameters can be measured on an hourly or daily basis, but the trends that indicate the source of aberrations may not be apparent for months. For example, Sky Christopherson, a former Olympic cyclist turned technology CEO, started having health problems despite a lifetime of fitness and healthy eating. Because of his familiarly with personal tracking as an elite athlete, he started collecting a range of biomarkers and environmental data to discover where he could make significant, positive contributions to his health, including sleep, diet, stress, exercise, and traditional physiological measures such blood pressure. One of the most significant causes of stress was related to sleep quality, and only after collecting data for nearly 9 months did he notice trends that varied with the season. His assessment was that the real issue was room temperature that varied with outdoor temperature, so he installed a water-filled cushion on his bed that actively regulated body temperature year round. Although he implemented other changes, the effects were profound. He not only reversed his health issues, but in the process set a world cycling record at the age of 35—a feat previously thought biologically impossible due to declining testosterone levels [42].

It would also be desirable to add a predictive function to the learning algorithm. The simplest method for predicting the next state is based on the probability calculated by the number of times each state has been outputted from the system. Lacking a model for the underlying system itself, this approach might in fact be the only reasonable method of prediction. Future activities include enabling hierarchical and orthogonal learners to detect patterns of patterns, detecting spatial patterns within the model, determining similarity measurements between patterns, and incorporating visualizations of the model to assist human decision-makers in the post-processing step to identify meaningful and more nuanced patterns.

Performance assessments would ideally be quantified relative to an individual baseline collected over time. As we saw with the confusion around the Yerkes-Dodson Law, to say simply classify someone as tired or stressed provides little correlation to performance. But if it were possible to know, for example, when a someone's critical thinking ability decreased by 25 %, a better decision as to how and when to address the symptom of fatigue could be made. Nor should these assessments occur only at the tactical level. Supervisors and leaders are just as likely, if not more so, to be suffering from lack of sleep and exercise, poor nutrition, and information overload that can impair decision-making.

Initially, such a system might sound too complex to be manageable, much less design. However, Kurzweil's view of complexity in his book, *How to Create a Mind*, is undoubtedly relevant. He points out:

> We might ask, is a forest complex? The answer depends on the perspective you choose to take. You could note that there are many thousands of trees in the forest and that each one is different. You could then go on to note that each tree has thousands of branches and that each branch is completely different. Then you could proceed to describe the convoluted vagaries of a single branch. Your conclusion might be that the forest has a complexity beyond our wildest imagination. But such an approach would literally be a failure to see the forest for the trees. Certainly there is a great deal of fractal variation among trees and branches, but to correctly understand the principles of a forest you would do better to start by identifying the distinct patterns of redundancy with stochastic (that is, random) variation that are found there. It would be fair to say that the concept of a forest is simpler than the concept of a tree. [43]

Of course, there are challenges to such an approach as well. As the volume of raw data from various sensors increases, the problem of finding underlying sources within the information becomes more difficult and time consuming. Increasing the spatial resolution increases the number of data channels. Increasing the temporal resolution increases the sampling rate. Such a system would likely require long periods of data collection and analysis, along with input directly from the user, before it was capable of reliably recommending appropriate augmentation strategies.

This suggests that assessment not only needs to happen outside the laboratory, but likely outside the workplace as well. Just as "digital natives" expect to be tethered to their computing devices, those who grow up, literally, with assessment tools will find it just as normal to incorporate sensors both on and off the job to enhance performance. The growing Quantified Self movement indicates this change is already underway. This means that although assessments based primarily on self-learning algorithms will take longer to refine, the end result should be operationally robust for users who begin using them long before entering the workforce. A combination of model-based and self-learning algorithms may make sense in the interim.

So far, we've looked at assessment of a user's physical and cognitive states as they perform tasks, but assessment must also predict the best augmentation strategy and timing, as well as an evaluation of its performance enhancement. Determining returns on investment for augmentation might initially come from population studies of augmentation methods, which are then refined over time as a user implements them. But as Topol describes, there are radical variations in terms of effectiveness, even with vigorously tested substances such as commercial drugs. Part of the problem is a tendency to treat the signal, not the underlying cause (if it's even known). As an example, he discusses the cost-benefit analysis of prescribing statin drugs like Lipitor which lower blood cholesterol and which therefore presumably prevent heart disease [44], "So almost all patients will have a great blood test result with Lipitor.

But only 1 out of 100 without prior heart disease but at risk for developing such a condition will actually benefit. It therefore seems that the predominant benefit is cosmetic, normalizing an out-of-range blood test, at the risk of engendering side effects." It's not that Lipitor, the most widely prescribed drug in the world, isn't effective at lowering blood cholesterol. It's that lowering blood cholesterol doesn't lower the risk of heart disease in most of the population, despite the fact that the correlation exists precisely because of large population studies. Since nearly all measures that we can conceive of today are surrogate measures, the role of assessment is not only to recommend augmentation, but to determine its efficacy.

Perhaps one of the biggest returns on investment for assessment comes from harnessing the power of feedback loops. Some of the most intractable health problems—obesity, diabetes, smoking—have shown progress with biofeedback, with improvement outcomes typically in the range of 10 % [45]. In real terms, that means an obese 40-year-old man would spare himself 3 years of hypertension and nearly 2 years of diabetes by losing 10 % of his weight. Reducing traffic speeds by 10 % from 40 to 35 mph would cut fatal injuries by about half.

What has prevented biofeedback from becoming a mainstay of HCI or other systems is the ability to effortlessly collect and track personalized data. Thus, the assessment tools we have been discussing can now deliver information not in the raw-data form in which it was captured, but in a context that makes it emotionally resonant because it quantifies the consequence of not changing. This is why assessment cannot be the domain of the machine alone, and the currency must be information, not data. The goal is to have shared perception and to make joint decisions about augmentation strategies and timing.

7.2.3 Augment

The rising role of technology in our lives has led to a deep dependency, but not always a harmonious one. From "crackberry" addictions to "digital sabbaticals", there is a kind of begrudging "you can't live with it and you can't live without it" attitude that pervades the relationship between man and machine. The Air Force's Technology Horizons report warns, "Although humans today remain more capable than machines for many tasks, natural human capacities are becoming increasingly mismatched to the enormous data volumes, processing capabilities, and decision speeds that technologies offer or demand; closer human-machine coupling and augmentation of human performance will become possible and essential [46]."

One of the greatest difficulties in developing more adaptive human-computer interactions is that they must accommodate a broad set of tasks, environments, and users. Unlike the pilots in Fitts' experiments, computer users today vary considerably in their cognitive and physical capabilities. This is why future adaptations and augmentations must happen at the individual level, though this is not without its own set of challenges. We define augmentation in a functional sense as enhancing human performance above baseline levels and/or boosting performance to baseline levels after a decrement. Although augmentation strategies may include modifications to the operator, the human-machine interface, or the machine itself, it is important to understand that the desired end result is always focused on improved task outcomes or human performance.

The options for augmentation, once an assessment has determined it is necessary, are extensive. Augmentation may come from increased use of autonomous systems, interfaces for more intuitive and close coupling of humans and automated systems, and direct augmentation of humans via drugs or implants to improve memory, alertness, cognition, or visual/aural acuity. Important considerations when choosing among those options are: (1) machine versus human enhancement, (2) task specific versus lifetime enhancement of the user, and (3) potential trade-off considerations.

(1) Machine versus human augmentation: Most augmentation strategies being considered or deployed today focus on the man-machine interface. For example, in high workload or stressful situations, tasks may move from manual to autonomous control, less critical information may be removed from the display to help the user focus, or the employment of multi-sensory signals to avoid visual overload.

Virtual partners or assistants are another option to machine augmentation. An early example was Microsoft's Clippy, the paperclip that offered advice and access to help topics if you appeared confused. The problem was that Clippy had very limited assessment or interaction with the user and frequently became an annoyance. Advances in the ability to understand natural language and assessment technologies make modern versions more appealing. For example, IBM's supercomputer named Watson, famous for beating the best human contestants at the game Jeopardy!, is now being used as an important diagnostician and consultant to physicians. Watson can help manage the flood of data coming into a patient's electronic record while simultaneously scanning for recent and relevant publications in the literature that might alert a doctor to new therapies or trends. According to Sloan-Kettering, one of the first hospitals to use Watson, it would take at least 160 h of reading a week just to keep up with new medical knowledge as it's published, let alone consider its relevance or apply it practically. In tests, Watson's successful diagnosis rate for lung cancer is 90 %, compared to 50 % for human doctors [47]. It is not just the access to medical literature that's important, but also the computer's lack of cognitive bias in assessing it. It's estimated that one-third of hospital errors are due to misdiagnosis, one of which is anchoring bias, the human tendency to over-rely on the first pieces of information offered when making decisions [48].

On the human side, pharmaceuticals can be used to repair decrements in cognitive function, such as from sleep deprivation or prolonged vigilance, or increase natural capacities. For example, Modafinil and similar alertness or vigilance support pharmaceuticals have been studied extensively by the Army to support aviation operations. Although Modafinil has demonstrated utility in several DoD operational contexts, research now implies that may also offer benefit to sleep-deprived senior decision-makers, with research showing improved planning among their test subjects [49]. More recently, a study of sleep-deprived physicians found the drug improved their cognitive flexibility while reducing impulsive behavior [50].

Interestingly, Modafinil has also been shown to enhance working memory, especially at harder task difficulties for lower-performing subjects [51]. The mode of action is not yet understood, but part of what seems to happen is that Modafinil enhances adaptive response inhibition, making the subjects evaluate a problem more thoroughly before responding, thereby improving performance accuracy [52].

There are many challenges to the use of pharmaceuticals for performance enhancement. Many, of course, have undesirable side effects. More importantly, the current system for the development and approval of pharmaceuticals is geared towards treatment of disease, not augmentation of otherwise healthy individuals. This means that nearly all drugs for enhancement purposes are being prescribed by doctors for "off label" usage. Since drug companies can't market non-therapeutic benefits, tests are rarely run to prove effectiveness and safety of the product in individuals. In order for pharmaceuticals to play a more large-scale role in HCI strategies, current Federal Drug Administration regulatory schemes would need to be reformed. Fortunately, there are many augmentation alternatives outside of pharmaceuticals.

(2) Task specific versus lifetime augmentation: Task specific augmentation often focuses on adaptive interfaces that modulate the speed, amount, and visualization of information. On the human side, techniques such as noninvasive brain stimulation have shown to be effective at improving performance across a wide variety of tasks, presumably by either raising neuronal membrane and force action potential in the case of transcranial magnetic stimulation or altering neuron excitability in a region in the case of transcranial direct current stimulation [53].

Long term augmentation strategies are increasingly being investigated, primarily changes in diet and nutrition. The quantity and quality of dietary choices and distribution of nutrients throughout the day greatly impact muscle performance, body composition, cognitive performance, and feelings of energy or exhaustion. In addition, there is a rapidly expanding body of research showing an ever-increasing linkage between commensal (native gut) bacterial and overall human phenotype, from increased obesity to cognitive metabolites and immune responses [54].

Not surprisingly, effective metabolism is key to both physical and cognitive performance. A wide range of non-pharmaceutical substances are currently under investigation for their ability to enhance cognitive performance through regular ingestion over the long term. Among these are several where the mechanism of action is largely metabolic, such as leucine, creatine, and coenzyme Q10. Recently researchers at MIT and elsewhere have shown substantial cognitive benefits by the use of a newly developed magnesium compound, magnesium-L-threonate (MgT). In animal studies it has been demonstrated that increasing brain magnesium enhances synaptic plasticity in the hippocampus and leads to elevated learning abilities, working memory, and short- and long-term memory [55].

Another promising metabolic target is ketone bodies. Under normal conditions, the brain is totally dependent upon the metabolism of glucose to supply its metabolic energy. However, during starvation, the body normally produces ketone bodies that can then supply the majority of brain energy needs. Not only can ketone bodies replace glucose as the major energy substrate for the brain, the metabolism of ketones instead of glucose, increases the energy contained in the major cellular energy transmitter adenosine triphosphate (ATP [56]). The increased metabolic energy contained in ketone bodies (as opposed to glucose) was recently exploited in a DARPA-funded project (with NIH and Oxford) where a ketone ester was developed that can be administered as a food and which can improve cognitive and physical performance in animals and improve physical performance in humans. An experiment conducted at Oxford University with 22 elite British rowers showed remarkable improvements in performance including 10 season's best, 6 personal best, and 1 world's record an hour after ingestion of the ketone ester. Additionally, in animal studies, substantial improvements in cognitive performance have been seen and human studies are currently underway (personal communication by K. Clarke during DARPA presentation of research results, July 2012).

Another potential target for long term enhancement resides in the microbiome. The enteric nervous system, a collection of neurons in the gut often called "The Second Brain" in the popular press, contains some 100 million neurons, more than either the spinal cord or peripheral nervous system. Approximately 90 % of the nerve fibers in the primary visceral nerve (the vagus) carry information from gut to brain, not the other way around. Recent evidence supports the view that triggers and signals from the gut affect our emotions, decision-making, response to stress, immune response, and learning and memory.

The enteric microbiome of non-human organisms living in the human gut is thought to impact cognitive performance and emotional resilience. A recent study showed that mice receiving *Lactobacillus rhamnosus* were less anxious, performed better on tests for learning and memory, and had lower cortisol levels after stressful situations [57]. Supplemented mice also had atypical mRNA levels for 2 GABA receptors involved in decision-making and learning/memory. Severing the vagus nerve attenuated the effect, implying that commensal bacteria in the gut can have a direct effect on neurotransmitter receptors in the central nervous system in normal, healthy animals.

Nor is the effect limited to animal studies. Oral ingestion of probiotics, which often include *L. rhamnosus*, have also been effective in reducing psychological distress in otherwise normal, healthy humans, while antibiotic use may disturb the microbiome flora population distribution for a poorly understood length of time. In a clinical trial, volunteers participated in a double-blind, placebo-controlled, randomized group study with probiotics administered for 30 days and assessed with the Hopkins Symptom Checklist (HSCL-90), the Hospital Anxiety and Depression Scale (HADS), the Perceived Stress Scale, the Coping Checklist (CCL) and 24 h urinary free cortisol (UFC). Results indicate probiotic administration significantly alleviated psychological distress in volunteers, particularly as measured by the HSCL-90 scale [58].

(3) Potential trade-offs: Augmentation is unlikely to serve as a one-size-fits-all solution, and it's clear that trade-offs may be a consideration with certain cognitive enhancement strategies. For example, genetically engineered mice with extra copies of the NR2B gene have improved memories and learn faster, but they are also more susceptible to addiction and feel pain longer than normal mice [59]. In a documented case of a human with naturally and profoundly enhanced memory, the man was able to remember vast amounts of text on a single reading, even in a unfamiliar, foreign language, but was almost entirely unable to grasp metaphors, as his mind was so fixated on particulars [60]. This anecdotal evidence is in line with computer models that show that memory is actually optimized by slight imperfections, as they allow one to see connections between different but related events [61].

The New York Times also ran a series of essays by students who were otherwise considered healthy but began illicitly taking prescription medications such as Adderall and Ritalin to maintain their edge at competitive schools. These drugs are psychostimulants intended for treatment of attention deficit hyperactivity disorder (ADHD) and narcolepsy, but the essays reveal that those without the condition find benefits too. Students reported the drug provided an almost tunnel-vision like focus, reducing fatigue while reportedly [62] increasing reading comprehension, interest, and memory. But many also found the drugs deliver some unpleasant side effects, especially as they are often abused, such as anxiety, depression, sluggishness, and social withdrawal.

The difficulty is that in most cases, the neural mechanisms underlying cognitive enhancement, particularly in healthy individuals, is poorly understood. With psychostimulants, for example, it was only in 2012 that low-level, cognition-enhancing doses were shown to exert regionally-restricted actions, elevating extracellular catecholamine levels and enhancing neuronal signal processing preferentially within the prefrontal cortex [63]. Little data is available on how effective these neuro-enhancing agents are for non-ADHD users or on long-term side effects. This mismatch between the demand for cognitive enhancers and funded studies on their use may stem from a historical tendency to regulate rather than educate when it comes to human enhancement. The ability to understand the genetic basis for the varying effectiveness of augmentations is improving with technologies such as genome-wide association studies, which have been used to predict drug response based on individual genetic variations [64].

One potential solution to mitigating these trade-offs is to think about augmentation cocktails that make smaller changes across a spectrum of abilities, with the hope the cumulative effect is greater than the sum of its parts. As neuroscience unravels how the brain generates behaviors and integrates multiple kinds of information, such as memories, sensory information, and decisions, more targeted augmentation strategies can be developed. Optogenetics [65], a technique that allows researchers to selectively express or silence neurons in a temporally precise fashion using pulses of light, offers a bit of both. Not only is the technique being used to systematically explore how neural circuits contribute to cognition and behavior, but it has also been used in live animal studies to directly control behaviors such as reward seeking [66]. Nearly all of these options, however, are relatively long-term before they can be integrated into commercial workspaces. In the meantime, we must be aware that trade-offs may exist and they may not be known prior to use. In this regard, the assessment part of the paradigm becomes even more important.

7.3 Future Perspectives

Aristotle once said, “Those who know, do. Those who understand, teach.” Had Aristotle lived during the computer age, he might have concluded his final statement with the word “simulate” instead. At the Dartmouth Conference of 1956, where artificial intelligence (AI) was essentially born, it was proposed [67] “Every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it.” Indeed, the driving assumption of artificial intelligence is that when we have created a machine that can pass the Turing Test with flying colors, we will have laid bare the underpinnings that make the human mind possible. And while significant progress has been made in the creation of machines capable of performing like humans and sometimes better, it is fair to say that a coherent theory of the mind, particularly one that links the underlying biology with behavior, has lagged far behind.

Thus when Newell, one of the founding fathers of AI, grew frustrated with the progress of cognitive science in 1980 (and again in 1990), his proposed solution was shaped by the governing principles of AI. Anderson and Lebiere [68] describe Newell’s thoughts on the matter as thus, “He would point to such things as the ‘schools’ of thought, the changes in fashion, the dominance of controversies, and the cyclical nature of theories. One of the problems he saw was that the field had become too focused on specific issues and had lost sight of the big picture needed to understand the human mind.” In response, Newell made two contributions that would ultimately drive the field of computational cognition. First, he proposed a set of functional criteria for the evaluation of cognitive theories that forced researchers to break out of what Newell feared was a kind of theoretical myopia, where models explained their own results but made little sense in the greater context of what was already understood or observed. One can debate whether Newell’s criteria are the right ones (and many have), but the result has been the development of increasingly sophisticated cognitive architectures that combine psychological, and more recently neuroscientific [69], knowledge with developments from artificial intelligence to produce a powerful general-purpose engine of cognition.

Newell’s second major contribution was his four bands of cognition—the biological, the cognitive, the rational, and the social—that delineate the outcomes of cognition based on the speed at which they typically occur [70]. Each successive band captures the human experience at roughly 3 orders of magnitude greater than the previous (see Table 7.1). Newell thought the cognitive band was most relevant to a theory of cognitive architecture. Given the immaturity of neuroscience or even biology at the time, Newell can hardly be blamed for focusing on the more observable (and programmable) aspects of cognition, but it has had important consequences, namely the instantiation of symbolism as the primary basis of cognitive architectures. In 1980, he stated “symbolic behavior (and essentially rational behavior) becomes relatively independent of the underlying technology. Applied to the human organism, this produces a physical basis for the apparent irrelevance of the neural level to intelligent behavior [71].” And although more

Table 7.1 Newell's time scales of human action

Scale (s)	Time units	System	World (theory)
10^7	Months		Social band
10^6	Weeks		
10^5	Days		
10^4	Hours	Task	Rational band
10^3	10 min		
10^2	Minutes		
10^1	10 s	Unit task	Cognitive band
10^0	1 s	Operations	
10^{-1}	100 ms	Deliberate act	
10^{-2}	10 ms	Neural circuit	Biological band
10^{-3}	1 ms	Neuron	
10^{-4}	100 μs	Organelle	

recent architectures have penetrated into Newell's biological band, as Anderson explains, "the approach in cognitive psychology has largely been not to actually model the biological processes but rather to describe them at some level of abstraction. This level is called the subsymbolic level [72]." To use an analogy to computer systems, the assumption has been that to understand intelligent thought, we need to focus on the software, not the hardware that performs the calculations.

The comparison between biological and information systems turns out to be quite useful for examining concepts such as constraints, tradeoffs, and layered architectures. However, it must be noted that the delineation between hardware and software in computer systems is quite obvious, whereas in biological systems (including the brain), chemistry is involved at every level. In either case, Doyle and Csete effectively argue that "robust yet fragile" control systems are the key to understanding such complex systems [73]. They point out, "This enormous, hidden, cryptic complexity, driven by robustness, is both the greatest initial obstacle in using advanced information and control technologies as metaphors for biology and also ultimately, the key to important insights and theories." Thus, in order to fully capitalize on complexity, one must develop (1) a thorough understanding of the protocols that drive interaction between layers and modules, and (2) an understanding of robustness trade-offs which drive complexity.

This first point calls into question whether Newell's bands of the cognition, based on temporal duration, are the right hierarchy for architecture evaluation. The fact that existing architectures such as ACT-R and SOAR work as well as they do suggests the symbolic layer of cognition was likely the right starting point. Although researchers are successfully transcending multiple timescales as they model more complex tasks [74], this does not mean that a time-based hierarchy is the most productive one. For example, Sun 75et al. [75] have proposed four layers—physiological, componential, psychological, and social/cultural. At first glance, these appear remarkably similar to Newell's, but in this case are based on phenomena that occur rather than their duration (Table 7.2).

Table 7.2 Hierarchy of levels from [75]

Level	Object of analysis	Type of analysis	Computational model
1	Inter-agent processes	Social/cultural	Collections of agents
2	Agents	Psychological	Individual agents
3	Intra-agent processes	Componential	Modular construction of agents
4	substrates	physiological	Biological realization of modules

What is important is that these levels interact with and constrain one another in such a way that they cannot be studied solely in isolation. We do not presume that Sun's or anyone else's hierarchical layers necessarily offer the correct level of detail, however we agree with Sun et al. that "The capability, at least in principle, to map collective phenomenological properties all the way down to neural properties (or other detailed level descriptions) is an essential aspect of an effective theory of cognition in a sociocultural context. In contrast, the ability of a high-level theory to accurately model high-level phenomena is a necessary but not sufficient condition for effectiveness." It is not clear how an architecture based on Newell's time-based bands can achieve this. Moreover, if Doyle and Csete are correct that biology is optimized for robustness instead of efficiency, this suggests that a focus on the protocols that connect layers and provide "constraints that de-constrain" may be a better approach. Alderson and Doyle clarify, "In protocol-based architecture (PBA), the protocols (rules of interaction that persist) are more fundamental than the modules (which obey protocols and can change and diversify). PBAs facilitate coherent and global adaptation to variations in both components and the environments on a vast range of time scales despite implementation mechanisms that are largely decentralized and asynchronous [76]."

This point is particularly important for human performance augmentation, since it suggests hard theoretical limits to optimization. We think the SAA framework presented in this chapter will provide essential information for exploring and perhaps elucidating those limits. We argue this idea must be at the heart of good HCI, since we desire to understand not just how an individual human thinks, but how well they perform at any particular point in time. Further, we need cognitive models that allow us to simulate the effect of various augmentation strategies before implementing them. Alderson and Doyle conclude that managing or perhaps preventing the "robust yet fragile" complexity spiral is a key challenge, stating "Indeed, the emergence of complexity can often be seen as a spiral of new challenges and opportunities that organisms and/or technologies exploit, but which also lead to new fragilities, often from novel perturbations." We propose it is impossible to fully address this issue for the purposes of human performance augmentation without simulating more fully the biological body in which the mind resides, particularly the neurotransmitters and metabolites that play a key role in regulating cognition. We think the SAA framework combined with new mathematics and applications of control theory offer much hope in this regard.

References

1. Pocock C (2010) USAF Chief outlines UAVs' manpower needs. AINonline, February
2. http://www.nih.gov/science/brain/. Accessed 31 Jun 2013
3. Hinze-Hoare V (2007) The review and analysis of human computer interaction (HCI) principles. arXiv preprint arXiv:0707.3638
4. Woods DD (1985) Cognitive technologies: the design of joint human-machine cognitive systems. AI Mag 6(4):86–92
5. Johnson M, Bradshaw JM, Feltovich PJ, Jonker CM, van Riemsdijk B, Sierhuis M (2010) The fundamental principle of coactive design: interdependence must shape autonomy. In: De Vos M, Fornara N, Pitt JV, Vouros G (eds) Proceedings of the 6th international conference on coordination, organizations, institutions, and norms in agent systems (COIN@AAMAS'10). Springer, Berlin/Heidelberg, pp 172–191
6. Shannon CE (1948) A mathematical theory of communication. Bell Syst Tech J 27:379–423, 623–656
7. Tanner WP, Swets JA (1954) A decision-making theory of visual detection. Psychol Rev 61:401–409
8. Miller GA (2003) The cognitive revolution: a historical perspective. Trends Cogn Sci 7(3):141–144
9. Proctor RW, Vu KL (2006) The cognitive revolution at age 50: has the promise of the human information-processing approach been fulfilled? Int J Hum-Comput Interact 21(3):253–284
10. Fitts PM (1954) The information capacity of the human motor system in controlling the amplitude of movement. J Exp Psychol 47(6):381–391
11. Papaxanthis C, Schieppati M, Gentili R, Pozzo T (2002) Imagined and actual arm movements have similar durations when performed under different conditions of direction and mass. Exp Brain Res 143:447–452
12. Seow SC (2005) Information theoretic models of HCI: a comparison of the Hick-Hyman law and Fitts' law. Hum-Comput Interact 20:315–352
13. Sanders MS, McCormick EJ (1987) Human factors in engineering and design. McGraw-Hill, New York
14. MacKenzie IS (1992) Fitts' law as a research and design tool in human–computer interaction. Hum–Comput Interact 7:91–139
15. Watson JB (1913) Psychology as the behaviorist views it. Psychol Rev 20(2):158–177
16. Kennedy D (2005) Neuroimaging: revolutionary research tool or a post-modern phrenology? J Bioeth 5(2):19
17. Proctor RW, Vu KL (2010) Cumulative knowledge and progress in human factors. Annu Rev Psychol 61:623–651
18. Galster SM, Johnson EM (2013) Sense-assess-augment: a taxonomy for human effectiveness. Technical report. United States Air Force Research Laboratory, Wright-Patterson Air Force Base
19. Feigh KM, Dorneich MC, Hayes CC (2012) Toward a characterization of adaptive systems: a framework for researchers and system designers. Hum Factors 54(6):1008–1024
20. Borin E (2013) The U.S. Military needs its speed. WIRED (online). Available at: http://www.wired.com/medtech/health/news/2003/02/57434. Accessed 24 May 2013
21. Murphy R, Woods DD (2009) Beyond Asimov: the three laws of responsible robotics. IEEE Intell Syst 24(4):14–20
22. Shanker T, Richtel M (2011) In new military, data overload can be deadly. NY Times, 16 Jan 2011 (Online)
23. Parasuraman R, Rizzo M (2008) Neuroergonomics: the brain at work. Oxford University Press Inc, New York
24. Parasuraman R, Christensen J, Grafton S (2012) Neuroergonomics: the brain in action and at work. NeuroImage 59:1–3

25. Topol E (2011) The creative destruction of medicine: how the digital revolution will create better health care. Basic Books, New York
26. Savulescu J, ter Meulen R, Kahane G (2011) Enhancing human capacities (Kindle locations 2753–2755), Kindle edition. Wiley Publishing, New York
27. Topol E (2011) The creative destruction of medicine: how the digital revolution will create better health care (Kindle locations 199–201). Basic Books, New York
28. Pickup L, Wilson JR, Sharpies S, Norris B, Clarke T, Young MS (2005) Fundamental examination of mental workload in the rail industry. Theor Issues Ergon Sci 6:463–482
29. (2012) Social media & workplace collaboration. (e-book) SilkRoad, p 3. Available through: http://pages.silkroad.com/rs/silkroad/images/Social-Media-Workplace-Collaboration-SilkRoad-TalentTalk-Report.pdf
30. Estepp JR, Christensen JC (2011) Physiological cognitive state assessment: applications for designing effective human-machine systems. In: EMBC, 2011 annual international conference of the IEEE. Engineering in Medicine and Biology Society, pp 6538–6541
31. Tsujino N, Sakurai T (2009) Orexin/hypocretin: a neuropeptide at the interface of sleep, energy homeostasis, and reward system. Pharmacol Rev 61(2):162–176
32. Cohen H, Liu T, Kozlovsky N, Zohar J, Mathe' AA (2012) The neuropeptide Y (NPY)-ergic system is associated with behavioral resilience to stress exposure in an animal model of post-traumatic stress disorder. Neuropsychopharmacology 37(2):350–363
33. Sah R, Geracioti TD (2012) Neuropeptide Y and posttraumatic stress disorder. Mol Psychiatry 18(6):646–655
34. Herzog N, Friedrich A, Fujita N, Gais S, Jauch-Chara K et al (2012) Effects of daytime food intake on memory consolidation during sleep or sleep deprivation. PLoS One 7(6):e40298
35. Hwang SW, Tao H, Kim DH, Cheng H, Song JK et al (2012) A physically transient form of silicon electronics. Science 337:1630–1644
36. Yerkes RM, Dodson JD (1908) The relation of strength of stimulus to rapidity of habit formation. J Comp Neurol Psychol 18(5):459–482
37. Hebb DO (1955) Drives and the C.N.S. (conceptual nervous system). Psychol Rev 62(4):243–254
38. Hancock PA, Ganey HC (2003) From the inverted-U to the extended-U: the evolution of a law of psychology. J Hum Perform Extreme Environ 7(1):3
39. Easterbrook J (1959) The effect of emotion on cue utilization and the organization of behavior. Psychol Rev 66(3):183
40. Ioannidis JPA (2005) Why most published research findings are false. PLoS Med 2(8):e124
41. Tapscott D, Williams AD (2010) Macrowikinomics: new solutions for a connected planet. Penguin, New York, US
42. Quantified Self (2012) Sky Christopherson on the quantified athlete (Video online). Available at: http://quantifiedself.com/2012/05/sky-christopherson-on-the-quantified-athlete/
43. Kurzweil R (2012) How to create a mind: the secret of human thought revealed, Kindle edition. Penguin Group, New York, US
44. Topol E (2011) The creative destruction of medicine: how the digital revolution will create better health care, Kindle edition. Basic Books, New York, p 22
45. Goetz T (2011) Harnessing the power of feedback loops. Wired Mag 19(7)
46. Dahm WJA (2010) Technology horizons: a vision for air force science and technology during 2010–2030, vol 1. USAF HQ, Arlington, VA
47. Steadman I (2013) IBM's Watson is better at diagnosing cancer than human doctors. Wired UK (online). Available at: http://www.wired.co.uk/news/archive/2013-02/11/ibm-watson-medical-doctor. Accessed 1 May 2013
48. Cohn J (2013) The robot will see you now (online). Available at: http://www.theatlantic.com/magazine/archive/2013/03/the-robot-will-see-you-now/309216/. Accessed 1 May 2013
49. Turner DC, Robbins TW, Clark L, Aron AR, Dowson J, Sahakian BJ (2003) Cognitive enhancing effects of modafinil in healthy volunteers. Psychopharmacology 165(3):260–269
50. Sugden C, Housden CR, Aggarwal R, Sahakian BJ, Darzi A (2012) Effect of pharmacological enhancement on the cognitive and clinical psychomotor performance of sleep- deprived doctors: a randomized controlled trial. Ann Surg 255(2):222–227

51. Muller U, Steffenhagen N, Regenthal R, Bublak P (2004) Effects of modafinil on working memory processes in humans. Psychopharmacology 177(1–2):161–169
52. Savulescu J, ter Meulen R, Kahane G (2011) Enhancing human capacities (Kindle locations 2660–2664), Kindle edition. Wiley Publishing, New York
53. McKinley RA, Bridges N, Walters CM, Nelson J (2012) Modulating the brain at work using noninvasive transcranial stimulation. Neuroimage 59(1):129–137
54. Burnet PW (2012) Gut bacteria and brain function: the challenges of a growing field. Proc Natl Acad Sci 109(4):E175
55. Slutsky I, Abumaria N, Wu LJ, Huang C, Zhang L et al (2010) Enhancement of learning and memory by elevating brain magnesium. Neuron 65(2):165–177
56. Sato K, Kashiwaya Y, Kean CA, Tsuchiya N, King MT et al (1995) Insulin, ketone bodies, and mitochondrial energy transduction. FASEB J 9(8):651–658
57. Bravo JA, Forsythe P, Chew MV, Escaravage E et al (2011) Ingestion of Lactobacillus strain regulates emotional behavior and central GABA receptor expression in a mouse via the vagus nerve. Proc Natl Acad Sci 108(38):16050–16055
58. Messaoudi M, Lalonde R, Violle N, Javelot H, Desor D et al (2011) Assessment of psychotropic-like properties of a probiotic formulation (Lactobacillus helveticus R0052 and Bifidobacterium longum R0175) in rats and human subjects. Brit J Nutr 105(5):755–764
59. Tang Y-P, Shimizu E, Dube GR, Rampon C, Kerchner GA et al (1999) Genetic enhancement of learning and memory in mice. Nature 401(6748):63–69
60. Lehrer J (2009) Neuroscience: small, furry … and smart. Nature 461:862–864
61. McClelland JL (1995) Constructive memory and memory distortions: a parallel-distributed processing approach. In: Schacter D (ed) Memory distortion: how minds, brains, and societies reconstruct the past. Harvard University Press, Cambridge, pp 69–90
62. DeSantis AD, Webb EM, Noar SM (2008) Illicit use of prescription ADHD medications on a college campus: a multimethodological approach. J Am Coll Heal 57(3):315–324
63. Berridge CW, Devilbiss DM (2011) Psychostimulants as cognitive enhancers: the prefrontal cortex, catecholamines, and attention-deficit/hyperactivity disorder. Biol Psychiatry 69(12):e101–e111
64. Topol E (2011) The creative destruction of medicine: how the digital revolution will create better health care, Kindle edition. Basic Books, New York, p 88
65. Bernstein JG, Boyden ES (2011) Optogenetic tools for analyzing the neural circuits of behavior. Trends Cogn Sci 15(12):592–600
66. Stuber GD, Sparta DR, Stamatakis AM, van Leeuwen WA, Hardjoprajitno JE et al (2011) Excitatory transmission from the amygdala to nucleus accumbens facilitates reward seeking. Nature 475:377–380
67. Halpin H (2004) The semantic web: the origins of artificial intelligence redux. In: Third international workshop on the history and philosophy of logic, mathematics, and computation (HPLMC-04 2005)
68. Anderson JR, Lebiere C (2003) The Newell test for a theory of cognition. Behav Brain Sci 26:587–640
69. Anderson JR, Fincham JM, Qin Y, Stocco A (2008) A central circuit of the mind. Trends Cogn Sci 12(4):136–143
70. Newell A (1994) Unified theories of cognition, vol 187. Harvard University Press, Cambridge
71. Newell A (1980) Physical symbol systems. Cogn Sci 4:135–183
72. Anderson JR (2002) Spanning seven orders of magnitude: a challenge for cognitive modeling. Cogn Sci 26(1):85–112
73. Doyle JC, Csete M (2011) Architecture, constraints, and behavior. PNAS 108(3):15624–15630
74. Myers CW, Gluck KA, Gunzelmann G, Krusmark M (2010) Validating computational cognitive process models across multiple timescales. J Artif Gen Intell 2(2):108–127
75. Sun R, Coward LA, Zenzen MJ (2005) On levels of cognitive modeling. Philos Psych 18(5):613–637
76. Alderson DL, Doyle JC (2010) Contrasting views of complexity and their implications for network-centric infrastructures. IEEE Trans Syst Man Cyber A 40(4):839–852