

A Fitness-Utility Model for Design Science Research

T. GRANDON GILL and ALAN R. HEVNER, University of South Florida

Current thinking in design science research (DSR) defines the usefulness of the design artifact in a relevant problem environment as the primary research goal. Here we propose a complementary evaluation model for DSR. Drawing from evolutionary economics, we define a fitness-utility model that better captures the evolutionary nature of design improvements and the essential DSR nature of searching for a satisfactory design across a fitness landscape. Our goal is to move DSR to more meaningful evaluations of design artifacts for sustainable impacts. A key premise of this new thinking is that the evolutionary fitness of a design artifact is more valuable than its immediate usefulness. We conclude with a discussion of the strengths and challenges of the fitness-utility model for the performance of rigorous and relevant DSR.

Categories and Subject Descriptors: F.0 [Theory of Computation]: General; H.1.1 [Models and Principles]: Systems and Information Theory; K.0 [Computing Milieux]: General

General Terms: Design, Theory

Additional Key Words and Phrases: Design science research, fitness, utility, evolutionary economics, sustainability, research paradigms

ACM Reference Format:

Gill, T. G. and Hevner, A. R. 2013. A fitness-utility model for design science research. *ACM Trans. Manage. Inf. Syst.* 4, 2, Article 5 (August 2013), 24 pages.
DOI: <http://dx.doi.org/10.1145/2499962.2499963>

1. INTRODUCTION

Current thinking in design science research (DSR) uses some form of *utility* as the primary research goal (e.g., Hevner et al. [2004, p. 80]). In this context, the close relationship of utility to practical *usefulness* is emphasized. The choice of usefulness as the pre-eminent dependent variable for DSR explains our particular interest in design artifacts—the concrete products of the design process. Equally important, usefulness ties DSR to earlier management information systems (MIS) research exploring appropriate dependent variables for information systems [DeLone and McLean 1992, 2003]. It also establishes a clear relationship between DSR and the influential technology acceptance model (TAM) for information systems, where usefulness plays a pivotal role in motivating use [Venkatesh and Davis 2000]. Given that usefulness provides DSR with both a clear focus and strong connections to well established research streams, does it even make sense to question if usefulness should always be our central criteria for evaluating the results of DSR?

In this research commentary, we re-open the question of the dependent variable in DSR. In a contrarian spirit, we consider a pair of alternative dependent variables: *design fitness* and *design utility*. In the case of fitness, we particularly focus on its biological meaning—the ability of an entity to reproduce itself and evolve from generation

Authors' address: T. G. Gill and A. R. Hevner, University of South Florida, 4202 East Fowler Avenue, Tampa, FL 33620; email: {grandon, ahevner}@usf.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2013 ACM 2158-656X/2013/08-ART5 \$15.00

DOI: <http://dx.doi.org/10.1145/2499962.2499963>

to generation. In the case of utility, rather than treating it as a synonym for usefulness, we employ its meaning in fields such as economics and decision sciences, where it serves as the basis for ranking decision alternatives. Naturally, usefulness plays an important role in determining both fitness and utility. Neither of these variables, however, is solely determined by usefulness. Indeed, we believe that understanding the relationship between the three variables via a new fitness-utility model complements current thinking and provides important insights into the nature and potential benefits of DSR in the MIS and information and computer technology (ICT) fields.

We begin by introducing the concept of a *design fitness landscape* and explain its relationship to the *design problem space*. We then explore the nature of our two proposed dependent variables, fitness and utility, as they are defined in biology, economics, and in the emerging interdisciplinary field of evolutionary economics. Subsequently, we consider how these concepts can be employed in the context of artifacts and designs; particularly in the rigorous evaluation of designs. Finally, the specific benefits and challenges of applying the fitness-utility model for DSR are discussed.

2. DESIGN FITNESS LANDSCAPES

In evolutionary biology, the term *fitness landscape* is used to describe a functional mapping between some abstract representation of an entity—such as a listing of attributes and traits or, even, as a DNA sequence—and its associated fitness that captures the entity's ability to survive, reproduce, and evolve from generation to generation. This concept can be generalized to design situations, whereby a design is represented as a collection of traits and its fitness represents the likelihood that all or some pieces of the design (which we informally refer to as *design DNA*) will continue to exist and evolve from generation to generation.

It is important to emphasize that the fitness landscape represents only one component of the broader design problem. For the purpose of the present article, we treat the design process as having the following three key elements.

- (1) *Design Fitness Landscape*. A mapping between attributes and fitness that exists in the real world, but which is not observable.
- (2) *Design Problem Space*. The collection of representation, rules, and mappings that exists, individual and collectively, within the mind or minds of the designers. We ignore the distinction drawn between problem space and design space made by some researchers (e.g., Purao et al. [2002]), instead combining them.
- (3) *Design Artifacts*. Tangible products of the design process that have moved from the problem space into the real world.

The relationship between these three elements is presented in Figure 1. A particularly important aspect of Figure 1 is the assumption that only one element of the design process is externally observable: the artifact. This assumption is not particularly limiting. Indeed, it is fully consistent with widely held perspectives relating to DSR (e.g., March and Smith [1995] and Hevner et al. [2004]). Where the present commentary offers an alternative perspective is in the role played by the rightmost two elements of the figure: *fitness* and *utility*.

3. FITNESS AND UTILITY

The central concept of this article is the idea that each and every artifact has an associated quality that we refer to as *fitness*. Fitness, however, is presumed to exist only in the abstract—it cannot be measured directly and its true value unfolds only over time. For it to be of any use to designers and researchers, we therefore need some construct within the problem space for estimating it. We refer to that construct as *utility*. Both of these concepts have a long history outside of DSR; fitness being central to

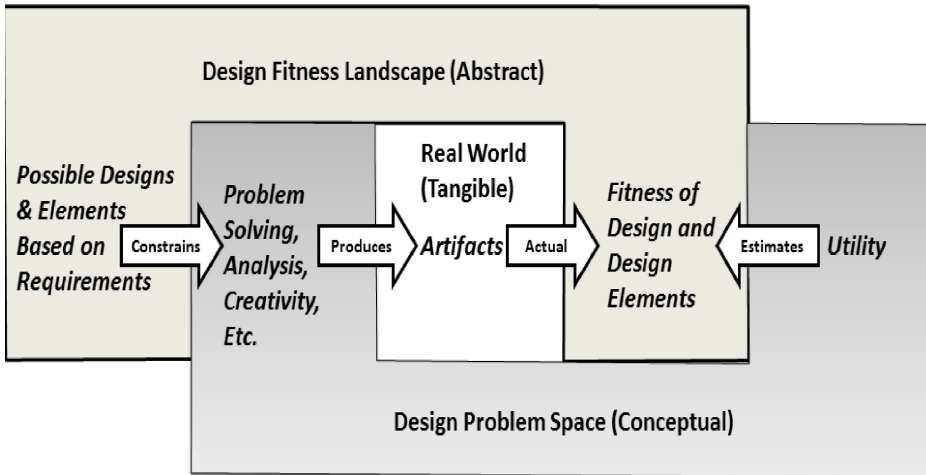


Fig. 1. Relationship between design elements.

evolutionary biology and utility being a fundamental building block of economics. The natures of these concepts, as well as the subtle differences that exist between them, are considered next.

3.1. Fitness

To understand fitness, it is useful to contrast two alternative definitions that might be used to characterize the “fitness” of an organism.

Fitness Definition #1. The fitness of an organism describes its ability to survive at a high level of capacity over time.

Fitness Definition #2. The fitness of an organism describes its ability to reproduce—completely or in part—and evolve over successive generations.

Which definition of fitness you are likely to prefer depends on your perspective. If the individual in mind were our personal physician, we would strongly prefer he or she focus on definition #1. Terms such as physical fitness, mental fitness, and emotional fitness all correspond to this general class of definition. If, however, the individual were an evolutionary biologist, definition #2 would be overwhelmingly preferred. An organism lacking the capacity to reproduce and evolve rapidly goes extinct. What is important about the distinction between definition #1 and definition #2 is that their outcomes are not necessarily correlated. This is graphically illustrated by the experience with human populations, as discussed in Example 1.

Example 1 (Two Versions of Fitness in Populations). At the end of the 18th century, Thomas Malthus proposed that any increases in the individual fitness (definition #1) of human populations would lead to a rapid increase in reproductive rate (a contributor to definition #2) that would quickly erase the gains in individual fitness and would, in the long run, reduce individual fitness since gains in food supplies tended to be arithmetic whereas changes in reproductive rates tended to be geometric [Gill et al. 1992, p. 6]. What has actually happened, however, is in stark contrast to predictions. After a period of adjustment, as individual fitness increases, evolutionary fitness has actually declined.

To illustrate this phenomenon, it is useful to consider two measures: life expectancy (a proxy for definition #1) and fertility rate (a proxy for definition #2). In an organism that employs sexual reproduction, the fertility rate represents the number of children each female of the species produces over her lifetime. In human populations—where the number of male babies is slightly higher than the number of female babies—a stable population requires a value slightly over 2. In much of the industrialized world, this value has fallen far below that stable value. For example, the 2006 U.N. Economic and Social Affairs agency estimated Japan’s 2000-2005 fertility rate at a shockingly low 1.29. During the same period, the U.S. had an estimated value of 2.04. Based on definition #1, the fact that Japan has the highest life expectancy in the world among major industrialized nations would imply high fitness. With respect to definition #2, on the other hand, such low birth rates suggest a population that is decidedly unfit from an evolutionary standpoint.

When considering design fitness, the choice of definition is also going to be context dependent. In the context of *routine design*, the immediate usefulness of the specific artifact being produced is likely to be of greatest concern, favoring the first definition. As the context of design moves closer towards research and development, we are likely to be more concerned about the ability of our artifacts to evolve under changing environments and circumstances, favoring the second.

No matter which definition is preferred, the true fitness of an organism or artifact is going to be unknowable in advance. Fitness landscapes change over time as a result of forces such as the collective impact on the environment of the entities, the impact of co-evolution of other entities, and the impact of unpredicted events that occur entirely outside of the systems being studied, popularly referred to as black swans [Taleb 2007]. Thus, since we cannot “know” or measure fitness directly, our best hope is to estimate it.

3.2. Utility

Similar to fitness, the term, *utility*, is used in a number of ways. When we consider the utility of a tool, we are normally referring to its usefulness. As the term is currently applied in the context of DSR, that is the prevailing meaning. For example, Hevner et al. [2004, p. 83] state: “The utility, quality, and efficacy of a design artifact must be rigorously demonstrated via well-executed evaluation methods.” This implies utility to be a characteristic of the design, one that exists in addition to other characteristics (e.g., quality, efficacy) that might well be part of a design’s fitness.

Economists, on the other hand, employ the term, *utility*, in a different way. Specifically, they posit each individual to have a utility function that can be used to rank choices in the context of decision-making. Each choice is expressed in terms of a set of attributes, for example, x_1, x_2, \dots, x_N , with the form of the utility function as follows:

$$\text{Utility} = u(x_1, x_2, \dots, x_N).$$

The assumption that individuals seek to maximize their personal utility is foundational to the field of economics [Gill 2008]. Each individual is assumed to have his or her own individual function that guides decisions and choices. These individual functions may, in turn, be influenced by cultural and social forces. In early economic theory, the prevailing assumption was that utility is determined by current consumption. More recently, however, it has been generally recognized that many factors contribute to individual utility beyond direct consumption, such as relative income, expectations, social context, and goals [Gill 2010].

To distinguish between the two usages of the term utility, we will refer to the first as *usefulness*. In this context, we apply the broadest meaning of the term—including

factors such as efficacy in performing the task (including performance), range of task cases performed, ease of use, ease of learning, and cost-benefit in the performance of a task. Essentially, we assume that any artifact characteristic that impacts task performance directly can be classified under the usefulness category. Therefore, if our choice of a tool was dictated strictly by usefulness, as just defined, then there would be little reason to distinguish between the two meanings of utility. When we employ the term utility in the remainder of the article, however, we assume its economic meaning and further assume that it consists of a complex function that varies according to individual and that it is not adequately captured by the single usefulness dimension.

3.3. Evolutionary Economics

Evolutionary economics and related fields such as evolutionary psychology examine human choice and preference through the lens of evolution. In the case of evolutionary economics, a central tenet serves to tie the concepts of fitness and utility together. The rationale is stated concisely as follows by Gandolfi et al. [2002, p. 97]: “Given the logic of natural selection, it is difficult to conceive how, for any living entity, a preference for maximizing fitness could fail to evolve.” Stated another way, our utility function provides us with an *estimate of fitness* for the choices that we encounter.

There are two implications of the fitness-utility relationship that are not immediately obvious. The first is that maximizing evolved utility will not always maximize the individual entity’s (Definition #1) fitness. For example, altruism—the desire to give without any expectation of reciprocation—is a widely observed trait in the human population [Fehr and Fischbacher 2003]. At the individual level, it is not clear how such behavior improves an individual’s chances for survival (Definition #1). By ameliorating short-term shortages of resources, however, it should increase overall species fitness (Definition #2). As a consequence, on an evolutionary time scale, we would expect populations predisposed towards altruism to increase in proportion compared to those that do not.

The second implication of the fitness-utility relationship is a pronounced tendency towards diversity on dynamic landscapes. Where fitness landscapes are subject to frequent change a population grouped on a single fitness peak is highly vulnerable. A major shock (e.g., a black swan) can lead to the extinction of the entire species. Having members broadly spaced across the fitness landscape—that is, diversity—increases the likelihood that some individuals will exist at positions where fitness is not seriously reduced (and may even be improved) where such a shock occurs. As illustrated in Example 2, a landscape subject to frequent shocks will therefore tend to evolve a subpopulation of individuals genetically predisposed to seek new peaks.

Example 2 (Explorers in the Population). Imagine a population that consists of two genetically determined types of individuals: explorers and settlers. Explorers are characterized by a tendency to search the fitness landscape for new peaks. Settlers prefer a safer approach to life: they “settle” for seeking local peaks that are occupied by others like themselves. Further suppose that explorers have a much lower survival rate than settlers—a situation that would certainly have been true for human explorers. The question: after many generations, would there be any explorers left?

The likely answer is yes, provided the fitness landscape is dynamic. On such a landscape, most of the explorers will not survive as they visit new locations. Nevertheless, a few will survive to find sites of acceptable fitness and, over time, may attract some settlers to those sites. When a shock to the landscape occurs that disrupts the fitness of the main population sites, however, at least some of the new sites are likely to continue to exhibit acceptable fitness. Since these new sites will necessarily have a higher percentage of explorers in their population, the overall percentage of explorers

in the population will rise. Long periods of stability will tend to reduce the percentage of explorers, while periods of frequent disruption will tend to increase their relative percentage (largely through reduction in the population of settlers). The fact that a particular trait can exist in a particular portion of the population, but not all of it, is an example of what biologists refer to as an *evolutionarily stable strategy* (ESS) [Hines 1987].

A sensible argument can be made that our current utility preferences may not serve as very good estimates of fitness. Numerous researchers have demonstrated that we, as human beings, are far from rational in our processes of choice [Ariely 2007]. There are a number of ways to respond to this.

- First, evolution is slow and, particularly over the past 250 years, changes in the environments we face as a consequence of the industrial and information ages have been so rapid that it would be inconceivable that our utility preferences could have kept up solely through human evolution via natural selection. Fortunately, our built-in genetic utility function also imbues most of us with a desire to learn and, as a consequence, our utility function can adapt to our changing environment through that mechanism, as opposed to natural selection.
- Second, it is actually very rare that we encounter tasks that have the same well-defined inputs as lab experiments (e.g., where probabilities are fully known) in our day-to-day life. Indeed, it is often the case that when we attempt to quantify such values in order to make our decision making more precise, we are in fact deluding ourselves.
- Third, even individuals who have done extensive research into the “irrationality” of our decision rules acknowledge that there are many contexts where these irrational decision rules prove to be beneficial [Ariely 2010].

The key point here is that utility can be treated as the mechanism by which we make choices when confronted with a fitness landscape. Obviously, it is not perfect. Rather, it represents an estimate-of-fitness that we can apply to make decisions presented by such a landscape. More importantly, as we learn more about the characteristics that impact design fitness, our individual utility functions can and will change. Change the designer’s utility function and you change the design process.

4. THE FITNESS-UTILITY MODEL IN DESIGN SCIENCE RESEARCH

The concepts of fitness and utility can readily be applied to the design of information and computing technology (ICT) systems. Returning to Figure 1, each artifact has an associated fitness that the designers estimate through their design utility functions. Artifacts perform two key roles in the design search process.

- (1) They provide evidence that a particular design candidate is feasible, has value, can be effectively represented, and can be built. This serves to help us better estimate the shape of the design fitness landscape.
- (2) They provide a mechanism for communication between designers and for retaining information that might be imperfectly stored during the design process (e.g., sketching a diagram for later recall).

Both of these roles allow designers to refine their estimate-of-fitness (i.e., utility function) through learning and collaboration.

Where design systems differ from biological evolution is in the role played by *intentionality*. The mechanisms of evolutionary change—such as production of new gene combinations through sexual reproduction and mutation—are posited to exert their influence with considerable randomness. While survival rates serve to cull the low

fitness organisms from the population, the actual construction of such organisms is unguided. In the design problem space, on the other hand, designers intentionally concentrate on areas of the design fitness landscape where promising candidates have been identified. What this means is that while utility serves as an estimate-of-fitness for design artifacts, it also feeds back into the fitness landscape itself since a low fitness evaluation for a particular design candidate will discourage further investigations into nearby regions of the design landscape. This, in turn, reduces the fitness of those regions since placing less effort into building artifacts based on a particular design will necessarily reduce the flow of future artifacts based on that design (which is how we define fitness).

Moreover, the shape of the utility function is likely to be guided by two forces: the nature of the evaluation artifacts being studied and by actual experience from artifacts developed for use. Thus, the experience of artifacts placed in practice has the ability to impact the design process just as evaluation artifacts in development do. Thus, the new fitness-utility model can reframe the goal of DSR as follows.

The goal of DSR under a fitness-utility model is to impact design processes so as to increase the likelihood of a continuous flow of high fitness design artifacts. This impact is accomplished in two ways: through the production of artifacts that demonstrate the feasibility of new designs and through improving the utility function that we use to estimate the fitness of these artifacts.

4.1. DSR Evaluation under the Fitness-Utility Model

The use of the fitness-utility model extends and complements current thinking and execution of DSR. Appendix A posits and explores how the new model might impact the seven DSR guidelines proposed by Hevner et al. [2004, p. 83]. The most interesting and extensive changes would be in our understanding and application of Guideline #3 concerning the evaluation of the design artifact. The evaluation would now be based on a more extensive and detailed utility function that estimates the fitness of the artifact.

The fitness-utility model recognizes a large number of characteristics (i.e., design attributes) that could potentially impact design fitness. These are illustrated in Figure 2. As the intersecting ellipses suggest, we continue to expect “usefulness” to play a major role in design fitness. The area within the fitness ellipse outside of the intersection with the usefulness ellipse reflects other characteristics that can impact fitness that are not a direct consequence of usefulness (although they may be correlated with it). The specific characteristics listed in Figure 2 are intended to serve as a preliminary list to be discussed in this section. These characteristics serve one or more purposes that help to support sustainable growth of designs over time.

- (1) They support the design’s ability to evolve incrementally;
- (2) They encourage experimentation by users and other designers; and
- (3) They are effective *memes* [Dawkins 1976; Dennett 1990; Gleick 2011], meaning that they contain ideas of a form that propagate and replicate.

We now consider each of the listed characteristics in Figure 2, again emphasizing that this list is exploratory and not intended to be complete. For each characteristic a brief argument is made as to why it should be included in a design artifact’s utility function. We begin; however, by considering the peculiar region labeled “Too Useful?” where design usefulness actually appears to undermine fitness.

4.1.1. Designs That are Too Useful? Figure 2 was drawn in a manner implying that usefulness does not *always* contribute to design fitness. We saw this phenomenon previously, in Example 1, where increased life expectancy (i.e., fitness definition #1) has

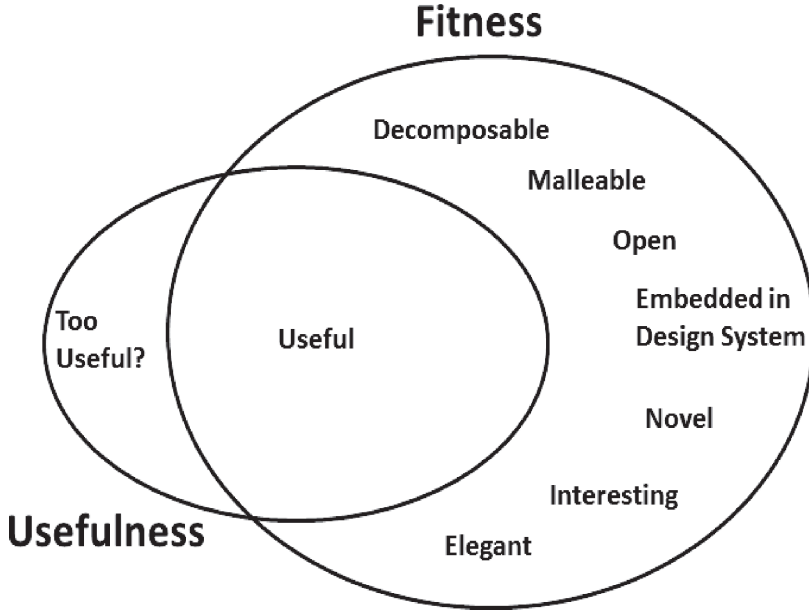


Fig. 2. Design candidate fitness characteristics and usefulness.

become associated with below replacement fertility rates (i.e., fitness definition #2). The analogy in DSR would be situations where a design artifact becomes so useful that it actually inhibits future design activity.

In abstract terms, as the complexity of a fitness landscape grows, so too grows the number of peaks where any incremental change leads to a loss of fitness [Kauffman 1993]. As technologies evolve, however, the landscape changes from a practical standpoint as new peaks become accessible enabled by new artifacts. For example, new devices became possible when the transistor replaced the vacuum tube. Unfortunately, the existence of a higher peak does not mean it will be explored by designers. Sometimes a fitness peak becomes so sticky that efforts to explore other peaks are abandoned or ignored. Such a strategy may make short term sense, particularly where marginal analysis is used by an agent already occupying a peak. It can be very expensive to transition from one design to a very different design. Over the long term, however, the highly useful proprietary design that cannot easily adapt to the changing landscape represents an evolutionary dead end. As such, it could not be considered a high-fitness design.

In fact, the tendency of organizations to stick with designs that have proven useful is a well-documented phenomenon known as the Innovator's Dilemma [Christensen 1997], described in Example 3. Many of the examples used to inform Christensen's research were drawn from ICT, including disk drives, printers, and minicomputers.

Example 3 (The Innovator's Dilemma). In *The Innovator's Dilemma*, Christensen [1997] describes a phenomenon that has been played out in many industries. An innovative company introduces a technology-based product that captures a substantial share of the market. Another competitor then introduces a product based on a different design—often radically different—that serves a customer base peripheral to that of the original company. The economics of redesigning the original product, combined with the limited impact of the new design, prevent the original company from responding; it continues to refine the original design.

Over time, however, the fitness landscape changes. Because the new design has less benefit of experience, improvements to it tend to occur at a more rapid pace than for the original design. Gradually, the new design builds market share, as continuing gains of efficiencies of scale and cost reductions make it applicable to a broader range of customers. Eventually, however, the new design reaches a point where its cost-benefits exceed those of the original design even for existing customers of that design. At that point, customers switch en-masse to the new design.

4.1.2. Decomposable Designs. The seminal work that launched the current study of design science is Herbert Simon's [1996] *The Sciences of the Artificial*. The second half of the book is largely devoted to explaining why systems tend to evolve from nearly decomposable subsystems. Indeed, even under the existing DSR goals, decomposability is likely to exert a strong influence on design quality and would therefore be evaluated as part of the design. In addition, such systems tend to be easier to construct, since work on individual components can be conducted separately. As an illustration, a study of the open source community, described in Example 4, found that four factors were identified by participants as encouraging contributions. Three of these were closely tied to the ability to work on pieces of the project and to make choices independent of the overall design.

Where decomposability can be particularly critical is in the reproduction and evolution of partial designs, relating directly to the ability of a design to evolve incrementally. Where a design cannot be decomposed into nearly independent subsystems, evolution of the design would tend to be a matter of all-or-nothing. Where a design is built upon separable systems or constructions, on the other hand, pieces of the design—strands of *design DNA* to use a biological analogy—may exhibit high fitness and evolve rapidly while others may remain static or be discarded. This can reduce the likelihood of the entire design becoming stuck on the “too-useful” peak.

Strands of design DNA can often survive long after their original artifact has been abandoned. For example, the brace and semi-colon syntax of the C programming language [Kernigan and Richie 1988] was heavily influenced by its immediate predecessor, the B programming language, which was, in turn, influenced by the BPCL language. Subsequent to its release, the syntax of C then influenced numerous later languages, including C++, Java, C# and many others.

Example 4 (Open Source Contributions and Decomposability). The open source software movement provides an interesting quasi-experiment in incremental design, since participants normally make contributions voluntarily. As a consequence, designs that attract such contributions would seem to be the most likely to evolve.

A study of what motivates open source developers [von Krogh et al. 2003, p. 1231] identified four key factors that would act as a “contribution barrier” if not in place. These were as follows:

- (1) ease of modifying and coding module;
- (2) the extent to which the potential developer can choose the computer language used to code for the module can vary;
- (3) ease with which to “plug” the module into the architecture; and
- (4) the extent to which a module is intertwined or independently working from the main code.

The last three of these are closely tied to the decomposability of the code base. Item 2 addresses the ability to code separate components in different languages. Items 3 and 4 could practically be used as definitions of decomposability in the context of software architecture.

4.1.3. Malleable Designs. Often enhanced by decomposability, the malleability of an artifact represents the degree to which it can be adapted by its users and respond to changing use/market environments [Gregor and Jones 2007; Williams et al. 2010]. MIS research has demonstrated that users frequently employ tools for unintended purposes. We would expect that such adaptations would allow designers to evolve artifacts to support these uses more effectively, making malleability a good example of a design characteristic that encourages experimentation. For example, the recognition that early spreadsheet users of VisiCalc were frequently using the tool to save data, such as contact lists, encouraged the designers of Lotus 1-2-3 to incorporate basic database capabilities into their spreadsheet product.

Design malleability can have many variations. For example, some researchers distinguish between *exaptation*, where a software developer (other than the original designer) exploits a component for an unintended purpose [Gregor and Hevner 2013] from *customization*, where a user adapts the component to unexpected purposes [Heineman 1998]. In the ICT domain, the ability to inherit code, for example, provides a mechanism by which developers can modify existing components without altering the original. Tools such as scripting languages are sometimes incorporated into application designs to provide power user-malleability.

User-malleability itself can be divided into levels, such as customization, integration and extension [Morch 1997]. Here customization refers to the ability of a design to be tailored to a user's preferences. Integration involves the ability to conveniently share the capabilities of one artifact with another, creating a resultant artifact with capabilities beyond those of either original design. The ability to create mash-ups of web components on a single page is an example of this capability. Extension means adding new capabilities to an artifact. Empirical research has documented numerous examples where users—particularly advanced “lead-users”—have exerted a strong influence on product design. The industries where such impact has been identified are broad, and include software development, industrial products and processes, and consumer products [von Hippel 2001, p. 84].

As previously noted in Example 4, ease of modification of source code reduces barriers to contribution in the open source community. In this community, the roles of developer and user are frequently co-mingled. For example, in one study, non-work need for code—meaning that the contributor expected to be a user of the application—was singled out as the principal motivator of open source contributions for a cluster of 27% of the respondents [Lakhani and Wolf 2003, p. 13].

4.1.4. Open Designs. Another design characteristic that has the potential to impact design fitness is the degree to which artifacts are open to inspection, modification, and reuse. Open designs—particularly when also imbued with decomposability and malleability—encourage further design evolution by making it easier both to see how an artifact is constructed and to modify existing components of the artifact. For example, an information system created as an open source application has a significant advantage over a proprietary design in terms of its ability to evolve rapidly based on changing user conditions, as illustrated in Example 5.

Example 5 (UNIX vs. LINUX). The potential for openness to impact design evolution is evident in the case of the UNIX and LINUX operating systems. What makes this example particularly relevant is that the two operating systems were, in fact, designed to be nearly identical from the user's perspective. The key distinguishing feature was that LINUX was always intended to be open, whereas the openness of UNIX has remained the subject of considerable debate.

The history of the two operating systems is summarized in the “OSI Position Paper on the SCO-vs.-IBM Complaint” [Raymond and Landley 2008]. The timeline begins in

1956, when AT&T signed a consent decree that prohibited it from selling computer technology (hardware or software) in response to an antitrust suit brought against it by the U.S. government. Because the company was and remained the largest user of computers in the world until it was broken up in 1984, its Bell Labs unit focused on developing technologies for its internal use.

In the late 1960s-early 1970s, the company designed an operating system intended to be portable across AT&T's many hardware platforms. Because it was prohibited from selling the code, it established UNIX as a standard and worked cooperatively with a number of academic institutions, such as UC Berkeley, who developed versions of the operating system for different platforms. In 1984, however, newly reorganized AT&T began to develop a commercial version of the UNIX operating system. In 1990, the company moved its UNIX operations into a separate subsidiary (UNIX Systems Laboratory, Inc.), which then sold shares to outside companies, beginning a complex trail of ownership that ultimately needed to be untangled in the court system.

The LINUX operating system was developed by Linus Torvalds to mimic the behavior of Minix, itself a textbook implementation intended to mimic the behavior of the UNIX operating system. A key factor distinguishing LINUX from the other operating systems was its unqualified origins as an open source project. That openness allowed the operating system to evolve rapidly. Major IT companies, such as IBM, chose to re-focus their development effort on this open system, sometimes even abandoning their own existing UNIX-based technologies, such as IBM's AIX, in the process. As a result LINUX frequently surpassed the original operating system it was based upon. For example, LINUX installations are used on over 91% of supercomputers, compared with 6% for UNIX [TOPS 2011].

LINUX's openness also facilitated its redesign for alternative platforms. For example, it served as the basis of Google's open Android operating system, installed on over half of all smartphones sold by mid-2011 [Gartner 2011], which itself evolved through four major versions in the first three years after its initial commercial release in 2008.

4.1.5. Embedded in a Design System. We would expect design artifacts that are the product of a sustainable design system environment to evolve more rapidly than artifacts that are produced in a context where design is an unusual activity. The particular purpose that such systems play is encouraging communication within and throughout the design process. A particularly effective illustration of such a system can be found in Apple, as described in Example 6.

This particular source of fitness can sometimes act as a counterweight to openness, as organizations with highly effective research and development activities may be reluctant to open up their designs and may use legal measures, such as patents and copyrights, to discourage unauthorized parties from evolving the original designs. Such measures not only can slow the progress of design evolution but can also increase the marginal cost of transitioning to a new design, thereby increasing the risk of becoming trapped on a peak whose design fitness is declining over time.

An effective design system can produce a stream of design artifacts, without ever succeeding in transforming these designs into use artifacts. A good example of this is Xerox's Palo Alto Research Center (PARC) which, during the 1970s and early 1980s, developed a set of artifacts that represent the core of today's desktop computing environments—the windows-based graphic user interface, the mouse, WYSIWYG, object-oriented programming (Smalltalk), and the Ethernet local area network [Parc.com 2012]—without translating the results into successful commercial products. Their openness in exposing their design DNA, however, led to many of their ideas being transformed into products by others, most notably at Apple.

A design system can also manifest itself as a community of users and designers, providing contributors with intrinsic motivation to contribute. Here openness may actually foster—rather than inhibit—the design system. This phenomenon has been observed in the open source community. For example, one study [Lakhani and von Hippel 2003, p. 937] found that the desire to promote the open source movement was the greatest source of motivation for frequent contributors to Apache user support forums. More directly, a related study found a cluster of 19% of all contributors that were motivated “primarily by obligation/community-based intrinsic motivations” [Lakhani and Wolf 2003, p. 13].

Example 6 (Design at Apple). In many ways, the design system at Apple appears to be the antithesis of openness, described as follows [Johnson 2010, p. 169].

Even as much of the high-tech culture has embraced decentralized, liquid networks in their approach to innovation, the company that is consistently ranked as the most innovative in the world—Apple—remains defiantly top-down and almost comically secretive in its development of new products. . . If open and dense networks lead to more innovation, how can we explain Apple, which on the spectrum of openness is far closer to Willy Wonka’s factory than it is to Wikipedia?

The answer to this paradox appears to be in the unique design system that Apple has established for product development. The design system for a use artifact at most companies might appear as follows [Johnson 2010, p. 170].

The designers come up with a basic look and feature set and then pass it on to the engineers, who figure out how to actually make it work. And then it gets passed along to the manufacturing folks, who figure out how to make it in large numbers—after which it gets sent to the marketing and sales people who figure out how to persuade people to buy it. . . the original idea gets chipped away at each stage of the chain.

At Apple, a process that is far less sequential and vastly more integrative is employed, described as follows [Johnson 2010, p. 171].

All the groups—design, manufacturing, engineering, sales—meet continuously throughout the product development cycle, trading ideas and solutions, strategizing over the most pressing issues, and generally keeping the conversation open to a diverse set of perspectives. . . the results speak for themselves.

There is a drawback to the closed design system employed by Apple, however. While it has proven extremely effective in innovating and popularizing new designs, it has also invited imitation by competitors that employ more open designs. Thus, Apple’s original lead in micro-computing was quickly lost to IBM open industry standard architecture (ISA) for PCs. Its early lead in windows-based computing was lost to Microsoft’s Windows, which was more open to third party developers and hardware manufactures. Most recently, its innovative iPhone has been outsold more than 3-to-1 by similar devices running Google’s Android operating system (see Example 5).

4.1.6. Design Novelty. A design may be considered novel if it originates from an unexplored region of the design fitness landscape. Once such a design candidate has proven viable, other design candidates from the same region may follow in an attempt to locate the local peak on the fitness landscape. Novelty alone is rarely enough to achieve high design fitness. In a sense, the role played by novelty in design fitness parallels

the trade-off between explorers and settlers described earlier in Example 2. While a particular novel design may be less individually fit than existing counterparts, where the landscape is dynamic the fitness of the population as whole benefits from having a subpopulation of designers seeking novelty for its own sake, thereby ensuring design diversity. Where genuine novelty is combined with other traits that make elements of its DNA effective memes, then the design may evolve rapidly. The proliferation of tablet designs that followed Apple's introduction of the iPad is an example of this.

Novel design artifacts present a particular challenge to traditional DSR. The creative process through which they are envisioned may not meet the criterion of usefulness and rigor suggested by the original guidelines and the potential benefits of the design may be hard to evaluate. A genuine new invention is a difficult goal for DSR research projects and we can expect few research contributions to be true inventions [Gregor and Hevner 2013]. However, exploration for new ideas and artifacts should be encouraged regardless of the hurdles.

4.1.7. Interesting Designs. Normally, a design artifact is created in order to explore or demonstrate some specific purpose. From time-to-time, however, an artifact may demonstrate unexpected emergent behaviors that are worthy of subsequent investigation and the creation of subsequent artifacts. An artifact may also be constructed in an unexpected way that intrigues other designers or design researchers. We characterize such designs as *interesting*.

While there is likely to be considerable overlap between designs that are interesting and designs that are novel, the two characteristics are not identical. We have framed the benefits of novelty in terms of contributing to the diversity of the design landscape. Broadly speaking, the benefit of an interesting design is its propensity to diffuse, to be an effective meme. In fact, requiring a design to be interesting may serve as a limitation on its novelty. Social scientists (e.g., Davis [1971]), for example, find research that largely conforms to existing expectations yet also incorporates an unexpected element is most likely to interest other researchers. Research that departs too far from existing paradigms, on the other hand, is likely to earn its author the title of "crackpot".

The degree to which a design is viewed as interesting is likely to be closely related to the concepts of stickiness [Gladwell 2000] and resonance [Gill and Bhattacharjee 2009]. Among the characteristics that make an idea sticky are simplicity, unexpectedness, concreteness, and credibility [Heath and Heath 2007]. These four characteristics seem equally applicable to a design artifact. In consequence, we would expect artifacts that are interesting by virtue of these traits will diffuse particularly well.

4.1.8. Design Elegance. In many areas of design, such as architecture, consumer products and apparel, there is an ongoing tension described as form versus function [Alexander 1964]. Function relates to practical usefulness. Form, in contrast, describes aesthetic elements such as appearance that do not necessarily serve a useful purpose, yet nevertheless increase the user's utility. The characteristic of an ICT design artifact that corresponds to form might best be referred to as *elegance*. Like quality, elegance is hard to define in a rigorous manner and yet characteristics that might be associated with it, such as compactness, simplicity, transparency of use, transparency of behavior, clarity of representation, can all lead to designs that invite surprise, delight, imitation, and enhancement. Equally important, they can cause a design artifact whose usefulness has yet to be demonstrated to endure. This last aspect of elegance is illustrated in Example 7, tracing the use of Boolean algebra.

Example 7 (Boolean Algebra). George Boole was a schoolmaster who lived from 1815–1864. Although he initially made his reputation in mathematics, he was most famous for developing a presentation of logical inference modeled after mathematical

operators, later known as Boolean algebra [Burriss 2010]. As a consequence of the elegance of his system—as opposed to its practical uses, which were largely unknown—Boole’s system was incorporated into many courses in philosophy and logic.

During the late 1930s, Claude Shannon, then an undergraduate at the University of Michigan, became acquainted with Boole’s work in one of those courses. Shannon later applied the algebra to the design of switching circuits as part of a Master’s thesis at MIT. As a consequence of that work, Boolean algebra was used extensively as a representation for switching circuits and then digital logic. Bits and pieces of the Boolean algebra design DNA persist in today’s information systems. Search expressions are often referred to as Boolean operators. The true/false data type is declared as `bool` in many languages, such as Java and C#. Thus, we see how an artifact largely preserved based upon its elegance ended up ultimately ending up becoming instrumental in the development of subsequent design and use artifacts.

4.2. Synthesis: Extended Characteristics

These seven characteristics in addition to usefulness (Figure 2) that impact design fitness represent a preliminary list, one that will doubtless require revision and clarification through future empirical and explanatory research. One way of thinking about the characteristics is in terms of related concepts in a variety of disciplines, such as genetics (e.g., mutation, recombination), evolutionary theory (e.g., diversity, longevity), technology acceptance theory (e.g., quality, consistency, motivation), and innovation theory (e.g., diffusion). If a design is going to evolve, it needs to be modifiable; designers need to be motivated to move it forward; it should contribute to the overall diversity of the design artifacts in existence; it should be in a form that encourages diffusion to other designers; and it needs to endure long enough so that a use can be found for it. The specific characteristics proposed here roughly map to these outcomes, as illustrated in Figure 3 (although overlaps certainly exist).

The combination of decomposability, malleability and openness can also lead to the emergence of a particularly desirable characteristic referred to as *antifragility* [Taleb 2012]. It exists at the end of a continuum that describes how an entity responds to external sources of stress: Fragile → Robust → Antifragile.

Fragile designs produce artifacts that fail in the presence of significant stress, and are therefore likely to be abandoned over time. *Robust* designs produce artifacts that are relatively indifferent to stress. *Antifragile* designs, in contrast, produce artifacts that are likely to improve as a consequence of the presence of stress—much as living systems often do (e.g., through evolution). Decomposability contributes to this characteristic by allowing for incremental improvements; malleability through placing elements of the design in the hands of users; openness through making improvements to the design accessible to other users. Taleb [2012] argues that antifragility tends to be a major source of longevity; in the context of design, this would translate to fitness.

5. DISCUSSION – PROS AND CONS OF THE FITNESS-UTILITY MODEL OF DSR

We assert that the fitness-utility model is a complement to, rather than a competitor of, the existing DSR paradigm. There are two key justifications for this assertion.

- (1) With its focus on reproductive fitness (i.e., definition #2) rather than individual artifact fitness (i.e., definition #1), the fitness-utility model is far less focused on immediate usefulness.
- (2) While the existing DSR model is particularly concerned with artifact creation and, as a result, is oriented towards researching the mechanisms of the design problem space, the fitness–utility model is more focused on artifact evaluation. Inasmuch as definition #2 fitness cannot be observed directly, it therefore needs to concentrate

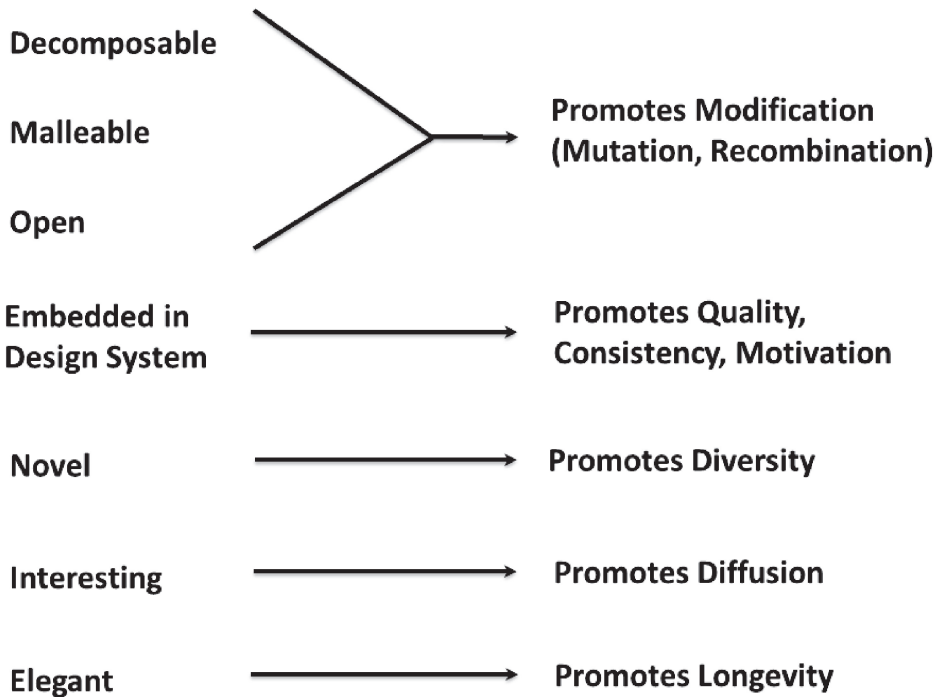


Fig. 3. Fitness characteristics and related outcomes.

on helping designers better estimate artifact fitness through informing their utility function.

We now discuss five potential advantages of DSR guided by the proposed fitness-utility model, followed by three key challenges.

5.1. Advantage: Makes the Researcher an Active Participant in the Design System

Because designers desire to publish research relating to innovative design artifacts, a great deal of DSR in ICT is in the form of action research [Sein et al. 2011]. Under the fitness-utility model, however, even the non-technical researcher strives to play an active role in the design system through increasing the set of artifacts representing achieved states in the design fitness landscape and, in so doing, perhaps making new states attainable that build upon these artifacts. Successful research will, as a matter of definition, lead to either an increase or decrease in the production of new artifacts based upon the specific design candidate or candidates investigated. The fitness-utility model would also be predicted to maximize the potential impact of individual research contributions by focusing on early stage design. Thus, if the researcher's goal is to impact the design problem space, emphasizing artifacts that subsequent designers can build upon, as opposed to artifacts that are immediately useful, may prove to be a promising path.

5.2. Advantage: Provides an Alternative Basis for Evaluating Research Impact

Today, within the MIS research discipline, the impact of research is generally measured through the estimated quality of the publication outlet and through subsequent citations by other researchers. The fitness-utility approach offers another alternative for DSR: charting the evolution of early-stage artifacts influenced by DSR research.

If the artifact continues to evolve and incorporate design DNA deemed favorable by the research, then impact—in the truest sense of the word—has been achieved [Nunamaker and Briggs 2011]. The same can be said of research that stifles the further evolution of design DNA deemed detrimental to fitness. For example, if particular design practice (e.g., allowing the user to enter free form text into a textbox that is then used to query a database) leads to a security threat (e.g. malevolent SQL injection), impactful DSR that identifies this as a low-fitness practice should reduce the frequency of the occurrence in later artifacts.

5.3. Advantage: Aligns with Dynamic Environments

A central premise of this commentary is that over time the evolutionary fitness of design artifacts becomes far more interesting than the use fitness of a particular artifact at a static point in time. The validity of this premise is likely to depend on the environment in which it is situated. For very static environments, for example, a particular use artifact may exist for a very long time. In such a world, the use fitness of the artifact is a matter of considerable interest. In a highly dynamic environment, on the other hand, the artifact's potential to evolve needs to be given much greater weight. Similarly, the benefits of encouraging diversity of designs as insurance against radical fitness landscape shifts, rather than solely for the sake of immediate usefulness, should be greater in such landscapes.

Our belief is that such dynamism describes most environments facing ICT designers today and forces such as globalization, energy conservation, social media, and advances in telecommunications will likely serve to increase environmental turbulence. This advantage is clearly related to the challenge of developing new research methods for defining and measuring design quality in terms of evolutionary fitness.

5.4. Advantage: Recognizes the Inherent Limitations of Intended Usefulness

Our research suggests that while usefulness is likely to be the best single predictor of artifact use (a finding consistent with most TAM research), it is not necessarily a very good predictor when applied by itself. In fact, a reasonable argument can be made that many of the most interesting (as per Davis [1971]) findings of MIS revolve around examples where an ICT artifact's impact was far different from the designer's intended use, as illustrated by Example 8.

Example 8 (Design vs. Use of SABRE System). One of the most famous systems in MIS lore is American Airlines' (AMR) SABRE reservation system. At the time of its original design, the system was viewed as a major technical achievement [Hopper 1990, p. 121] but was principally focused on the internal need of AMR to develop a more efficient approach to booking and assigning seats on its aircraft.

In the late 1970s, however, deregulation fundamentally changed the U.S. airline industry, with routes, ticket sales channels, pricing policy and competition changing dramatically. AMR was able to adapt SABRE to the new environment, using it as a competitive weapon. As travel agents became an indispensable part of the continually changing air-travel market, SABRE was exploited in many different ways to obtain preferential booking at a time when competition between carriers had become intense. Its use contributed to the downfall of a number of competing carriers. It generated profits at a time when losses had become the norm for air carriers.

What is particularly intriguing about the SABRE story is that the system's success in the early 1980s had little to do with its original design objectives. Rather—in fitness landscape terms—it happened to occupy a position whose fitness increased dramatically as a consequence of external change (i.e., airline deregulation).

5.5. Advantage: Encourages Collaboration between MIS Researchers and Designers in Other Fields

The fitness-utility approach specifically targets clients in the design communities supplying the resources necessary for further design evolution. In early-stage ICT design research, these communities will likely contain a preponderance of researchers in technical fields such as computer science and engineering including many academics. Thus, we will have a strong incentive to collaborate with these communities if we are to exert impact. Where we may be able to contribute most effectively is in our understanding of the potential unintended consequences of artifacts employed in an organizational setting [Niederman and March 2012]. Having observed these consequences in the field and studied them in our literature, we are in a unique position to provide perspective to designers who may otherwise become overly focused on technical performance and intended use in static work environments.

5.6. Challenge: Current Research Standards Do Not Reward Design Impact

Given that researcher rewards, including salary and promotion and tenure, tend to be closely tied to measured research impact based on numbers of quality publications and citations, the fact that the fitness-utility model offers another approach to measuring impact over time—tracing how artifact design DNA changes as a consequence of research findings—may not be appealing to academic researchers. The longitudinal nature of such research presents particular challenges to researchers (e.g., nontenured professors, industrial designers with market windows) in need of short-term rewards. In the absence of institutional change with respect to how impact is defined, it may be hard for the fitness-utility model to gain traction.

5.7. Challenge: The Framework for Evaluating Design Fitness Is Not Well Researched

Earlier in this article, we proposed a number of non-use characteristics (Figure 2) that seemed likely to impact design fitness. This list was largely induced from examples and can in no way be considered complete, rigorously derived, or rigorously supported in the absence of future research. Unfortunately, there is little research into the characteristics that provide good estimates of design fitness as we have defined it. Stated another way, our design utility function is largely unexplored. This naturally presents a substantial obstacle to any research that attempts to estimate the fitness of a particular artifact. The largely unexplored forces driving fitness and utility are in stark contrast to the much better established approaches to evaluating design usefulness. Although the field laments its lack of theoretical base, constructs, and generalizability [Hevner et al. 2004, p. 99], it has a plethora of these when contrasted with the fitness-utility model. Add to this the fact that immediate usefulness is likely to seem a more concrete research objective than fitness, and the researcher is likely to have a much easier time designing research under the existing paradigm.

5.8. Challenge: Building Rigor for Fitness-Utility Research Requires Alternative Research Methods

It may be argued that the previous challenge actually represents a considerable opportunity for future research into the factors that lead to fitness. Such research, however, is likely to substantially differ in character from the main body of existing research in the ICT disciplines [Chen 2011]. To understand fitness, you need to look backward in time in order to trace the evolution of an artifact. Indeed, it may take years to validate the actual fitness of an artifact, a necessary step if the characteristics contributing to fitness are to be identified systematically. Thus, historical research methods are likely to play a much greater role than is the case in most contemporary ICT research (see

Table I. Summary of Usefulness and Fitness-Utility Models

Characteristic	Usefulness Model	Fitness-Utility Model
Focus	Useful artifacts	Artifact reproduction and evolution (fitness) and the choice mechanisms guiding artifact design (utility)
Applicable artifacts	Construction and use	Feasibility and evaluation
Unit of study	Entire artifact	“Design DNA” within artifact
Time horizons	Short and medium-term	Long term
Source of rigor	Careful evaluation of intended use and expected performance	Systematic evaluation of non-usefulness factors that may contribute to fitness and the potential for unintended consequences
Most likely external clients	Developers and use clients	Researchers and R&D clients
Source of models	Study of current artifacts in the field	Study of historical progression of artifacts based upon a particular design candidate
Particular value offered by IT research	Understanding the organizational context in which artifact development and use takes place	Understanding the role played by unintended consequences in typical artifact implementation; broad perspective on factors that influence artifact success
Desired impact of research	Improved design and development of useful artifacts and better understanding of the factors that make an artifact useful	Improving fitness of desirable design DNA and suppression of undesirable strands; better understanding of the factors that increase real-world artifact fitness leading to improved choice between alternative design candidates

Mokyr [2002]). In addition, fitness landscapes in general tend to be rugged [Kauffman 1993], meaning that interdependencies between variables prevent decomposability; interaction effects dominate main effects. Such ruggedness can confound traditional statistical techniques [Gill and Sincich 2008]. What this means is that data analysis techniques most preferred by empirical researchers may prove largely inapplicable in the analysis of sources of fitness.

6. CONCLUSIONS

Several times in this commentary, we have posited that the fitness-utility model for design science research is better viewed as a complement to the existing usefulness model, rather than as a competitor. As summarized in Table I, these two variations of DSR (the usefulness model vs. the fitness-utility model) focus on different objectives, are most applicable to different artifacts, tend to examine different units of analysis, are appropriate for different time horizons, are likely to employ different research methods, and will tend to be of greatest interest to different client constituencies. We have already noted that high levels of usefulness may actually inhibit artifact evolution. There is likely to be quite a bit of causality here—an organization making a large investment in designing and deploying a use artifact may be unlikely to view the tendency to evolve rapidly as a major benefit. To the contrary, such a manager is most likely to appreciate an artifact that is highly useful and is likely to remain that way as long as possible. For this clientele, maximizing the fitness of design, as we have defined it, is more likely to be scary than desirable. Moreover, an understanding of the factors contributing to usefulness is central to the fitness-utility model. Many factors outside of usefulness may contribute to fitness, but usefulness or potential usefulness remains at the center of most design settings. What we hope is that the analysis presented in this commentary will inspire further research into the remaining factors that influence design fitness.

As future research directions, we hope to advance these ideas by describing selected historical case studies and existing DSR projects in which evaluation is based on the fitness-utility model. Potential candidates for longitudinal IT design studies under the lens of the fitness-utility model are the following.

- *Decision Support Systems*. The first definitions of decision support systems (DSS) were provided by Scott-Morton [1967]. The subsequent evolution of DSS into Executive Information Systems (EIS) and Group Decision Support Systems (GDSS) has a long history of design innovation in the IT community.
- *Data Mining Methods*. Agrawal et al. [1993] developed the first full conceptualization of mining databases for association rules as well as an efficient method for discovering them. This paper generated and influenced a whole new field of research that is loosely described as the field of business intelligence.
- *Cleanroom Software Engineering*. The concepts and principles of Cleanroom software engineering were developed in IBM by Mills and colleagues [Mills et al. 1986, 1987]. The goals of the Cleanroom process are to produce software with a certifiable level of reliability. While the full scale application of the Cleanroom process has been limited, many of its central design ideas (e.g., incremental development) have evolved into industry best practices. However, other parts of the Cleanroom DNA (e.g., formal statistical testing) have not evolved into wide spread use.

Appendix. Fitness-Utility Model Impacts to DSR Guidelines

In this appendix, we examine how the fitness-utility model might impact the seven DSR guidelines proposed by Hevner et al. [2004, p. 83].

Guideline #1: Design as an Artifact

Guideline #1. Design-science research must produce a viable artifact in the form of a construct, a model, a method or an instantiation. The fact that DSR is constrained to deal with the concrete by this guideline is important in distinguishing it from behavioral research. Moreover, the original guidelines are sufficiently broad in their definition that meta-designs and nascent design theories [Gregor and Jones 2007; Walls et al. 1992] would, themselves, constitute artifacts and would therefore—quite rightly—fall under the DSR heading. The fitness-utility approach, however, emphasizes improving our ability to estimate the fitness of design artifacts, without necessarily producing such artifacts. Thus, the impact of the fitness-utility model on this guideline would be as follows:

Impact to Guideline #1. Fitness-utility DSR seeks to impact design through enhancing our ability to evaluate the fitness of design artifacts.

Guideline #2: Problem Relevance

Guideline #2. The objective of design-science research is to develop technology-based solutions to important and relevant business problems. While we agree with the continued importance of relevance, the problem with the existing statement of guideline #2 is that, from a practical standpoint, it tends to constrain the time horizons for design research. We often cannot foresee what problems will be relevant for the future of IT. The challenge this unpredictability presents to DSR is that if you try to anticipate the important long term problems that a design will solve, it will be nearly impossible to get them right. Thus, being overly problem-focused demands a shorter term outlook. Another way of looking at the issue is to use the analogy of constructing a puzzle. At the beginning of a puzzle, as in a design process, you have a collection of pieces

that can only be put together in certain ways. True “problems”—in the form of missing pieces—tend to be discovered near the end of the assembly, when the gap is identified. If we require that our DSR project solve an important problem, we may need to wait until we know what is missing. In other words, we may not know the important problem we should solve until sometime during the research process. With this caveat, the key implication of the new model to DSR Guideline #2 is the consideration of “high-fitness” prior to technology-based solutions, thereby implicitly asserting our preference for solutions likely to evolve.

Impact to Guideline #2. The objective of fitness-utility DSR is to improve our ability to estimate the fitness of technology-based artifacts.

Guideline #3: Design Evaluation

Guideline #3. The utility, quality and efficacy of a design artifact must be rigorously demonstrated via well-executed evaluation methods. If the fitness-utility model is applied to DSR, then the evaluation criteria are where the utility function is to be shaped. (A thorough examination of DSR evaluation using the fitness-utility model is a principal contribution of this commentary.) That would require an impact statement to the original design guideline along the following lines:

Impact to Guideline #3. The fitness of a design artifact must be estimated using a utility function that involves the full range of characteristics that impact the likelihood that the artifact will replicate and evolve.

Guideline #4: Research Contribution

Guideline #4. Effective design-science research must provide clear and verifiable contributions in the area of the design artifact, design foundations and/or design methodologies. With respect to this guideline, the fitness-utility approach and the original approach are relatively similar. As originally stated, however, it is not clear that research that leads to better understanding of utility (i.e., estimating the fitness of a design artifact) would be included under the design heading. For this reason, impact might be:

Impact to Guideline #4. Effective DSR impacts the design problem space through contributions in the areas of constructing design artifacts, design fitness, design foundations and theories, and/or design methods.

Guideline #5: Research Rigor

Guideline #5. Design-science research relies upon the application of rigorous methods in both the construction and evaluation of the design artifact. A particular challenge associated with the use of the term rigor is that it is perceived to be generally “understood” but is rarely defined. One definition that has been proposed (e.g., Gill [2010]) treats research rigor as consisting of three related elements: (1) the investigation is systematic, (2) a thoughtful balance is struck between the risk of accepting that which is false (Type 1 error) and rejecting that which is true (Type 2 error), and (3) challenging questions are posed. By this definition, the current guideline would tend to place considerable obstacles in the way of early stage design artifacts, such as: (i) systematic search of the design fitness landscape is generally impossible, (ii) current standards of empirical research in the social sciences tend to lean heavily towards avoiding Type 1 error [Ziliak and McCloskey 2008] making rejection of novel ideas more likely, and

Table A.1. Existing vs. Fitness-Utility Design Science Research Guidelines

Guideline	Hevner, et al. [2004, p.83] Description	Fitness-Utility Impact
Design as an artifact	DSR must produce a viable artifact in the form of a construct, a model, a method or an instantiation.	Fitness-utility DSR seeks to impact design through enhancing our ability to evaluate the fitness of design artifacts.
Problem relevance	The objective of DSR is to develop technology-based solutions to important and relevant business problems.	The objective of fitness-utility DSR is to improve our ability to estimate the fitness of technology-based artifacts.
Design evaluation	The utility, quality and efficacy of a design artifact must be rigorously demonstrated via well-executed evaluation methods.	The fitness of a design artifact must be estimated using a utility function that involves the full range of characteristics that impact the likelihood that the artifact will replicate and evolve.
Research contributions	Effective design-science research must provide clear and verifiable contributions in the area of the design artifact, design foundations and/or design methodologies.	Effective DSR impacts the design problem space through contributions in the areas of constructing design artifacts, design fitness, design foundations and theories, and/or design methods.
Research rigor	Design-science research relies upon the application of rigorous methods in both the construction and evaluation of the design artifact.	DSR requires that the construction and evaluation of design artifacts be investigated employing a level of rigor appropriate to the nature and stage of design.
Design as a search process	The search for an effective artifact requires utilizing available means to reach desired ends while satisfying laws in the problem environment.	The search for an effective artifact requires utilizing available means to reach desired ends while satisfying laws in the problem environment and constraints imposed by the design fitness landscape.
Communication of research	Design research must be presented effectively both to technology-oriented as well as management-oriented audiences.	Design research must be communicated to those communities most likely to supply the resources required for future design using communication channels appropriate to each community.

(iii) early stage design artifacts often leave challenging questions—such as scalability and relative benefits compared to alternative designs—largely unanswered. Rather than abandoning rigor altogether, the impact of the new model on this guideline could be stated as follows:

Impact to Guideline #5. DSR requires that the construction and evaluation of design artifacts be investigated employing a level of rigor appropriate to the nature and stage of design.

Guideline #6: Design as a Search Process

Guideline #6. The search for an effective artifact requires utilizing available means to reach desired ends while satisfying laws in the problem environment. There is little need to change the spirit of this guideline, which captures perfectly the process of search in a fitness landscape. A slight impact to the wording is desirable, since the fitness-utility model assumes we are searching for high fitness artifacts constrained by a design fitness landscape.

Impact to Guideline #6. The search for an effective artifact requires utilizing available means to reach desired ends while satisfying laws in the problem environment and constraints imposed by the design fitness landscape.

Guideline #7: Communication of Research

Guideline #7. Design research must be presented effectively both to technology-oriented as well as management-oriented audiences. This guideline once again illustrates the preference for late-stage design research in the original conception of DSR. Management-oriented audiences, in particular, are unlikely to be impressed by designs whose usefulness has not been demonstrated. The fitness-utility approach would take an entirely different perspective. Where the goal is to exert impact on the design fitness landscape, what makes sense is to target those communities most likely to initiate the next iteration of the design process through supplying resources, which would naturally include time, intellectual effort, facilities, and money.

Impact to Guideline #7. DSR must be communicated to those communities most likely to supply the resources required for future design using communication channels appropriate to each community.

Summary

A summary of the fitness-utility model impacts to the DSR guidelines proposed by Hevner et al. [2004] is presented in Table A.1.

ACKNOWLEDGMENTS

The ideas shared in this commentary have benefited greatly from the insightful comments provided by participants in seminars in which the fitness-utility model was presented. The authors also gratefully acknowledge the improvements suggested by the reviewers and the editors.

REFERENCES

- Agrawal, R., Imielinski, T., and Swami, A. 1993. Mining association rules between sets of items in large databases. In *Proceedings of the ACM SIGMOD Conference* (Washington DC, May). ACM, New York.
- Alexander, C. 1964. *Notes on the Synthesis of Form*, Harvard University Press, Cambridge, MA.
- Ariely, D. 2007. *Predictably Irrational: The Hidden Forces that Shape Our Decisions*, Harper Perennial, New York, NY.
- Ariely, D. 2010. *The Upside of Irrationality: The Unexpected Benefits of Defying Logic at Work and at Home*, Harper Collins, New York, NY.
- Burris, S. 2010. George Boole. E. N. Zalta, *The Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/archives/sum2010/entries/boole/>.
- Chen, H. 2011. Design science, grand challenges, and societal impacts. *ACM Trans. MIS* 2, 1.
- Christensen, C. 1997. *The Innovator's Dilemma*, Harper Collins, New York, NY.
- Davis, M. 1971. That's interesting!: Towards a phenomenology of sociology and a sociology of phenomenology. *Philo. Soc. Sci.* 1, 309–344.
- Dawkins, R. 1976. *The Selfish Gene*. Oxford University Press, Oxford, UK.
- DeLone, W. and McLean, E. 1992. Information systems success: The quest for the dependent variable. *Inf. Syst. Res.* 3, 1, 60–95.
- DeLone, W. and McLean, E. 2003. The DeLone and McLean model of information systems success: A ten-year update. *J. Manage. Inf. Syst.* 19, 4, 9–30.
- Dennett, D. C. 1990. Memes and the exploitation of imagination. *J. Aesthet. Art Critic.* 48, 2, 127–136.
- Fehr, E. and Fischbacher, U. 2003. The nature of human altruism. *Nature* 425, 785–791. DOI:10.1038/nature02043.
- Gandolfi, A., Gandolfi, A., and Barash, D. 2002. *Economics as an Evolutionary Science: From Utility to Fitness*. Transaction Publishers, New Brunswick, NJ.
- Gartner. 2011. Gartner says sales of mobile devices grew 5.6 percent in third quarter of 2011; smartphone sales increased 42 percent. <http://www.gartner.com/it/page.jsp?id=1848514>.
- Gill, R., Glazer, N., and Thernstrom, S. 1992. *Our Changing Population*. Prentice-Hall, Englewood Cliffs, NJ.
- Gill, T. G. 2008. A psychologically plausible goal-based utility function. *Inf. Sci.* 11, 227–252.

- Gill, T. G. 2010. *Informing Business: Research and Education on a Rugged Landscape*. Informing Science Press, Santa Rosa, CA.
- Gill, T. G. and Bhattacharjee, A. 2009. Whom are we informing? Issues and recommendations for MIS research from an informing science perspective. *MIS Quart.* 33, 2, 217–235.
- Gill, T. G. and Sincich, T. 2008. Statistical illusions in the rugged fitness landscape. *Inf. Sci.* 11, 197–226.
- Gladwell, M. 2000. *The Tipping Point*. Back Bay Books, New York.
- Gleick, J. 2011. *The Information*. Pantheon Books, New York.
- Gregor, S. and Hevner, A. 2013. Positioning and presenting design science research for maximum impact, *MIS Quart.* 37, 2, 337–355.
- Gregor, S. and Jones, D. 2007. The anatomy of a design theory. *J. AIS* 8, 5, Article 2, 312–335.
- Heath, C. and Heath, D. 2007. *Made to Stick*. Random House, New York.
- Heineman, G. 1998. Adaptation and software architecture. In *Proceedings of the 3rd International Workshop on Software Architecture (ISAW'98)*. 61–64.
- Hevner, A., March, S., Park, J., and Ram, S. 2004. Design science in information systems research. *MIS Quart.* 28, 1, 75–106.
- Hines, W. 1987. Evolutionary stable strategies: A review of basic theory. *Theoret. Pop. Biol.* 31, 2, 195–272.
- Hopper, M. 1990. Rattling SABRE—New ways to compete with information. In *Harvard Business Review: Revolution in Real Time: Managing Information Technology in the 1990s*. HBS Press, Boston. 113–127.
- Johnson, S. 2010. *Where Good Ideas Come From: The Natural History of Innovation*. Riverhead Books, New York.
- Kauffman, S. 1993. *The Origins of Order: Self Organization and Selection in Evolution*. New York, NY, Oxford University Press, Inc.
- Kernigan, B. and Ritchie, D. 1988. *The C Programming Language*. 2nd Ed. Prentice-Hall, Englewood Cliffs, N.J.
- Lakhani, K. and von Hippel, E. 2003. How open source software works: “Free” user-to-user assistance. *Res. Policy* 32. 923–943.
- Lakhani, K. and Wolf, R. 2003. Why hackers do what they do: Understanding motivation and effort in free/open source software projects. MIT Sloan Working Paper No. 4425-03. DOI:10.2139/ssrn.443040.
- March, S. and Smith, G. 1995. Design and natural science research on information technology. *Decision Support Syst.* 15, 251–266.
- Mills, H., Linger, R., and Hevner, A. 1986. *Principles of Information Systems Analysis and Design*, Academic Press, Inc.
- Mills, H., Dyer, M., and Linger, R. 1987. Cleanroom software engineering. *IEEE Soft.* 4, 5, 19–25.
- Mokyr, J. 2002. *The Gifts of Athena: Historical Origins of the Knowledge Economy*, Princeton University Press, Princeton, N.J.
- Morch, A. 1997. Three levels of end-user tailoring: Customization, integration and extension. In M. Kyng and L. Mathiassen Eds., *Computers and Design in Context*. MIT Press, Cambridge, MA.
- Niederman, F. and March, S. 2012. Design science and the accumulation of knowledge in the information systems discipline, *Manag. Info. Syst.* 3, 1.
- Nunamaker, J. and Briggs, R. 2011. Toward a broader vision for information systems, *Manag. Info. Syst.* 2, 4. Parc.com. 2012. About Parc. <http://www.parc.com/about/>.
- Purao, S., Rossi, M., and Bush, A. 2002. Towards an understanding of the use of problem and design spaces during object-oriented system development. *Inf. Org.* 12, 249–281
- Raymond E. and Landley, R. 2008. OSI position paper on the SCO-vs.-IBM complaint, Revision 2.8. <http://catb.org/esr/hackerlore/sco-vs-ibm.html>.
- Scott-Morton, M. 1967. Computer-driven visual display devices – Their impact on the management decision-making process. Unpublished PhD thesis, Harvard Business School.
- Sein, M., Henfredsson, O., Purao, S., Rossi, M., and Lindgren, R. 2011. Action design research. *MIS Quart.* 35, 1, 37–56.
- Simon, H. 1996. *The Sciences of the Artificial* 3rd Ed.. MIT Press, Cambridge, MA.
- Taleb, N. 2007. *The Black Swan: The Impact of the Highly Improbable*. Random House, New York.
- Taleb, N. 2012. *Antifragile: Things that Gain from Disorder*. Random House, New York.
- TOPS. 2011. Top 500 Statistics. <http://i.top500.org/stats>.
- Venkatesh, V. and Davis, F. 2000. A theoretical extension of the technology acceptance model: Four longitudinal field studies. *Manage. Sci.* 46, 2, 186–204.

- von Hippel, E. 2001. Innovation by user communities: Learning from open source software. *Sloan Manage. Rev.* Summer. 82–86.
- von Krogh, G., Spaeth, S., and Lakhani, K. 2003. Community, joining, and specialization in open source software innovation: A case study. *Res. Policy* 32, 1217–1241.
- Walls, J., Widemeyer, G., and El Sawy, O. 1992. Building an information system design theory for vigilant EIS. *Inf. Syst. Res.* 3, 1, 36–59.
- Williams., K., Chatterjee, S., and Rossi, M. 2010. Design of emerging digital services: A taxonomy. In A. Hevner and S. Chatterjee Eds., *Design Science in Information Systems Research*. Springer, New York, 235–253.
- Ziliak, S. and McCloskey, D. 2008. *The Cult of Statistical Significance*. University of Michigan Press, Ann Arbor.

Received May 2013; accepted May 2013