The Ubiquitous Button

Lars-Erik Janlert, Umeå University

Why are buttons so common in contemporary artifacts and yet so often a source of irritation and trouble? Could we, by reinstating the natural mode of operation with traditional mechanical systems, dispel our confusions and remedy our confirmation deficiencies? Probably not.

You are likely familiar with the following situation: Standing in front of a coffee machine, you are uncertain whether you should just press the “coffee” button momentarily, or if you are supposed to keep pressing it until your cup is filled. The second button option is like a buzzer: Coffee keeps pouring only as long as you press. The

first option, however, is at least as likely in a modern machine: You initiate a process that will be automatically completed, rather than continuously control the process yourself. Without disputing the benefits of automation, one may wonder how such modern-day uncertainties arose. Might the latest development of the buttons themselves have something to do with it?

Old-style buttons and switches often have just two functionally relevant states. They are used for making some function of the controlled artifact operative or non-operative, as in switching an electric light or a motor on or off. A flip switch may have up

Insights

- Deeper examination of the button can help us understand interactivity better.
- As complexity increases, there seems to be a shift in button design from controlling to expressing.
- Our basic expectations of symbols may conflict with promoting trust in our control devices through transparency of mechanism.

COVER STORY

and down as its two relevant states, a rotary switch a horizontal and a vertical state, and a push button a raised and a depressed state. These states are readily observable for the user. In many cases both states are also *persistent*; that is, they will not change without user intervention. In other cases, just one of the states is persistent; in a buzzer button, for example, the depressed state depends on the constant pressure of the user's finger.

Old-style buttons and switches with several functionally relevant states are also common: either a limited number

mapping between button states and artifact states may be unknown or inconstant. One example of the latter is when you have several switches that independently control the same lamp, as in a staircase or a corridor where you usually have switches at both ends. In somewhat more autonomous and complex artifacts, the mapping is incomplete: Much goes on inside that is not reflected in detail in the states of the buttons (or displays) on the outside. You start the artifact by pushing a button, and then it goes through different stages of processing.

transitory confirmations, of course. Then there is the often all-important visual confirmation: Provided there is light, you can see the orientation of a rotary switch, the position of a toggle switch, and the displacement of a push button—all persistent confirmations. Haptic, audial, and visual feedback of old-style buttons and switches have in common that they are basically straightforward consequences of the functional mechanism itself. Of course, the mechanism may still be designed to produce (or emphasize or attenuate) certain effects on the senses. If you don't want a loud snapping noise when you turn the light off in the children's bedroom, a quiet rocker switch is preferable to the traditional toggle switch.

Now, compare this with modern buttons, which tend to use a wider variety of technologies. Understandably, we have strived to get away from mechanical constructions, more expensive to make and generally more prone to wear and malfunction than electronically based or supported solutions. A modern button operated by bringing it into (or out of) a depressed state may have a very small offset compared with old-style mechanical push buttons, which means that “snapping” tactile and/or audial feedback becomes even more important for confirming that you have performed a proper button press. But rather than displacement, it may be that you need to apply a certain amount of pressure to put it into the alternate state, the particular threshold ranging from soft to hard. If the button is of the now common “touch buttons,” it should be enough to touch it lightly, or in some cases just bring a finger or hand close to it, to change its state. There are a variety of technological solutions, and for a user it may not be apparent exactly what will work: How much force is required? Can it be operated with gloves on? Can a pen be used? And so on.

Even more common today are the *virtual* or graphical buttons and switches, buttons presented as images on a graphical display, operated directly on it (as with a touchscreen) or with the help of some pointing device.

Typical for these modern buttons is that they have just a single, persistent, observable, and functionally relevant state in their basic design. At the same time, they are frequently not used in

of distinct states, as in the knobs of the burners in an old electric stove, or a continuous range of states, as in the knobs of the burners in a gas stove and many modern electric stoves. Again, the states are observable and typically also persistent. With a continuously variable range of states, some devices have just a single persistent state—the *quiescent* state—like, for example, a gas pedal.

Even when the states are observable, as they are in old-style buttons and switches, it is not always possible for a user to infer the state (or some state parameter) of the artifact it controls from the state of the button: The

Old-style buttons and switches typically confirm a successful change of their state in several modalities simultaneously. They provide haptic feedback; for example, a spring-loaded push button with two persistent states will cease to resist your finger if you manage to get it “hooked” in the depressed state, and you may also feel the depression of the button relative to the surrounding plate or an adjoining button. Audial feedback is often prominent and particularly reassuring, usually in combination with haptic feedback: You both hear and feel it snap into position. These are just momentary,

the manner of a buzzer button—to continuously control something—but rather in the manner of the traditional push button with two persistent states—to select some parameter or to start or stop a process. That means it is the *state change* of the button that matters rather than the state in itself—in practice often two consecutive state changes: from the quiescent state to the non-persistent state and back to the quiescent state again. The “click” is the typical modern button interaction.

CONFIRMATION AND INTENTION

When state change rather than state, per se, is in focus, even if the successful operation of the button is momentarily confirmed in a satisfactory manner, you might still wish for some lasting indication of what you have done. You may be able to rely on memory to know the state at a later point; of course, that would not be a proper state of *the button* as defined earlier (which requires being observable by the user), but you might still think of it as a state of the button that is *hidden*.

Obviously, hidden button states are not ideal: They add to the cognitive load of users and are vulnerable to operation errors, particularly in situations where there are several simultaneous users or where users come and go. The common remedy is the so-called *pilot light* that in effect can transform a button with a single persistent state into a button with two or even several persistent states; for example, push it once and the button lights up; push it again and the pilot light is switched off; alternatively, the light changes to a different color. Elevator buttons are early examples of buttons with pilot lights. When you push the button for the floor where you want to get off, the button lights up and stays lit to indicate the elevator will stop at that floor. Other users can see whether their floor button is already lit or if they need to choose their floor.

This brings up the more general question of the *depth of confirmation* of successful action. Confirmation that a button or switch has been successfully operated is shallow. It usually leaves open the possibility of malfunction deeper down the causal chain of events and may not definitively confirm the intended change of the artifact state. Ultimately, of course, it is not even the artifact state that proves success, but

the intended result of artifact use. For example, if a user is operating an electric pump for emptying a tank of water, then actually seeing, hearing, and feeling the water gushing out from the tank should be the ultimate confirmation of success. At a slightly shallower level, hearing the hum of the electric motor would confirm the pump is running, or rather the motor is running, since there could still be something wrong with the transmission from the motor to the pump or some mechanical defect in the pump rotor that meant water would still not be pumped, or some blockage in the outgoing piping. The indicator light of the button labeled “Emptying tank” will usually be only the most shallow of confirmations: that the button press was recognized.

The depth of confirmation has a parallel in the notion of the *depth of intention* of the user. This is an adaption of a term introduced by the Norwegian philosopher Arne Næss for discussing a kind of mismatch between how the creator and the interpreter of a linguistic expression understand it [1]. Here, depth of intention is applied in a more general sense to *actions* (of which spoken words or “speech acts” can be considered one variety). In pressing the pump button, did the user intend to empty the tank, or just to start the electric motor, or just to press the button? The last alternative is reasonable if it is his first day on the job and his foreman just told him, “Press that button.” When a regular user operates the button, his or her depth of intention might typically be on the level of emptying the tank, and the user might or might not be alert to other, deeper confirmations and signs of failure, aware of the possibility of something not going as intended. But, again, we cannot be sure of that; perhaps the user is just thoughtlessly following a list of operating instructions or is completely unconcerned about consequences.

Of course, it is possible to transform the pilot light into an indicator of state or process deeper down, for example by connecting it to a sensor that detects when the pump rotor is moving, or to a sensor that measures the flow of water in the outgoing pipe. It may not always be clear to a user what exactly an indicator or display associated with a button means or how deep the loop of confirmation goes; the border between “pure” pilot light and artifact state indicator is vague.

The more invisible, inaudible, untouchable, and inaccessible the “inside” and/or the “product” of the artifact, the more users need to rely on controls and instruments. The general trend has long been to rely on more complex machinery and thicker layers of insulation between it and the user and the world, and often the effects, the “product” of the artifact, is hidden or remote. Another factor is the sometimes important and unavoidable time delay between the action and the confirmation; the deeper down, the longer the delay. The proper operation of the button can be confirmed instantly, but the motor might take a few seconds before it has revved up enough for us to hear the sound, and it may take some additional time before the water has traveled through the piping and is actually flowing.

In view of the complexity of modern machinery, and in view of the increasing number of casual encounters with an increasing number of devices and appliances, maybe you do not really want confirmation of a particular artifact state or process state, into which you are likely, in many cases, to have little insight. Perhaps you would rather be satisfied with having successfully conveyed the intended direction or information to the artifact, and trust it will do what it is supposed to do, hopefully whatever is best to do.

It should be said the transition from directly effective switches to symbolic switches does not just serve to accommodate lazy users and obfuscate the relation between user and artifact. It definitely can have important advantages. One example is the on/off button of a CD player: Instead of simply and brutally cutting the power, the on/off button usually rather *signals* to the artifact the user wants to shut it off, which makes it possible to automatically perform certain routines, such as

retracting the tray and not leaving it sticking out precariously, before actually cutting the power.

This adds to the uncertainties of the modern button pusher: Am I actually doing this now and have full responsibility, or am I just making a humble request, and the machine can be counted on to save me from any serious blunders and hazards? Or is it something in between? In clicking the “Coffee” button, do you feel confident that the machine will stop pouring in time? Does it assume a certain cup size (and is this cup size the correct one?) or does it sense the level inside the cup? Might it pour coffee even if you fail to set down a cup? Such everyday button anxieties may appear petty and inconsequential, but certainly not all of our modern artifacts are as harmless as coffee machines.

In the end we cannot just blame the button. These uncertainties would not arise were it not for the variety and complexity of the choices and functions offered by our modern artifacts.

WHY ARE BUTTONS EVERYWHERE?

With functionally more capable artifacts follows greater internal artifact complexity, which calls for more controls (unless operations are automated). At present, many of those controls are buttons. Buttons seem to offer basically robust and uniform manual operation that does not require great dexterity or subtlety.

Buttons seem to suit the overtly digital character typical of many electric devices and appliances, as well as computer artifacts. Power and various functions are turned on and off. Could it be that the fundamental notion of closing and breaking an electric circuit created the framework for the design of electric appliances? We may have come to generally think in such terms, thus propagating and reinforcing their use.

This can be contrasted with current ideas about “always on” devices going to “sleep” and being on “standby” rather than being shut off (“dead” as opposed to “live”), and our increased expectations of nuanced function.

Knobs, controls that are operated by turning, invite analog semantics and more nuanced control. Their range of possible physical states is both wider (equaling the length of the knob’s periphery), more visible (orthogonal to the operator’s view), and can be read with higher precision, compared with ordinary push buttons.

When there is a lack of space, however, which is typical of many of the new and small but very capable and complex digital artifacts (the ratio of artifact size to artifact complexity has a general tendency to decrease in the general quest for improved artifacts that digital technology supports and encourages [2]), buttons have the important advantage of allowing a smaller footprint compared with knobs (as well as sliders, toggles, levers, etc.), before becoming too small to handle. You can pack buttons very tightly.

Their operation may be obvious, easy, and uniform—explaining much of their attraction—but there are so many of them and they look so much alike. Which one should you push? One consequence is that small digital devices studded with buttons tend to look similar whatever their purpose. Is it a wireless phone, a TV remote, or a pocket calculator?

We know we can make something happen by pushing a button, but not necessarily what that will be or which button to use. This is where the celebrated notion of *affordance* lets us down: Rightly designed affordance may help us to understand the operation, but not the function [3]. When there are few functions, there is still no big practical problem (what can you do with a door beyond open or close it?), but when there are many or the need for detailed control is great, it may become a major concern.

In view of their similarity (which of course is good from the operational point of view), it is understandable why buttons on digital devices practically always are accompanied by symbols. Alas, the symbols take up precious space, too, and take time to decipher.

But the operation of modern buttons is not as uniform as that of traditional buttons: The basic mode of operation is

The uncertainties of the modern button pusher: Am I actually doing this now and have full responsibility, or am I just making a humble request, and the machine can be counted on to save me from any serious blunders and hazards?

in the case of a piano), this is not very often the case with modern button use, where you are not so interested in continuously controlling a process as with setting and adjusting various parameters.

THE FORLORN IDEAL OF THE OBJECT SYMBOL

If we really wanted to turn the clock back and return to the clarity and transparency of the good old days when control devices were mechanical—for example, a door locked by a simple hook latch: You can check whether it is locked by seeing or feeling the position of the hook, and you lock and unlock it by operating the hook—then a good place to start would seem to be the notion of the *object symbol* introduced by Donald Norman and Edwin Hutchins in 1988 [4] and defined by Norman as follows:

“When the object in the artifact is both the means of control (for execution of actions) and also the representation of the object state (for evaluation), then we have the case of an object symbol” [5].

Object symbols “represent the natural and frequently occurring mode of operation with mechanical systems,” they point out, but they were lost in the transition to electronic systems, more by accident than by design. If that is so, maybe we could reinstate the natural order by making an effort to implement and use object symbols in the new technologies? Before we start on a project like that, however, we should examine the key properties of the object symbol. To do that I will start with considering what we generally expect of external symbols—the kind of symbols that are designated and used to represent various states of affairs as a substitute for and complement to dealing directly with the referred entities themselves; paradigmatic examples include spoken and written languages and mathematical and logical notations. Most of the symbols we find in a user interface, such as icons and menu items, have a similar character.

Among the basic expectations we have of symbols in this sense are (in no particular order):

- they are *lightweight*—easy, cheap, and safe to manipulate—compared with their referents;
- they *can be at a distance* from their referents;
- they *can be counterfactual*—they can symbolize states of affairs other than the

stretched in various dimensions to get additional control, mainly with regard to the amount of pressure, the duration of the press, and combinations and patterns of presses. The usual motive for this is that there would be too few buttons of the ordinary kind in the artifact to achieve the intended level of user control, either due to lack of space or because the designer wants a clean, nice-looking design and maybe hides away some rarely used functions or settings. In a modern camera, pressing lightly on the shutter button usually activates focusing and light metering; pressing harder takes the picture. Digital kitchen timers may suddenly increase the speed of counting up the minutes (for instance, from intervals of a minute to intervals of 10) when you keep holding the button. Wristwatches and mobile phones often have buttons that perform different functions if you press them for a longer time; also, hard-to-remember combinations of simultaneously pressed buttons take on special meanings. We are not very surprised anymore if a double-click achieves something different from a single click.

Such modifications undermine the

button's reputation of being simple to operate. Also, the degree of user control may suffer in the effort to limit or reduce the number of controls and prioritize the use of buttons over other control devices. Dimmers are sometimes now equipped with a single push button. As long as you press it, the intensity of the light will increase until it reaches the maximum; then it will go down at the same constant speed until reaching the minimum and continue this up-and-down cycle until you let go of the button. Like in a stock market rally, it is difficult to decide when the maximum has been achieved. It is also awkward to adjust the light intensity downward (or upward) when you are in the cycle's uphill (or downhill) phase.

The new, versatile buttons can be compared to keys on pianos and other musical instruments, where *how* you press a key normally is significant: The speed and force with which you hit it, how long you keep it depressed, and so on, can affect the result. But whereas the significant operational parameters of these keys are usually analogous and roughly proportional to their immediate effect on the result (the greater the force, the greater the amplitude, for example,

actual and present; and

- they are *perspicuous*—it is easy to extract from them what for the user’s purposes is the most relevant information—compared with their referents.

The first condition, being lightweight, ensures the relative practicability of symbol manipulation: If making a blueprint of a building involved handling symbols as heavy and unwieldy as the building elements themselves, we would be sorely disappointed.

The second condition, separability, allows symbols to reach beyond their own existence here and now, their spatiotemporal particularity. A real estate agent can bring along the blueprint to prospective buyers or include it in an advertisement. Symbols are regularly used to provide us with information about things we cannot presently perceive.

The third condition, counterfactual freedom, enables us to have a blueprint for a house that is yet to be built, a house that used to exist but is no more, a house that will in fact never be built, a planned modified version of an existing house, and so on. For planning purposes, it is obviously a must. For dynamically changing referents, counterfactual freedom is needed to keep a record of what went on before (you could also say that it falls under the separability condition: that symbols can be at a distance in time from their referents).

The fourth condition, perspicuity, is where the design of the symbols can make an important difference: first, in singling out which information about the referent will be the most relevant for the user’s purposes; second, in finding a form for the symbols that most conveniently and efficiently conveys this information to the user. The blueprint picks out information about floor space, floor plan, the size of walls, doors, windows, and their relative positions. It also conveys that information rather

efficiently, and it will for example be easier to see how the floor plan is organized or to find the length of a wall or the area of a room from the blueprint than from the physical building itself (if it exists).

In failing to satisfy one or more of these conditions, some of symbols’ capacity to serve as tools for thinking—our ability to use them to represent states of affairs other than the actual situation, our immediate present—is lost. In other words, symbols enable us to track remote events and states, to recall past events and states, to predict or plan future events and states, and indeed even to *fantasize* (disregard any connotation of impossibility and useless daydreaming). To uphold the division between thinking and acting essential to *Popperian* and *Gregorian* creatures in Daniel Dennett’s broad sketch of cognitive evolution, at least the first three conditions seem crucial [6].

Alas, object symbols fail on each of these points. Per definition they emphatically disappoint the third expectation: An object symbol is unable to represent a counterfactual state of affairs. And when objects represent themselves or a larger artifact of which they are a proper part, they cannot be at a distance from their referents since they cannot be separate from themselves, thus disappointing the second expectation. For similar reasons the first and fourth expectations are disappointed. They cannot be easier to manipulate than themselves, nor can they be more perspicuous than themselves.

This is not to say that symbols that fail to meet the basic expectations are useless. A radar screen in an air control center represents airplanes and their movements from a distance: It tracks what is going on in real time. It is useful because the events may be hard or impossible to follow by the naked eye at a distance (conditions two and four). Similarly, the click

rate of a Geiger counter symbolizes the present radioactivity. It is useful because you cannot see (or hear or smell) radioactivity (condition four).

But it really should come as no surprise that a physical object is a rather poor symbol of itself. Even a *duplicate* is a rather unsatisfactory symbol of the object. It is not lightweight, but it can at least be at a distance. It can be counterfactual to the extent that its lack of lightness does not prevent us from modifying it into various states that are different from that of its referent. It may satisfy the condition of perspicuity if the referent object is more difficult to get at—being remote and/or being in a situation that is in some sense critical. One dramatic example is NASA’s figuring out how to put together the improvised “mailbox” for removing carbon dioxide in the Apollo 13 incident (“Houston, we’ve had a problem”). Some states that—using “proper,” lightweight symbols—we may imagine, we may have no current means of achieving and are thus not able to symbolize, constraining our fantasizing ability.

Taking object symbols as an ideal for interaction with digital artifacts in general means that they may help the user in observing and tracking (besides enabling the user to perform operations on the artifact), but otherwise will not be particularly supportive of the user’s thinking. For example, they will not help the user to form and convey intentions without having to take direct action. It is true that physical actions sometimes serve a primarily cognitive role: You think by doing, by manipulating; this is what David Kirsh and Paul Maglio refer to as *epistemic actions* [7]. But, again, the use of epistemic actions is limited by the lightweightness condition: The cost in terms of effort spent and potential damage caused can in many cases not be sufficiently compensated by subsequent remedying actions, unlike in a “normal” symbol system satisfying all the basic requirements. With an object symbol, you “just do it”—there is no holding back, no time for reflection, hesitation, assessment, and possible abortion. You do not have what Hegel believed was the crucial distinction between humans and animals: the ability to dissociate thought and action—that is, to resist an impulse—and before you have that ability, there is in effect no distinction between thought and action. It is an indivisible whole, and thus *no thinking* is

Symbols enable us to track remote events and states, to recall past events and states, to predict or plan future events and states, and indeed even to *fantasize*.

taking place, properly speaking [8].

Let us try a simple, low-tech, and cognitively unsophisticated example like the knobs of an electric stove (surely we have no high expectations of stoves as props for thinking). Imagine that the knobs of the stove are object symbols: Not only can the user control the heat by turning the knob, but the current temperature is also simultaneously indicated by the current angle of rotation of the knob. Two practical problems immediately present themselves, as an effect of violating the first condition: If the stove has ordinary electric heaters, the logic of object symbols will require the user to apply torque to the knob for as long as it takes the stove to reach the desired temperature—not very convenient! And if the symbol really works both ways, how does the user express desired artifact states except by constantly working the controls? What, for example, stops the stove from getting cooler, slowly allowing its knobs to turn in the opposite direction to indicate lower and lower temperature? In many ways it seems much easier to make interfaces to virtual worlds than to the real world, where you cannot adjust the physics to suit the desired logic of an object-symbol interface.

Norman's definition refers to "the object in the artifact," and the concept is illustrated with the use of physical levers. We can think of, for example, a lever that controls a chimney valve: The lever could simply be a handle-shaped extension welded onto the valve itself. If the valve is inconveniently high up in the chimney, however, the lever might rather be connected to the valve by a chain or a string. The lever would still satisfy the definition, given that we understand it to be an object "in" the whole chimney-valve artifact. But a string might conceivably break, get stuck, slip, or come loose. There may be circumstances under which the lever would cease to function as a proper object symbol (possibly without the user becoming aware).

It appears that Norman and Hutchins want to contrast the self-evidence, reliability, and equating of action and effect of a mechanical linkage to the inscrutability, flimsiness, and conceptual and perceptual separation of action and effect of an electronic

linkage. However, from the above example it is not difficult to see that it is a matter of degree how robust the linkage is and appears to the user, and that there is no clear boundary between object symbol and non-object symbol in that sense. What matters is how far the user *trusts* the linkage. If, for example, earlier pilot lights and other light signals were deemed unreliable because of the relatively short and variable lifespan of traditional light bulbs, new LED-based technology has considerably improved trust. Now if a pilot light is not lit, it is much less likely because it is broken. Generally, the relative costs of false positives and false negatives will also vary depending on the application, affecting the user's degree of trust.

Regarding the *transparency* of the linkage of symbol and referent and the perspicuity of the relation between action and feedback, the tight coupling of the object symbol, basically the identity relation, is in principle unattainable in a complex artifact. While trust has mainly to do with reliability, transparency has more to do with complexity. The ideal of object symbols was formed with traditional low-complexity artifacts as paradigm. Digital artifacts are complex, and their complexity tends to increase rather than decrease. On the positive side, reliability can be improved and maintained even under these circumstances (the *perceived* reliability may be another matter). On the negative side, high complexity inevitably bars transparency in the direct sense of object symbols. While it may be possible to achieve a strong superficial association between action and feedback, the exact nature and meaning (including the limits and exceptions) of that connection will inevitably turn more vague and elusive as artifacts grow more complex.

FROM CONTROLLING TO EXPRESSING

The button and the object symbol can be seen as early manifestations of the *control* thought style in interface design [2]. As artifact complexity increases, what can be called the *expressive-impressive* thought style [2] is becoming more common and the control thought style reserved for less complex artifacts and for complex artifacts considered to

require less detailed control by a user. This is probably what we see reflected in the changed role of low-level devices like the button in modern artifacts, illustrated earlier by the shift from using the button to control coffee flow to expressing a specific request for a cup of coffee.

We may have reason to be concerned about the depth of intention and depth of confirmation. Whereas the expressive-impressive thought style would seem to lead to increased depths of confirmation, the depth of intention may still differ considerably between different operators of an artifact. Clearly, it is impossible to intend all the ramifications of our actions, but we may well be in a process of trying to or being compelled to extend our scope of intention beyond immediate outcomes. Consider, for example, how environmental concerns are looming larger and larger in everyday actions and decisions.

ENDNOTES

1. Naess, A. *Interpretation and Preciseness: A contribution to the theory of communication*. Det Norske Videnskapsakademii Oslo, Oslo, 1953.
2. Janlert, L.-E. and Stolterman, E. Faceless interaction—a conceptual examination of the notion of interface: past, present and future (manuscript submitted for publication).
3. Djajadiningrat, T. et al. Tangible products: Redressing the balance between appearance and action. *Pers. Ubiquit. Comput.* 8 (2004), 294–309.
4. Norman, D. and Hutchins, E. Computation via direct manipulation (Final Report: ONR Contract N0001485C0133). Institute for Cognitive Science, University of California, San Diego, La Jolla, CA, 1988.
5. Norman, D. Cognitive artifacts (Report 9002). Department of Cognitive Science, University of California, San Diego, La Jolla, CA, 1990.
6. Dennett, D. *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. Simon & Schuster, New York, 1995.
7. Kirsh, D. and Maglio, P. On distinguishing epistemic from pragmatic action. *Cognitive Science* 18 (1994), 513–549.
8. Hegel, G.F. *Vorlesungen über die Philosophie der Geschichte*, 1837.

✉ **Lars-Erik Janlert** is professor of computing science and cognitive science at Umeå University. His research ranges from knowledge representation to interaction design, new media, and philosophy of computing and information.
→ lej@cs.umu.se