Essay-4 human error

Precis: While the human factor initially was about the shortcomings of human performance compared to the perfection of machines (epitomised by the Fitts' List) it soon became associated with one specific feature, namely the unpredictability and unreliability, that turned into the concern for human error, in the sense that variable and unreliable performance of humans quickly became the favourite candidate for causes of UUOs. This essay, will present a broader view of human factors. The particular human error view is the topic of a following essay.

Humans as presumably sentient beings feel anxious whenever something unexpected happens, particularly if it also brings unacceptable outcomes. One thousand years ago the Muslim scholar Ibn Hazm (994-1064) wrote that "The chief motive of all human actions is the desire to avoid anxiety." In the 19th century the great German philosopher (Friedrich Nietzsche (1844-1900) elaborated "To trace something unfamiliar back to something familiar is at once a relief, a comfort and a satisfaction, while it also produces a feeling of power. The unfamiliar involves danger, anxiety and care -the fundamental instinct is to get rid of these painful circumstances. **First principle - any explanation is better than none at all.**" In their attempts to find a socially acceptable cause people usually ,and involuntarily, make an efficiency-thoroughness trade-off and settle for the cause that represents the default explanation in the current age of safety, cf. Table 1. This is why the preferred cause in the second age of safety became human error. It was first of all convenient, it was also immediately understood and accepted by others. And it finally corroborated with the experience everyone have that they sometimes act in a way that leads to UUOs.

Professors Andrew Hale and Jan Hovden (Hale & Hovden,1998) defined the second age of safety thinking as the age of human factors(Figure 1 and Table 1).

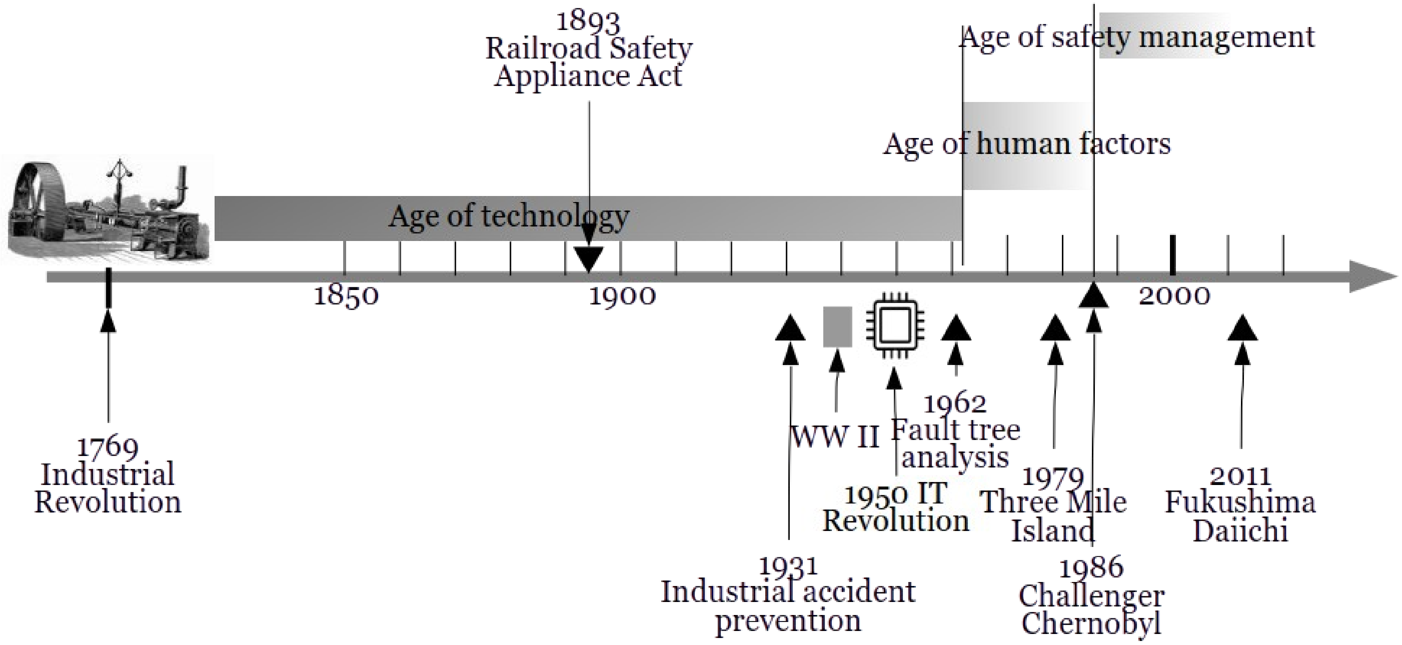

Figure 1: The three ages of safety based on (Hale & Hovden, 1998).

Table 1: Characteristics of the three ages of safety (from Hale &Hovden, 1998).

Age of safety thinking	Typical (default) cause	Typical (default) response	Characteristic mode of causality	Initiating event
I. Technology	Failure of technology	Replace, repair and improve	Linear, Monocausal	Watt's steam engine 1769
II. Human factors	Human factor "human error"	Blame, train, design, automate	Linear, Monocausal	Multiple accidents during World War II
III. Age of safety management	Organisational failure	Audits,standardisation regulation	Linear, Polycausal	Challenger & Chernobyl (1986)

The original version of the domino model (Heinrich, 1959) included Fault of Person as the second domino piece, defined as in Figure 2 and Table 2:

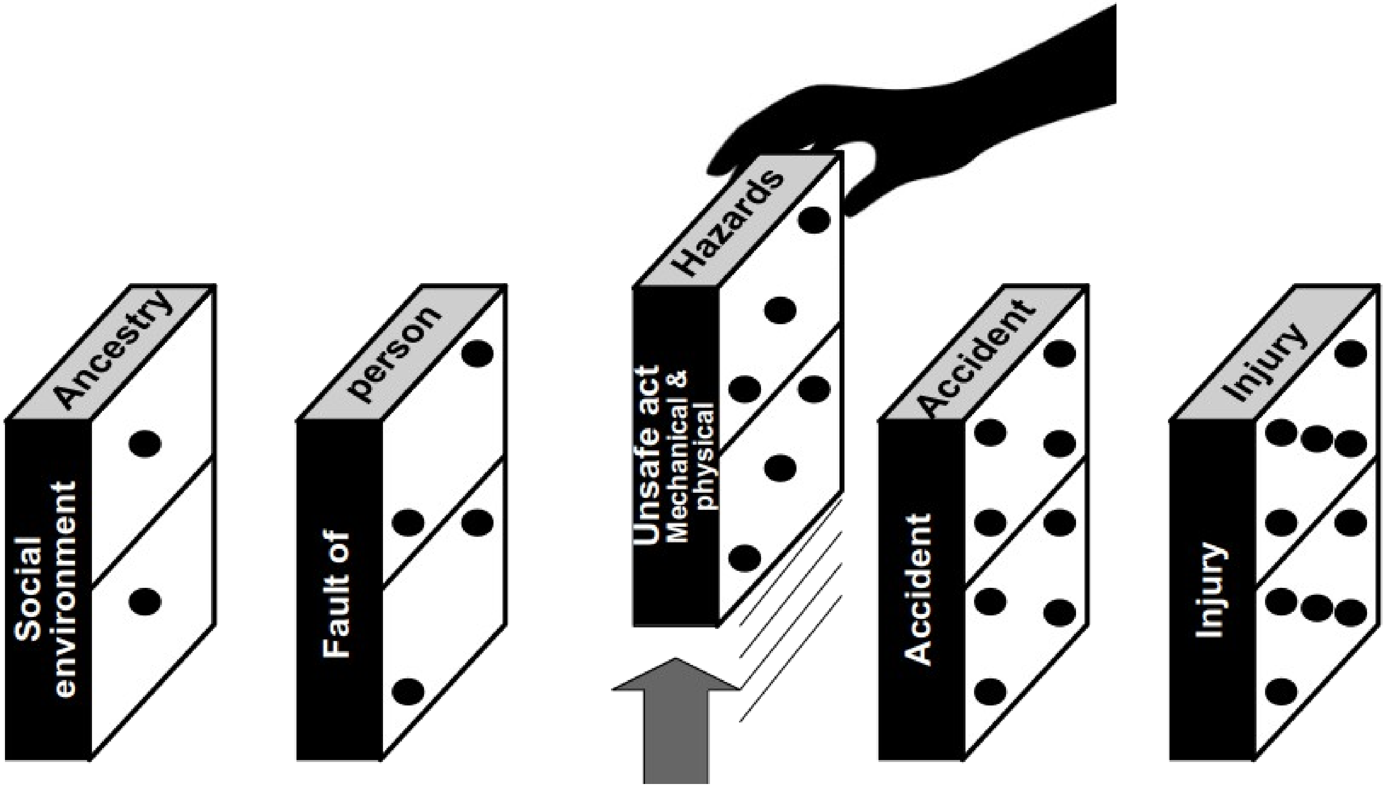

Figure 2: The five factors in the accident sequence (Heinrich, 1959)

The reasoning principle of the model of the accident sequence is simple. Indeed the very name “accident sequence reveals all. Each domino piece is the assumed cause of what happens next. What is interesting, but not made clear is whether the first domino piece ”Ancestry and social environment” truly is the first; the Pre-socratic philosopher Leucippus of Miletus (c. 480–c. 420 BC), is credited with the statement that ‘Nothing happens in vain, but everything from reason and of necessity.’ So neither ancestry nor the social environment can be assumed to have appeared out of the blue. Indeed, the undesirable traits of character, are directly said to be “passed along” which means they must have arisen somehow at an earlier, but undefined stage. The same uncertainty applies to the last piece “injury” Heinrich himself had introduced the concept of hidden costs, which are the hard to detect consequences that follow an injury. So the accident sequence does not really begin by the first domino piece, nor end by the fifth and last. It is indeed possible to continue in both directions for as long as one likes, or as far as the imagination allows.

Although Heinrich’s use of “fault of person” and “unsafe performance of persons” obviously refer to humans, it cannot be considered the beginning of human factors. The contribution of humans at a place of work is unquestionably essential, without human performance nothing can be done, automated machines can do many things, but they cannot yet themselves decide what they should do. They have to be instructed and motivated by humans. Until now we still use machines for our purpose, and woe the day when machines begin to use us for their purposes. Heinrich’s terminology does not correspond fully to the concept of the human factor as it is used today, it merely introduced the harmful suggestion that something done by a person may be part of the explanation of an accident.

Ever since human factors were adopted by safety studies as described in the third essay, there has been a bias in the preference for using human error to at least partly explain why work and actions that usually result in EAOs, occasionally, but very rarely produce UUOs. The bias is the assumption that there exists something called human error about which meaningful questions can be asked - and answered! human error thereby gets the status of a concrete phenomenon similar to, for instance, decision making. It is, however, obvious that human error does not refer to something observable, in the way that decision making does. Decision making represents a psychological function and denotes the activity of making a decision as well as the assumed cognitive functions behind that (and there are many different theories of decision making. There are many theories of decision making, ranging from normative decision theoey (*homo economicus*) to various descriptive theories, such as Simon's satisficing, (Simon, 19xx)Lindblom's muddling throug (Lindblom, 19xx) and Klein's naturalistic decision making (Klein, 19xx) as well as mixed approaches (Tversjy, 1972). (In the first sense it is observable in an everyday meaning of the term. In the latter it has to be inferred from observed behaviour, but the inferences need not be very elaborate, and there is no shortage of theories of decision making, ranging from formal theories of rational decisions to more descriptive accounts such as satisficing (Simon, 18xx) "muddling through (Lindblom, 19xx) and naturalistic decision making (Klein 2xxx). As well as mixed normative and descriptive approaches (Tversky, 1972).

Compared to decision making, there are practically no theories of human error, and there obviously cannot be a normative theory about the ideal way to make a human error.There are, of course quite a few descriptive theories Best known of these is, without a doubt James Reason's terminological distinction between mistakes, defined as deficiencies or failures in the judgmental and/or inferential processes in the selection of an objective or in the specification of the means to achieve it, irrespective of whether or not the actions directed by this decision-scheme run according to plan. (Reason, 1990) additionally distinguished between slips and lapses, and violations as illustrated by Figure 3. These terminological proposals do, however, not constitue an articulated theory of human error as such. result from a failure in the execution and/or storage stage of an action sequence, regardless of whether or not the plan which guided them was adequate to achieve its objective.where the plan or intention to act is inadequate, even if the action goes exactly as intended and performance deficiencies, where the plan is satisfactory, but the execution fails. With further distinctins between slips, lapses, and violations. These terminological proposals do, however, not constitue an articulated theory of human error. Reason's terminology was later combined with the Shill-bases, Rul-based, Knowledge-based (SRK) framework developed by Rasmussen (19xx), resulting in the following typology

- Skill-based (SB): slips and lapses, usually errors of inattention or misplaced attention
- Rule-based (RB): mistakes, usually a result of picking an inappropriate rule caused by misconstrued view of state, over-zealous pattern matching, frequency gambling, deficient rules
- Knowledge-based (KB): mistakes, due to incomplete/inaccurate understanding of system, confirmation bias, overconfidence, cognitive strain.

The SRK framework, is however itself without any theorerical or psychological foundation, so although the combination of the Srk framework with Reason's error types was called a Generic Error modelling System (GEMS), it is actually not a (descriptive) model of human error, and cannot replace a theory either, impressively looking flowcharts notwithstanding.

The distinction between error phenotypes and possible error genotypes is more systematic, yet not a theory as such. The phenotypes can be derived from the experience from FMEA and FMECA, and therefore has an undisputed empirical basis, even though it is still far from being a theory, let alone a model of human error. somewhat more systematic treatment of human error, yet not a theory as such. (Reason, 1990) deacribed, the acronym GEMS notwithstandingwas neither a theory or model of human error, but essebtially a systematic classfication system Figure 3 does not in any way explain how human errors happen, but just how they can or should be categorised. Violations hav e recent yearws attracted renewed interest, but now go by the name of non-compliance.

***Positive and negative connotations*.**

Most of the terminology used to describe the occurrence of human error unsurprisingly have negative connotations, cf. Heinrich. Terms such as violations, deviations, and non-compliance are all negatively loaded. This is in good agreement with the safety legacy. Human error is mainly used to explain UUOs, and since UUOs have a negative valence, it is only to be expected that the same goes for their putative causes. This is indeed one of the symmetries or biases in Anshin.

Multiple meanings of human error

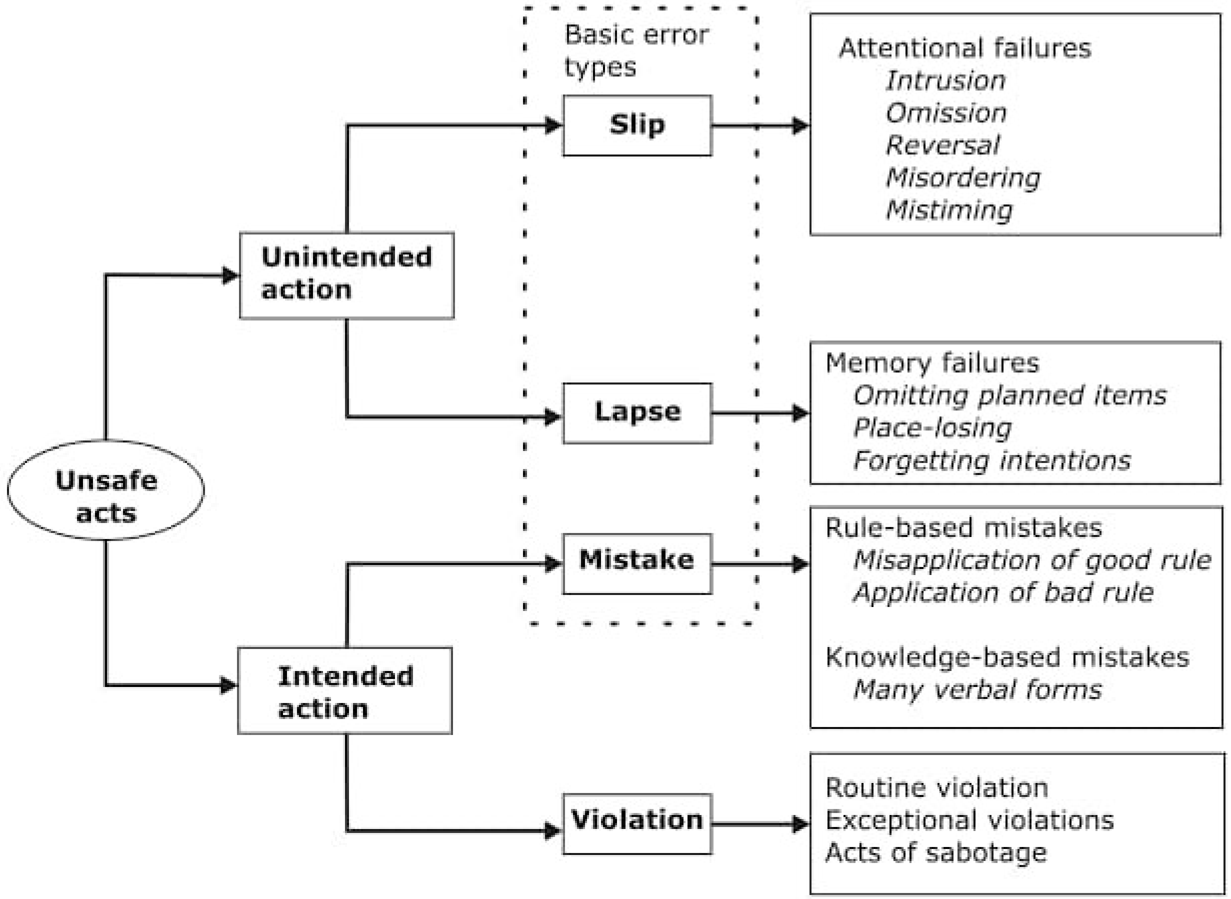

Figure 3: (Reason, 1990) error types

The Generic Error-Modelling System (GEMS) integrates, within the same framework, the different error mechanisms (slips, lapses and mistakes) and the three levels of performance (Skill, Rule, Knowledge (SRK)). The integration of these two dimensions allows us to:

Errors can result from operating at wrong level:

humans are reluctant to move from a RB to KB level even if rules aren't working

But changing to work as KB instead of b is not an intentional action, any more than swithing for Kahneman's ystem-1 to System-2

Satisficing and muddling through are hsabits, but not choices. Done but not chosen, just like ETTO.

Mistakes: (e.g., following a bad checklist)These are broken down into:Slips: Observable, unintended actions (e.g., flipping the wrong switch).Lapses: Internal memory failures (e.g., forgetting to complete a step).Violations: Intentional deviations from rules, procedures, or safety

In both cases decision making is regarded as a function. As a rule I will assume that functions, such as decision making, can be detected or seen in a straightforward way by an observer - although they may not be directly observable from a more stringent philosophical point of view.

"Human error" is, however, not a function, but a cause (or, to be precise: an assumed cause). We can use the term in a functional sense, as when we say that someone is making a mistake or an error. But in neither case is the "human error" an activity, nor the result of an intention. It is simply a contradiction of any reasonable definition to say that a person can make an error intentionally. It would therefore also be

meaningless to call it a function. It may be argued that "human error" characterises the outcome of an action rather than the cause. But classifying an outcome as a "human error" is a misuse of the terminology. What is meant is rather that the outcome was caused by a "human error". Neither can the "human error" be the activity that leads to the outcome. We cannot classify an activity as being a "human error", since that would assume that making the error was intentional. As that is not the case, it will be more correct to classify the activity as a failure to accomplish the intended outcome.

As a cause, "human error" must be inferred from observations rather than observed directly, just as causality, according to David Hume is metaphysical rather than physical. Other examples of such non-observables are "goal", "memory", etc. Consequently we must specify the observations from which the inferences are made. These observations will normally be about a particular performance or segment of a performance. We may observe the performance of a person, classify it as being incorrect, and determine the cause to be a "human error". But in no case can we observe the "human error" directly.

Since "human error" is inferred, it is not necessarily unique. Another way of saying this is by noting that "human error" is just one explanation out of several possible for an observed performance (or more precisely, a part of an actual performance description, cf. Hollnagel , Pedersen & Rasmussen, 1981). The analysis is normally carried just far enough to find a plausible explanation. If an explanation, which refers to the technological parts of the system, cannot be found the category "human error" is used instead (cf. Rasmussen, 1981) and also (Hale & Hovden, 1998, and also Chapter 9). It is only when the analysis is carried beyond this point that we may realise that an explanation in terms of "human error" is insufficient.

Since my major source of experience is operators in control of a complex process (for instance, a nuclear power plant), I will assume that the focus is on what conventionally was called human-machine systems, now referred to as socio-technical systems, or joint cognitive systems (Hollnagel & Woods, 2xxx) we deal with is a Man-Machine System (MMS) that functions as a process control system. By an MMS I mean a system that is composed of one or more operators and one or more machines (usually computers) that are designed to support the control of the process. A particular example of this approach is the Cognitive Systems Engineering (cf. Hollnagel & Woods, 1983; Hollnagel & Woods, 2005; Woods & Hollnagel, 2006).

When the performance at work is observed (and evaluated) a mismatch may be detected between the actual and the intended system states, or between the achieved results and the goal. The detection of this mismatch presumes that a description of the intended system state (or goal) is available. The mismatch is assumed not to be random, and therefore to have an identifiable cause. Finding the cause amounts to accounting for the observed difference betaween expected and actual performance. If faults in the technological parts of the system cannot be found, the solution is generality to assign the variance (or residual variance) to people at the sharp end, hence to use "human error" as an explanation.

The detection of this mismatch is thus the observational basis for inferring the existence of a "human error". It should be noted that if there is no observed mismatch, then there is little or no motivation to look for a cause. Variations in performance do not necessarily lead to undesired outcomes, hence mismatches. In fact they are usually an indispensable part of work that goes well. They may, for instance, be detected and corrected, at an early stage or the environment can be sufficiently friendly and forgiving. There will consequently be many cases of performance variability that remain unnoticed. From the point of view of a theory of "human error" they are, however, just as important as the cases where a mismatch is observed, and should therefore be accounted for by it. The opposite of what reason claimed.

The crucial point thus is a mismatch between intended and actual outcomes of an action. If the functional analysis is carried one step further, it will show that the cause of the mismatch can be located either in the selection of the goal for the action (the formation of the intention) or in the execution of the actions designed to achieve that goal. One may even distinguish between a larger number of categories by using one of the models of human decision making, or a theory of human performance. But this actually reduces the need for a specific theory of "human error", since the observed discrepancies instead can be explained by referring to, for instance, a performance theory. That may furthermore have the virtue of focusing on the situation and context in which the MMS must function, and the interaction between its inherent characteristics and the environmental constraints.

Observed mismatches in performance are always caused, in the sense that they can be analysed until the necessary and sufficient conditions, at least according to the underlying model or theory for their occurrence have been established. In some cases they may be classified as random, but that just means that the natural performance variability is sufficient to account for the mismatch, hence that no definite "other" cause has been identified or is require by the explanatory framework in use.

Since errors are not intentional, and since we do not need a particular theory of errors, it is meaningless to talk about mechanisms that produce errors. If the insatiable need for causes compel us to look for causes or "mechanisms", we should rather be concerned with whatever may be used to explain everyday work that goes well. If we are going to use the term psychological mechanisms at all, we should refer to possible "faults" in the functioning of how our minds usually function rather than to "error producing mechanisms". We must not forget that in a theory of action, the very same mechanisms must also account for the correct performance which is the rule rather than the exception.

Even though we do not yet have a sustainable "theory of error", it makes sense to distinguish between endogenous and exogenous causes for observed performance mismatches. There are certainly cases where the mismatch can be attributed to external causes, such as a bad interface design, interruptions, lack of operational support, incorrect information, time pressure etc. Similarly, there are cases where the causes are of an internal rather than external nature. In most cases the cause is,however, best described as a mixture. Stress, for instance, is often caused by (situationally) unreasonable demands to the operator. And deficiencies in the design of the interface may often be compensated by the adaptability of the operator (cf. Taylor & Garvey, 1959). Replacing a "theory of error" with a general theory of human action increases rather than reduces the importance of both internal and external causes, and emphasises the need to carry the analysis as far as possible.

To conclude, a theory of error must be a theory of the interaction between human performance variability and the situational constraints.

12.2: Terminology and taxonomies

Other dimensions can easily be found, and several complete taxonomies are available. One good example is the CSNI taxonomy (cf. Rasmussen et at., 1981), which is an attempt to characterise the situation where a mismatch occurs, rather than the "human errors". In this taxonomy "human error" is simply one of the many possible causes for a reported incident. Other taxonomies can rather easily be suggested once a proper theoretical background has been established. The choice of a taxonomy must depend on the purpose of the description, e.g., whether one wants to reduce the frequency of reported incidents, or improve the understanding of human decision making.

Before the key terms are defined, it is important to make sure that they are properly selected. One can, of course, make a pot pourri of terms that are normally used to characterise situations where humans make mistakes or errors, and then define them, e.g., by using a recognised dictionary. But if the definitions are to serve a purpose, it is essential that they have a common basis, for instance a theory. By the same rationale it also is essential that the terms have a common basis.

To repeat what has been said above, I believe we should attempt to come forward with a theory for "Human Action" rather than "human error", and that this should be used for selecting and defining the key terms. Such a theory is not yet available, but I will nevertheless attempt to give a definition of some of the terms the organisers have listed, using intentional action as a basis.

Error: Undefined. This term should be substituted by "action" or "activity".

Mistake: Incorrect selection of goal state; incorrect goal decision.

Fault: Incorrect selection of action to reach a goal, or incorrect execution of that action.

Slip: Unintentional substitution of a correct performance segment (action) with an incorrect one.

Accident: External disturbance of intended performance.

Cause: Accepted explanation for some performance characteristic, normally a performance mismatch.

Reason: Subjective explanation of goal state or intention.

Responsibility: Attribution of cause for the mismatch to a specific part of the MMS.

Assuming that we aim to establish a theory of human action rather than "human error", the predictions must be about actions. They must specifically be about the variability of human action that leads to mismatches. We can, of course, make a count of the instances where an operator makes a mistake, i.e., where the cause of the mismatch is attributed to a "human error". But that does not mean that it is sensible to attempt to assess the reliability of the operator, even if we refrain from considering the operator in mechanistic terms. Making such a count furthermore assumes that a meaningful measurement has been defined.

12.3: Predictions of "error"

It is obvious for anyone who has worked with the reliability aspect of the human operator, that the occurrence and frequency of human errors depend more on the interaction with the environment than on any stable inherent characteristic of people. Similarly, quantitative measures, such as error rates, will therefore be inadequate and even misleading. Instead we need detailed and consistent descriptions of the conditions where mismatches occur. These qualitative descriptions may eventually be used as a basis for more straightforward measurements.

With regard to the specific questions relating to prediction, it will at our present state of knowledge only be the frequency of mismatches and typical causes that can be predicted. We know from experimental psychology, particularly the studies of attention and performance, that there are important regularities, as diurnal variations, situation dependencies, etc. Even though most data come from simplified laboratory situations, there is no reason to assume that they cannot be applied to realistic work situations. This has been confirmed, for instance, by studies of shift-work. It is also highly plausible that there are significant individual differences in “error proneness”.

To summarise, making predictions requires an adequate definition of what the predictions are about. Unless frequencies and probabilities are sufficient, one must have a theory, or at least a set of good hypotheses, in order to make the predictions. It is furthermore logical that predictions cannot be about causes, unless we assume a strictly deterministic world. Consequently, the predictions must be about outcomes, i.e., observed mismatches, and possibly the actions leading to them. In the sense that “human errors” are causes, we can therefore not make predictions of human errors.

From a practical point of view the most important question is how mismatches can be prevented. One clue to this is found in the cases where mismatches do not occur, either because they are detected and corrected by the operator, or because the system is sufficiently forgiving. It would be reasonable to look further into these possibilities for preventing mismatches, hence reducing “human error”

There are probably very many ways in which a place of work can be designed to facilitate the detection and correction of unintended performance variability. A good working theory of human action is invaluable in this respect, since it makes it possible to indicate more precisely when and how interventions to change the course of action can be made. It is probably better to design for general detection and correction rather than for specific prevention. The experience from all types of process control clearly shows that Murphy’s law cannot be beaten.

However, even if the best of systems has been designed, there will remain a basic variability in human performance that will lead to mismatches when the circumstances are right (or wrong, rather). If the operator was turned into an automaton (or even replaced by one), we might produce an error-free system, provided the degree of complexity was sufficiently low. But that is precisely the crux of the matter. The mismatches may occur not just because of mistakes made by the operator during operation, but also because of mistakes made by the designer during earlier phases. These mistakes would not be contained unless a theory of human action was applied to literally every aspect of the system.

It is not enough to go behind human error (HE) (as advised by Woods et al.(2010) because that still tacitly acknowledges that “human error” is a legitimate concept and phenomenon, otherwise one could not go behind it. It is instead necessary to leave the idea of human error completely and go beyond it. People have been obsessed with “human error” since the time of Cicero who in the year 43 BCE in a speech against Marcus Antonius uttered the famous words "errare humanum est", but errare means both to stray and to err which in this context is quite ironic, because it is staying on the path of simple linear reasoning, rather than straying from it that is the “cause of the error” (yet to argue so unfortunately remains within the paradigm). “Human error” as a result of simplified linear reasoning is, perhaps, the only kind of “human error” that is real. By continuing to do so we also disregard the second part of Cicero's famous statement, which was "nullius nisi insipienestis in errore perseverare" (Sometimes also rendered as "sed in errore perseverare turpe est"), meaning "but only a fool perseveres."(The way out of this is the solution on the other hand fairly simple as explained in this short poem by the Danish polymath Piet Hein: "The road to wisdom? Well, it's plain and simple to express: err and err and err again, but less and less and less."

It is possible to argue that it is being seduced or sucked in by oversimplified cause-effect thinking(unfortunately an inseparable part of most investigation methods) that is wrong – without that “human error” would not exist! So it is staying on the path rather than straying from it that is the “error”. The worst example is probably the attempt of St. Thomas of Aquinas to prove the existence of God by reasoning backward to the first cause of why we are here in the world. It clearly shows the drawbacks of thinking or reasoning within a constrained paradigm. Staying on the path, hence continuing what is wrong,whether it is going backward or forward is precisely what Cicero warned about, but that part has rather conveniently been forgotten.

12.3: causality RIP

12.4: Chapter 12: Commentary

12.5: Chapter 12: References

Hollnagel, E., Pedersen, O. M. & Rasmussen, J. (1981). Notes on human performance analysis (Risø-M-2285). Risø National Laboratory: Roskilde, Denmark.

Hollnagel, E. & Woods, D. D. (1983). Cognitive systems engineering. New wine in new bottles. International Journal of Man-Machine Studies, 18, 583-600.

Hollnagel, E. & Woods, D. D. (2005). Joint cognitive systems: Foundations of cognitive systems engineering. Boca Raton, FL: CRC Press

Rasmussen, J. (1981). human errors. A taxonomy for describing human malfunctioning in industrial installations (Risø-M-2304). Roskilde, Denmark: Risø National Laboratory.

Rasmussen, J.,et al. (1981). Classification system for reporting events involving human malfunctions (Risø-M-2240), SINDOC(81) 14. Roskilde, Denmark: Risø National Laboratory.

Taylor, F V. & Garvey, W. D. (1959). The limitations of a procrustean approach to the optimisation of man-machine systems. Ergonomics, 2, 187-194.

Woods, D. D. & Hollnagel, E. (2006). Joint cognitive systems: Patterns in cognitive systems engineering. Boca Raton, FL: CRC Press.

Table 2: Details from the accident sequence (Heinrich, 1959).

Name of domino piece	Explanation
1. Ancestry and social environment.	Recklessness, stubbornness, avariciousness, and other undesirable traits of character may be passed along through inheritance. Environment may develop undesirable traits of character or may interfere with education. Both inheritance and environment cause faults of person.
2. Fault of person.	Inherited or acquired faults of person; such as recklessness, violent temper, nervousness, excitability, inconsiderateness, ignorance of safe practice, etc., constitute proximate reasons for committing unsafe acts or for the existence of mechanical or physical hazards.
3. Unsafe act and/or mechanical or physical hazard.	Unsafe performance of persons, such as standing under suspended loads, starting machinery without warning, horseplay, and removal of safeguards; and mechanical or physical hazards, such as unguarded gears, unguarded point of operation, absence of rail guards, and insufficient light, result directly in accidents.

Yet despite the widespread concern for and ditto importance of what humans do at a place of work, there are no simple definitions of what human factors are. There are many accounts available of human factors as a discipline as an area of R&D, but no definitive authoritative account, although (Meister & Enderwick, 2001) is a good place to begin. Few, however, realise that the very term human **factors**, endorses the view that a system can be understood by decomposition, by breaking it into its constituent parts, following the millenia old tradition of atomism, that is central to the safety legacy. Doing that raises the question of what is meant by the term system.

What is a system?

The traditional and widely-shared view of humans as parts or components of work systems, applying a physical mechanical analogy that goes back at least to Julien Offray de la Mettrie (1996) whose treatise “Man a Machine” was published originally in 1747 not forgetting that the Golem (Wiener, 1956) appeared already in the Bible, according to this way of thinking humans are intricate mechanisms or machines, but machines nonetheless hence equal parts or components of the larger work system where they can be treated as if they were mechanical parts,This was recognised when (Taylor & Garvey, 1959, p. 187) wrote

> ”It is now common to regard the human operator of a machine and the machine itself as two elements in one overall man-machine system. ”The pilot and his plane, the helmsman and his craft, and the lathe operator with his lathe are examples of such systems.”

Humans are in this way described both as a mechanical component themselves and as a genuine part of a larger socio-technical or workplace. Although the Golem was not a machine in the sense we use the term today, but more like Galatea and Pinocchio – a humanoid object made from clay (Golem) marble (Galatea) or wood (Pinocchio) that miraculously became alive. The temptation to think of humans as a machine only grew stronger and became nearly irresistible after digital computing machinery became a reality around the middle of the 19th century, spearheaded by Turing (1950), followed by von Neumann (1958), and Arbib (1987), and concluded most famously by Ashby (1960). These and many other papers beginning by the mathematician George Boole’s writing about the Laws of Thought (Boole, 1854) made it philosophically and scientifically legitimate to describe the human operator as a machine and inadvertently supported the outrageous idea that human information processing could become a replacement for psychology (Lindsay & Norman, 1972), Newell & Simon 1963). The comparison between humans and machines was made on technological premises and therefore rarely advantageous to humans. An early expression of that view was provided by Paul Fitts, the undisputed pioneer of Human Factors Engineering, which he himself just named human engineering since, the primary concern being efficiency rather than reliability or safety Fitts argued that

> ”the final consideration which needs mention is relative fallibility of a man to a machine” (Fitts et al., 1951). p. 6)

since then it has been commonly agreed that humans at best are fallible machines. Fitts also noted that

> ”We have been very much occupied in perfecting the machines and tools which the worker uses in the economic arts, and went on “We have hardly attempted to improve the worker himself.” (Fitts, et al., 1951 P. iv).

The generic solution that Paul Fitts was instrumental in developing became known as human factors engineering, although he just called it human engineering with the stated purpose to engineer the human to ensure a better fit with the engineered machines. The three main remedies which from the beginning were endorsed by human factors were, **training**, **design**, and **automation**. Fitts’ concerns led him to propose a rigorous method to compare humans and machines, that we now known as the Fitts’ list which he introduced in the following way:

> We begin with a brief analysis of the essential functions ... We then consider the basic question: Which of these functions should be performed by human operators and which by machine elements? (Fitts, 1951 p.x).

This approach later achieved fame as the MABA-MABA list after the initials in (Men Are Better At) — (Machines Are Better At) (Dekker & Woods, 2002).

System definition

The notion of a human-machine system (HMS) became very popular initially called just a Man-Machine System (MMS). The term referred to the conventional definiion of a system as

> “a set of objects together with relationships between the objects and between their attributes” (Hall & Fagen, 1969, p. 81) – or even as anything that consists of parts connected together. In this definition, the nature of the whole is arbitrary, and the boundary of the system is therefore also arbitrary.

This view is illustrated by the following delightful quote:

> “It is legitimate to call a pair of scissors a system. But the expanded system of a woman cutting with a pair of scissors is also itself a genuine system. In turn, however, the woman-with-scissors system is part of a larger manufacturing system – and so on. The universe seems to be made up of sets of systems, each contained within a somewhat bigger, like a set of hollow building blocks.”(Beer, 1959, p. 9)

Taking a different approach Cognitive Systems Engineering (CSE)(Hollnagel & Woods, 1983) defined a system by how the boundary depends on the system’s functions, leading to the following

Table 3: *Definition of boundaries in Cognitive Systems Engineering (CSE)*

	functions that are essential for the ability of the JCS to maintain control.	Functions that are of no consequence for the ability of the JCS to maintain control.
Objects that can be effectively controlled by the JCS.	1.Objects are included in the JCS	2. Objects may be included in the JCS

Objects that cannot be effectively controlled by the JCS	3.Objects are not included in the JCS	4. Objects are excluded from the description as a whole.

The lack of an authoritative account of what human factors are has consequences also for the history of human factors – when did the concern for the human factor begin. Some date the beginning to 1898 to the efforts of Frederick W. Taylor, a mechanical engineer who promoted the idea of Scientific Management (Taylor, 1911).
Taylor's Scientific Management, known as Taylorism is interesting because the purpose unashamedly is to increase productivity, to make humans more efficient components in a place of work. This is obvious from the three main principles Taylor proposed(Taylor, 1911 p. 14):

1. Analyse tasks to determine most efficient performance.
2. Select people to achieve best match between task requirements and capabilities
3. Train people to ensure the specified and required performance but nothing more and nothing less than that,
4. Insure compliance by economic incentives

The first principle alone, makes clear that the primary concern was productivity rather than the well-being of the workers.

Some ambitiously date the beginning of human factors centuries earlier to Leonardo da Vinci's famous drawing of the Vitruvian man, dated 1490.

Others realisticallyput the start much later, after the end of the Second World War, (Hale & Hovden, 1998) in their proposal of three ages of safety, for instance wrote:

> "Just as the second age of human factors was ushered in by increasing realizations that technical risk assessment and prevention measures could not solve all problems, so were the 1980s characterized by an increasing dissatisfaction with the idea that health and safety could be captured simply by matching the individual to technology. [which otherwise was the foundation of human factors engineering]. The 1990s are already well into the third age of safety, where management systems are the focus of development and research." (Hale & Hovden, 1998, p. 130).

It is far more reasonable to date it to the period after the end of the Second World War, where the discrepancy between the systems designed to snake optimum use of technological capabilities and the innate or natural human ability to fit into these newly designed systems became so large that it constituted a problem both for the efficiency by which these systems functioned, and for their reliability, that it became necessary to do something about it. (Hale & Hovden, 1998) in their proposal of the three ages of safety, saw it primarily as a safety problem, recognisable by the tendency to consider human error the default socially accepted cause as described further in the following essay. But an explicit concern for the human factor became a problem already at the start of the first age of safety thinking, that Halæe & hovden, 1998) called the age of technology marked by the widespread use of the steam engine that ushered in the second uíndustrial revolution. (It is actually possible to distinguish among even more ages or stages in safety thinking. (Hollnagel, 2026, for instance proposes no less than seven stages or eras, including, of course the three covered by Hale & Hovden, 1998). Much of today's fashionable discussion about industry generstions, for instance, Industry 5.0, with no justification for using decimals in the numbering, miss that the first industrial revolution was the change to agriculture, and that industry 5.0, therefore ought to be industry 6.0. The transition to agriculture 12.000 years ago required as much use of novel technology as the "proper" industrial revolution in 1769, although as ploughs pulled by oxen rather than large. noisy steam engines. The industrial revolution in the 18th century logically forced a focus on the human factor, because people from then on had to work in conditions (factories) designed by others rather than as artisans in conditions gradually developed by themselves to suit their own needs.

Training as a human factors solution

Human factors engineering from the very beginning clearly needed ways to engineer and improve the human factor, mostly to ensure a better fit between humans and the machines they had to work with but also to make humans perform more reliably like machines or technological artefacts. Training, design, and automation were solutions adopted from the start that have been used ever since, although in various proportions since they are not mutually independent. **Training** was used to shape humans to meet the requirements of technology and was initially seen as an ideal and inexpensive solution, which unlike design and automation had a pedigree stretching back thousands of years, which was already widely practiced by most industries and therefore familiar. **Design** was used to ensure that the requirements of machines corresponded to the "natural" abilities of humans, as determined e.g., by the Fitts' list. And **automation** was the ultimate solution of replacing humans by technology, thereby enabling, if not ensuring a perfect match between system components.

It did, however, not take long before it became obvious that training was not the perfect solution to engineer the human factor and overcome the problems, due to performance variability, and insufficient reliability. Training was by some seen as a "Procrustean" approach referring to the legendary Procrustes, a robber of Attica, who had an inn on the sacred road between Athens and Eleusis (home of the Eleusian mysteries). The inn only had a one iron bed, where Procrustes invited everyone who passed by to spend the night. Procrustes became famous for making the visitors fit the bed, rather than the other way around. If visitors were shorter than the bed, Procrustes would stretch them until they were long enough, and if they were too long he would make them fit the bed by having parts of their legs cut off. Taylor and Garvey, 1959) used this analogy to criticise training when they wrote that

"Two rather different human factor approaches may be distinguished in efforts to optimise the performance of man-machine systems.

One approach tries to standardise performance by "Making people shorter" , i.e., limit or constrain what people should do, to limit and constrain their performance and to use less than their full potential. The other effort approach is to "make people longer" to extend or stretch human capabilities to meet task demands through additional specialised training, and or technological support

The Achilles heel of human factors, as a discipline, is that it blindly uses the machine analogy as an explanatory and analytical principle for what a human is, and because of this technological bias humans were from the very beginning of human factors seen as inefficient variable, and unreliable, a view that led to the perceived need of human factors engineering cf. the Fitts quote above, and an analogy which still is widely accepted, humans were therefore seen as a liability in consequence of which the human factor, and specifically human performance variability became problems that had to be solved, "We have been very much occupied in perfecting the machines and tools which the worker uses in the economic arts. We have hardly attempted to improve the worker himself." (Fitts, 1951, p. iv). (It was this postulated need to improve or engineer the worker that gave rise to human factors engineering which at some point was shortened to just human factors. The starting point was then as now the widely accepted view that the human operator was unreliable and inefficient in comparison to machines and that these imperfections made it impossible fully to exploit the advantages new technologies seemed to offer. The Main driving forces behind the ever-growing complexity of human-machine systems and work environments, in the 1940s and today, are summarised in Figure 1.

In Figure1, the technology potential is mostly a euphemism for saying that someone has spotted an opportunity to save money in recent years often seen by the introduction of remote work in many industries such as remote control of offshore oil rigs, because transporting operators by helicopter from land to an off-shore platform is both risky and expensive. Other examples are the use of remote towers in air traffic management (Papenfuss & Friedrich, 2016), remote control of ships sailing on regular routes, and even robotic surgery (Nyssen & Blavier, 2019). So in relation to NASA's famous Faster better cheaper slogan (Paxton, 2007), cheaper is practically always the highest priority for industries (whereas for NASA faster was more important. Technology is in most cases a solution looking for a problem and is usually adopted because it seems advantageous to do so rather than because there is a well-defined need. Most managers are what the father of Cybernetics Norbert Wiener (1964), called gadget worshippers, defined as people:

who "regard(ed) with impatience the limitations of mankind, and in particular the limitation consisting in man's undependability and unpredictability"

Today Wieners' concept of gadget worshipping has practically become a design philosophy of its own, called solutionism described as:

An intellectual pathology that recognizes problems as problems based on just one criterion: whether they are "solvable" with a nice and clean technological solution at our disposal. (Morozow, 2013a & b).

New technology may well be introduced with a specific and reasonable purposes in mind, even if they more often are economic rather than ergonomic, but any change soon falls prey to what Larry Hirschhorn called The Law of stretched systems which states the following:

Every system is stretched to operate at its full capacity; whenever improvements are made, for whatever reason, they will be used to achieve a new intensity and tempo of activity — to stretch the system a bit further. (Hirschhorn, 1997)

Training used to limit performance to match demands

The old idea that people should have just the competence needed to do the job, but no more because it would be a waste of effort, got a new lease on life with the introduction of Scientific Management (Taylor, 1911). Scientific management hoped that performance variability could be reduced and work efficiency improved by ensuring that people did exactly what they should do, neither more nor less requirements of the system, The first solution of "making people short" resembles the third principle of Taylorism, listed below from (Taylor, 1911 p. 14)

1. Analyse tasks to determine most efficient performance
2. Select people to achieve best match between task requirements and capabilities
3. Train people to ensure the specified and required performance but nothing more and nothing less than that,
4. Insure compliance by economic incentives

Training used to extend performance to meet demands

Training can also be used to "Make people longer " ie, to extend or stretch their capabilities and skills beyond what they would naturally or normally do through long and specialised training, This has unfortunately has become increasingly necessary as we relentlessly build human-machine systems and work environments that no longer are intuitive despite ambitious attempts such as ecological interface design (EID) (Vicente & Rasmussen, 1992) and SOAR (Laird, Newell & Rosenbloom,1987), cf. Figure 2. Few people can naturally or intuitively fly a modern aircraft, or control a nuclear power plant, and perhaps not even drive a present day EV. On the other hand few of us would presumably be able to drive a Ford T. It would be too simple and instrumentation and controls would be unrecognisable to us.

Shortcomings of training

Either solution provides only a temporary relief at best, The problem with both is that people inevitably and involuntarily will revert to their "natural" level of skills and performance whenever something unexpected happens, as it is bound to do according to Murphy's Law (Bloch, 1977, p. 11) because even though of any specific type of disturbance always can be made acceptably low, the possibility, that something will happen will still be equal to 1. But this does not mean that the probability is also 1, as there is a crucial difference between possibility (Zadeh, 2014) and probability. (Dubois & Prade, 1989), and when this happens people will either neglect or forget the constraints and therefore do more than they were supposed to do (e.g., intervene, manage by exception (Dekker & Woods, 1999) , hence do too much or do it too early), or to their former level of comfort and competence and forget what they have been taught and therefore do less than they were supposed to do (e.g,regress to habitual, intuitive or default actions, wrong procedure, do too little and do it too late) This will inevitably increase the difference between what people actually do. (Work-as-Done, or WAD) and (Work-as-Done, or WAD), and what they were supposed to do (Work-as-Imagined or WAI), which is the opposite of the intended outcome In either case the predictable and totally involuntary return to the "normal or "natural" way of performance renders the system incapable of functioning as intended. Despite the allure of the machine analogy and the promises of Skinnerian behaviorism, humans cannot be programmed as mindless robots. It was presumably never a problem for Procrustes if people regained their original length during the night. The stretching might well cease to have an effect, but legs would not spontaneously begin to grow. Training has since the early years been an essential solution for human factors engineering as such and the popularity has hardly diminished over time, despite insufficient evidence that it actually works. Training can, of course, be supplemented by recruiting, i.e., selecting people who naturally have the required competencies as already suggested by the second principle of Scientific Management.

Since the systematic study of Human-Machine Interaction (HMI) arose from the need to solve practical, technological problems, the basis for the description was the engineering view of men and machines. The technical and engineering fields had developed a powerful vocabulary to describe how machines worked and it therefore seemed natural to apply the same vocabulary to how people worked, i.e., as a basis for modeling human performance (e.g. Stassen, 1986). I will refer to this approach as the forced automaton analogy.

The automaton analogy denotes how one can think of or describe a human being as an automaton or a machine. A particular case is the use of the information processing metaphor (Newell & Simon, 1972, Simon (1972 and Newell, 1990). But the automaton analogy can be found in practically every explanation of human performance , e.g. by behaviourism or psychoanalytic theory (cf. e.g. Weizenbaum, 1976).

Worker-as-imagined

Training and design both imply an idea not just of Work-as-Imagined – but also of Worker-as-Imagined, just as design implies the idea of work-as-imagined (to be discussed later). Scientific Management was refreshingly honest about a worker-as imagined and the extreme case is, of course, the way Procrustes tried to make the physical size of his visitors fit the iron bed(He might also have posted a sign saying “bed requires people who are 178 cm long”) or whatever unit was in use at the time. No one would of course dream of doing that today, we rely on people doing it themselves by bending and stretching leading to physical effects such as back pain. But we nevertheless do it when it comes to the cognitive sizes as (Taylor & Garvey; !959) pointed out. Trying to make sure that everyone has the same physical or cognitive size is an example of standardisation.

The allure of standardisation

Training can also be used to “Make people longer" i.e., to extend or stretch their capabilities and skills beyond what they would naturally or normally do through long and specialised training, This has unfortunately become increasingly necessary as we relentlessly build human-machine systems and work environments that no longer are intuitive despite pretentious attempts such as ecological interface design (EID) (Vicente & Rasmussen, 1982) and SOAR (Laird, Newell & Rosenbloom,1987), cf. Figure 2. Few people can naturally or intuitively fly a modern aircraft, or control a nuclear power plant, and perhaps not even drive a present day EV. On the other hand few of us would presumably be able to drive a Ford T. It would be too simple and instrumentation and controls would be unrecognisable to us. Standardisation is generally attractive because it improves predictability, if something conforms to a standard in size, measure of functioning there is little uncertainty. If physical components are standardised, it becomes easier to assemble them, cf. Womack, et al. (1990). In terms of efficiency-thoroughness trade-off efficiency can be increased locally because thoroughness was taken care of earlier

The irony of design

If training is unable to meet its stipulated purpose, namely to ensure that human performance, or what the human operator does correspond to the requirements of the technology of the human-machine system, then the logical alternative is to design the technology so that the requirements correspond to how people naturally perform, requiring neither too much, nor too little. This is why design was the second main element to ensure effective human-machine interaction (HMI) and the second tool in the human factors’ toolbox. Cf. (Hollnagel,1993)

To improve the performance characteristics of these man-machine combinations, one has the choice of either trying to alter the man so that he fits the machine better, or modifying the mechanisms to fit the man. The latter approach is that of the engineer and the new specialist called the human engineer, while efforts to adapt the man to the machine characterise the work of the training expert.(Taylor & Garvey, 1959, p. 18.)

The automaton analogy is useless and even potentially misleading for HMI,HCI studies, (Hollnagel, 1992). I will not argue that the automaton analogy is ineffectual as a basis for describing human performance per se; I simply take that for granted. (This point of view has certainly not always been generally accepted and often not even clearly stated even when it was used extensively, for instance, by the mainstream of American Cognitive Science in the 1970s; it is nevertheless a view which is fairly easy to support.) I even believe that the automaton analogy is useless for machines in the context of human-machine systems where the functioning of the machine must be seen together with the functioning of a person.

The forced automaton analogy suggests that it is legitimate to think of or describe a human as an automaton or a machine. A particular case is the use of the information processing metaphor (Newell & Simon, 1972, Simon (19729 and (Newel, 1990) —- or even worse, preposterously claiming that a human is an information processing system (as for instance Simon(1972) and Newell(1990) both did.)

In the field of human-machine systems HMS, and in particular Human-Computer Interaction (HCI) the automaton analogy is, however, useless and even misleading. Instead I will go even further and argue that the automaton analogy is useless, and possibly harmful even for machines when the context is man-machine systems, i.e., when the functioning of the machine must be seen together with the functioning of a human being. (Hollnagel, 1992).

The designer's problem can be considered in the situation where there are two (or more) options available to solve a specific problem, e.g. displaying information about the state of a sub-system. If both of these options match the requirements (e.g. with regard to information requirements or ease of diagnosis) and constraints (e.g. with regard to maintaining safety or productivity), then both alternatives are equally good. However, the designer would usually prefer that there was a further constraint or requirement, so that a choice could be made between the two options. One often used solution to that problem is design guidelines, which can assist the designer in the following ways:

- Identify the general essential constraints that must be met by the design, e.g.

- Identify the specific essential constraints that hold for specific conditions that may occur, e.g.
- Describe the options for design and the consequences of each option vis-a-vis the essential functional requirements,
- Identify the specific supplementary constraints that may apply to each situation (i.e., constraints which comply with the functional requirements, but relate to additional requirements that only need be fulfilled if the conditions allow),(quoted from Booth, 1989)

The forced automaton analogy

An automaton can in general be described by a set of inputs, outputs, internal states and the corresponding state transitions; a classical example of this is the Turing machine (Turing, 1938). More formally, a finite automaton is a quintuple (e.g. Arbib, 1964; Lerner, 1975):

$A = (I, O, S, \lambda, \delta)$

where I is the set of inputs,

O is the set of outputs

S is the set of internal states

$\lambda: S \times I \rightarrow S$ is the set of rules for determining the next state, and

$\delta: S \times I \rightarrow O$ is the set of rules for determining the next output.

A machine, or a program, can in general be described as a finite state automaton or a state machine, in terms of the quintuple defined above (a set of inputs, outputs, internal states, and state transitions; all programming languages are actually based on that assumption. In order for the machine to work and to produce a predefined output, it must get a correct input. For HCI this means that the human user must respond in a way that corresponds to the predefined categories of input. If not, the machine, hence the joint cognitive system (Hollnagel & Woods, 1983), will not function appropriately; the HCI will become unreliable or worse ineffective.

Assume, for instance, that the automaton is in a state Sj. From that state it can progress to a predefined set of other states (Sk, ..., Sm) or produce a predefined output (Oj) only if it gets the correct input (Ij). If the system in question is a joint human-machine system then the input to the automaton is provided by the user. Therefore, the user must therefore provide an input that the automaton can interpret. Any other response by the user will lead to one of two cases:

- The input may not be understood by the automaton, i.e., transition functions λ and δ are not defined for the input. In this case the automaton may either do nothing or move to a default state.
- The input may be misunderstood by the automaton, for instance if the semantics or the syntax of the response have not been rigorously defined. In this case the automaton may possibly malfunction, i.e., be forced to a state which has not been anticipated. This happens easily if the input was a physical control action or manipulation, e.g. like switching something on or off, opening a valve, etc.

If the Human-Machine System (HMS) therefore is to function properly the user must provide a response that falls within a limited set of possible responses (the set of recognisable inputs). But the determination of the user's response is at least partly determined by the information available. The output from the machine constitutes the input to the user and is partly determined by the previous input i.e., what the user did, which means the HMS is a non-trivial system (von Főrster & Pőrksen, 2002). The content and structure of the machine's output must therefore be correctly understood by the user, i.e., correctly interpreted and mapped onto one of the predefined answer options. In order to do that the designer is forced to consider the user as a finite state automaton, as a machine. We know from practice that people may interpret information in many different ways depending on the context. We also know that there is no way in which we can possibly account for this infinity of interpretations (Work-as-Imagined or WAI will always be different from Work-as-Done or WAD! we therefore assume that the user interprets the information in a limited number of specified ways — and designers of course take every precaution to prevent misinterpretations. In other words, designers try to force the user to function as a finite automaton and therefore think of the user as a finite automaton. A good example of that is the graphical user interfaces, which basically serve to restrict the user's degrees of freedom.

This means that the starting point of thinking of the machine as a finite automaton willy-nilly forces the designer to think of the user as a finite automaton as well because there is no other conceptualization or model of the user that will fit the requirements of the design. Yet this is clearly unacceptable, no matter how sophisticated we assume the automaton to be (even if the number of states is exceedingly large and the state transitions are stochastic or multi-valued). It probably is the case that whatever analogy we use for one (machine or man) we will have to use for the other (man or machine a kind of forced anthropomorphism, perhaps) as well. Because we want to retain some distinct human elements in the description of the user we are forced by this argument to apply the same elements to the description of the machine (a kind of forced anthropomorphism). One proposal for an alternative description is the notion of a cognitive system as defined by Hollnagel & Woods (1983):

A cognitive system produces intelligent action, that is, its behaviour is goal oriented, based on symbol manipulation and uses knowledge of the world (heuristic knowledge) for guidance Furthermore a cognitive system is adaptive and able to view a problem in more than one way. A cognitive system operates using knowledge about itself and the environment in the sense that it is able to plan and modify its actions on the basis of that knowledge. (Hollnagel & Woods, 1983, p. 589)

Design is telling stories about the future

A further concern with design as a solution, is that design is tantamount to telling stories about the future (Roesler, et al., (2001). Design is about shaping something, an artefact, a work process, a task or job, a machine, or a tool, that not yet is, in particular shaping how it is going to be used or how it should be used to perform as intended. To do so it is necessary to make assumptions, not only about Work-as- Imagined. But also about the larger context or World-as-imagined. And telling stories about the future invariably changes the future, including how working conditions will be, and how people respond to the artefact (Merton,1936). What kind of adjustments and trade-offs will people use to ensure the artefact functions as they expect despite the way it has been designed Designing an artefact cannot avoid also being Cognitive Task Design (Hollnagel, 2003) Telling stories about the future again highlights the potential and actual differences between Work - as .Imagined and Work-as-Done. In summary, the irony of design is that it is telling stories about the future which thereby affects or changes the future. Changing the premises for design decisions introduces a recursion that creates very difficult problems, as recursion usually does. The specific problem here is that it proposes improvements to a situation and set of working conditions that are imprecisely known. If design fails as an effective tool, another problem to be solved is the difference between Work-as-Imagined (WAI) and Work-as Done (WAD).

Work-as-Imagined and Work -as-Done

Two terms that frequently appear in contemporary approaches to safety and work management in general are Work-as-Imagined (WAI) and Work-as-Done (WAD) (Shorrock, 2020) (Hollnagel, 2017). They also played an important role in the initial discussions about resilience engineering, as described by (Dekker, 2006), although the origin can be found much earlier in the French ergonomic tradition (Ombredane & Faverge, 1955) (Leplat, & Hoc, (1983).

The meaning of the two terms is — presumably — obvious.

- Work-as-imagined (WAI) represents the various assumptions, explicit or implicit, that people have about how their or others' work should be done. WAI is clearly related to the concept of requisite imagination (Adamski & Westrum, 2003).
- Work-as-done (WAD) represents how something is actually done, either in a specific case or more routinely the characteristic or habitual way of carrying out an activity, and therefore what ought to have been imagined.
- There will for a number of reasons always be a difference between how work is 'imagined' or thought of and how work is actually done. This may or may not be problematic.
- The solution to this difference is to try to understand what determines how work is done and to find effective ways of managing that to keep the variability of WAD within acceptable limits. (but not by constraints and compliance, by making people "shorter".

This Difference erodes the very basis from which the design is made, the functional requirements and constraints as well as the matching criteria and design rules so that it is nearly inevitable that design will not solve the problem it was supposed to address but makes it worse – by creating more uncertainty it ironically increases the need of human performance variability to fill out the gaps between design (WAI) and reality (WAD).

There is first of all a need to consider and describe work as it should be done during the design of machines and tools, not least since human-computer interaction nearly always is part of it. It may be easier to describe for a robot. There is also a practical need to think about how work should be done as part of managing and scheduling actual work under real conditions, for instance, to ensure that enough people with the requisite competencies are available when needed to meet the expectations of employers, customers and clients, and the needs of the task. And there is finally a need to think about how work should have been done when events are being analysed post hoc — which usually means some kind of accident investigation. Unfortunately, this is often reduced to inventing stories about the past or proposing explanations in terms of counterfactual conditionals — such as “if only they had done X, then Y would not have happened, or if only we had had more of, or Z (safety culture/situation awareness/trust /communication, then Y would not have happened”. From a scientific and practical perspective it is regrettable that the need to explain and understand WAD is so obvious and sometimes even mandatory, when something has gone wrong , while it is practically non-existent when “nothing“ happened (the dynamic non-events proposed byWeick (1987) because everything and everyone just worked as it/they should. (Hollnagel, Laursen & Sørensen, 2022)

Conclusions about design as a solution

There is an irreducible difference between WAD and WAI, why is it irreducible? The consequence of that is the designed systems function not because of the design, but because of human performance variability that compensates for design deficiencies. The irony of human factors and even more of safety science (Hollnagel, 2013) is eliminating human performance variability if it actually was possible would not solve the problems of imprecision, variability and lack of speed, but only make them larger, and confront us with the fact that we cannot do without that the human factor which we so desperately want to get rid of it.

Conclusions

Human factors engineering from the beginning relied on three main tools namely, training, design, and automation and each of these harbor at least one irony. In her (1983) paper Lisanne Bainbridge presented and discussed three ways, called ironies, in which automation of industrial processes could expand rather than reduce the problems of inefficiency and insufficient reliability. For each of these three ironies the outcome may at times be so different from what was intended, that any hoped for improvements never arrive.

Can we ever imagine how work is done?

Returning to the question above, the answer is the typical human factors reply of “Yes, but ...” The answer is on the one hand affirmative, because we certainly can imagine how work is (to be) done if we try, especially if we pay attention to what actually happens instead of relying on what we imagine or wish should happen or counterfactual hypotheses of what should have happened. On the other hand,the provisory “but” signifies that we should not expect ever to achieve a perfect match. The solution is neither to force WAD to comply with WAI — as in the Zero Accident Vision and (Zwetsloot et al., 2001) and Lean — nor to constrain WAI so that it corresponds to WAD as Scientific Management unsuccessfully tried to do. Work-as-Done is a moving target because internal and external working conditions (demands, and resources) rarely are stable or predictable. The solution is rather to try to understand what determines how work is done and to find effective ways of managing that to keep the variability of WAD within acceptable limits. (but not by using constraints and compliance to make people “shorter”).

not by constraints and compliance, by making people “shorter”.

Performance variability as a liability

Human performance variability, individually and/or collectively is necessary to adjust the ways and means to fit the current conditions as they actually are, and not as they have been imagined at design-time. Yet it is normally described by negatively loaded terms for example, as deviations, violations or non-compliance (of which there are several types, cf., Hudson et al., 2008)., or even as just the utterly useless “human error” (Le Coze, 2022). The negative descriptions of performance variability conveniently disregards that it can be found at both the sharp-end and the blunt-end, and that resorting to such categories usually means that the analyst has made an efficiency-thoroughness trade-off. There is a further small irony here, namely that performance variability by management usually is seen as positive (at least by management itself), and called flexibility or innovation, even though it by its nature does not differ from the performance variability at the sharp end.

Performance variability as an asset

The greatest problem in describing performance variability as an asset is that there are few if any terms that can be used to do it. We are blessed with nearly countless error taxonomies, but have practically no words or terms to characterise and therefore talk about that which goes well, the dynamic non-events that Weick (1987)referred to. The most notable exceptions are the concept of efficiency-thoroughness trade-off (ETTO, Hollnagel, 2009). Another candidate term is serendipity (Merton & Barber, 2011)

The Ironies of automation

In the 1983 paper Bainbridge described three ironies of automation

The first irony is that the more advanced a control system is, the more crucial the contribution of the human operator will be

The first irony is that the designer who tries to eliminate the operator still leaves the operator to do the tasks which the designer cannot think how to automate. This is another consequence of the limitations of Work-as-Imagined

The third irony is that the automatic control system has been put in because it can do the job better than the operator, but yet the operator is being asked to monitor that it is working effectively

A further issue of automation that Bainbridge did not address at the time, because it had not yet been recognised as a problem, is due to the substitution principle

The Fourth irony of automation

The substitution principle expresses the common assumption that artefacts are neutral in their effects and that their introduction into a system therefore only has intended and no unintended consequences. The basis for this principle is the concept of interchangeability, which of course has proved its value as the basis for large scale industrialisation. Thus if there are a number of identical parts, such as light bulbs or pumps, it is possible to replace one by another without unwanted side-effects. In general, however, substitutability only works when parts are not interacting and when there is no appreciable tear and wear. If parts are interacting, they constitute a system with dependencies, which almost by definition invalidates the substitution assumption.

Automation was seen as a possible replacement for humans because it is based on components with known failure rates. Humans, on the other hand, are generally seen as fallible and unreliable, as ‘proved’ by countless examples of ‘human error’. The fallacy of this argument should by now be so obvious that it hardly needs to be belaboured. The substitution principle expresses the common assumption that artefacts are neutral in their effects and that their introduction into a system therefore only has intended and no unintended consequences. The basis for this principle is the concept of interchangeability, which of course has proved its value as the basis for large scale industrialisation (Womack, Jones & Roos, 1990) Thus, if there are a number of identical parts, such as light bulbs or pumps, it is possible to replace one by another without unwanted side-effects. In general, however, substitutability only works when parts are not interacting and when there is no appreciable tear and wear. If parts are interacting, they constitute a system with dependencies, which almost by definition invalidates the substitution assumption.

If the substitution principle is dubious in the case of technological components, it is even more suspect in the case of a substitution of functionality of which the extreme is replacing humans by automation. Carroll & Campbell (1988p. 4), for instance, noted that “(n)ew tools alter the tasks for which they were designed, indeed alter the situations in which the tasks occur and even the conditions that cause people to want to engage in the tasks” (; see also Sarter, Woods & Billings, 1997). In relation to risk reduction, this means that a substitution of functions changes the basis for the risk assessment, often in a fundamental way. It is consequently not warranted to claim that the substitution is value neutral (Besnard & Hollnagel 2012) unless both short and long term consequences of the change are fully taken into account.

When training is unable to meet its stipulated purpose, namely to ensure that human performance, or what the human operator does correspond to the requirements of the technology of the human-machine system, then the logical alternative is to design the technology so that the requirements correspond to how people naturally perform, requiring neither too much, nor too little. This is why design became the second main element to ensure effective human-machine interaction (HMI) and the second tool in the human factors’ toolbox. cf. (Hollnagel,1993)

The design of HCI is based on a number of assumptions. Some of these are explicitly stated in design guidelines. Others are hidden in the design and possibly concealed even for the designer. A particular instance is the assumptions about the reliability of human performance --- and cognition --- and about which things can go wrong. It is important for designers of HCI to know more about human reliability, regardless of whether the HCI is for industrial or academic applications. Lack of knowledge may deceive designers to rely on their personal experience or intuition. That is, however, not a valid basis for the design and the resulting system is therefore likely to be inadequate. This caution is pertinent for human reliability as well as for the, unfortunately, many other aspects of man-machine interaction that do not receive the attention they rightly deserve.

The design of Human-Computer Interaction (HCI) is based on a number of assumptions. Some of these are explicitly stated in design guidelines. Others are hidden in the design and possibly concealed even for the designer.

A particular instance is the assumptions about the reliability of human performance – and cognition – and about which things can go wrong (Swain and Guttman, 1983, Cojazzi et al, 1883, Kirwan, 1994, Hollnagel, 1998 Boring(2012). l. It is important for designers of HCI to know more about human reliability, regardless of whether the HCI is for industrial or academic applications. Lack of knowledge may tempt designers to rely on their personal experience. That is, however, unlikely to constitute a valid basis for the design and the resulting system is therefore likely to be inadequate. This caution is pertinent for human reliability as well as for the, unfortunately, many other aspects of man-machine interaction that do not receive the attention they rightly deserve.

Thus the fourth irony is that the automatic replacing the human operator with automation fundamentally changes the human-machine system. and the work that the operator does, hence the premise for what should be automated. . The fourth irony is therefore a paradox and in that way more, and more serious than an irony such.

Performance variability or performance agility?

The term performance variability was introduced partly to avoid the negative connotations of deviance , violations, and non-compliance. But even variability is tainted, since it clashes with the ideal of steady machine-like functioning. If something is variable it is also not completely predictable, which is not a desired quality. An alternative to performance variability could be agility: the state or quality of being agile.It is human agility that allows people to overcome the temporary glitches and problems that every tool use involves

Performance agility as putty?

The ideal way to build a house, a bridge or a wall, is to make sure that every element or stone fits exactly to its neighbour. This can be seen for instance from the Pyramids, Roman bridges, or the walls around Japanese castles, in e.g., in Osaka or Himeji. But this a slow process that requires enormous work and infinite care.The alternative to use in stones that each is carefully shaped to fit with neighbouring stones, is to use prefabricated elements that do not necessarily fit exactly, such as the bricks used to build houses today, and then to fill in the gaps with putty. In this case the components only have to fit reasonably well to each other, but not precisely so. The putty will cover any gaps and help to provide the appearance of a whole, especially when some time has passed. Since it is impossible to fit humans exactly to match the machines, lest we resort to the Procrustean solution, and also impossible to fit the machines exactly to the individual human tailor made tools like tailor made clothes are far too expensive. Instead some kind of putty is needed to ensure that the human-machine system works. In the case of human factors the putty is human agility or performance variability. The deep irony of human factors is that we attempt to remove this putty, because it mistakenly has been seen as a liability rather than the asset it really is. There are interesting historical reasons why this is so, as this paper has tried to explain. But there are even more compelling reasons to stop doing so.

When WAI is compared to WAD, there is always a small gap or discrepancy. The two never match completely (Hollnagel, 2017). So just as when bricks are put together to build a wall or a house, something is needed to fill out and cover the gaps. The alternative would be to ensure that the bricks fit perfectly The design of HCI has generally suffered from the lack of a methodology or discipline corresponding to what can be found e.g. in the field of software engineering, although some attempts have been made (e.g. Dowell & Long, 1989). Design can roughly be defined as the art (or skill) of matching constraints with possibilities, as shown in Figure 1. If there are no constraints, the design can be anything that the system or the technological options allow. If there are too many constraints, then possibly the system capabilities are insufficient, i.e., there are no options that match the constraints. And if either constraints or possibilities are ill-defined, design becomes an unwieldy process.

The designer's problem can be considered in the situation where there are two (or more) options available to solve a specific problem, e.g. displaying information about the state of a sub-system. If both of these options match the requirements (e.g. with regard to information requirements or ease of diagnosis) and constraints (e.g. with regard to maintaining safety or control), then both alternatives are equally good. However, the designer would usually prefer that there was a further constraint or requirement, so that a choice could be made between the two options. One often used solution to that problem is design guidelines, which can assist the designer in the following ways:

whatever the mode of operation. because it is the conceptual putty that fills out the cracks between WAI and WAD. We have been too obsessed with the old view of humans as a liability and source of error and variability. The truth is that without this variability, no system would be able to function. It is the putty for the inevitable discrepancy between WAI and WAD !

If neither training, nor design can ensure that the human becomes the required and desired reliable component in the workplace,then the only remaining solution, and also the third tool in the human factors toolbox is automation, i.e., having a (reliable) machine take over part of the human's work and activities. In this case the ironies have already been described by Bainbridge (1983), and can therefore simply be summarised here:

Conclusions

The irony of training is

The irony of design is

the ironies of automation are, and the overarching irony of human factors is

The irony of human factors is that by trying to eliminate or at least neutralise the human factor, meaning the variability of human performance we expand the problem we try to solve. Because human agility is the putty that hides the cracks between Work - as Imagined and Work - as Done. The irony of that is that human factors are doing themselves a disservice. It is reducing and if possible eliminating human performance agility it is eliminating the quality that is necessary for human-machine systems to work.

Literature

Adamski, A. J., & Westrum, R. (2003). Requisite imagination: The fine art of anticipating what might go wrong. In Handbook of cognitive task design (pp. 193-220). CRC Press.

Ashby, W. R. (1960) Design for a Brain. The origin of adaptive behavior. (Second Edition). London: Chapman & Hall, Ltd.

Arbib, M. A. (1964). Brains, machines and mathematics. New York: McGraw-Hill.

Arbib, M. A. (1987). Brains, machines and mathematics. New York: McGraw-Hill.

Besnard, D., & Hollnagel, E. (2014). I want to believe: some myths about the management of industrial safety. Cognition, Technology & Work, 16, 13-23.

Arbib, M. A. (1964). Brains, machines and mathematics. New York: McGraw-Hill.

Besnard, D., & Hollnagel, E. (2014). I want to believe: some myths about the management of industrial safety. Cognition, Technology & Work, 16, 13-23.

Boring, R. L. (2012). Fifty years of THERP and human reliability analysis INL/ CON-12-25623).Idaho Falls, ID; IdahoNational Lab.(INL), (United States).

Bloch, A. (2003). Murphy's law. Harmondsworth UK: Penguin.

Boole, G. (1854). An investigation of the laws of thought: on which are founded the mathematical theories of logic and probabilities (Vol. 2). Walton and Maberly.

Carroll, J.M., Campbell, R.L. (1988). Artefacts as Psychological Theories: The Case of Human-Computer Interaction. User Interface Institute, IBM T.J. Watson Research Centre, Yorktown Heights, NY.

Cojazzi, G., Pedrali, M. & Cacciabue, P. C. (1993). Human performance study: Paradigms of human behaviour and error taxonomies ISEI/TE/2443/93). JRC Ispra, Italy: Institute for Systems Engineering and Informatics.

Booth, P. (1989). An introduction to human-computer interaction. London: Erlbaum.

Bainbridge, L. (1983a). Ironies of automation Astomatica, 19(6) pp. 775 779,

Bainbridge, L. (1987). Ironies of automation. In J. Rasmussen, K Duncan, and J. Leplat (Eds.) New Technology and Human Error, Wiley, Chichester, pp. 276-283. Analysis,design and evaluation of man—machine systems (129-135). Pergamon.

Dekker, S. W., & Woods, D. D. (1999). To intervene or not to intervene: The dilemma of management by exception. Cognition, Technology & Work, 1, 86-96.

Dekker, S. W. A., & Woods, D. D. (2002). MABA-MABA or abracadabra? Progress on human—automation co-ordination. Cognition, Technology ¢» Work, 4, 240-244.

Dekker, S. W. A. (2006) Resilience engineering: Chronicling the emergence of confused consensus. In E. Hollnagel, D. D. Woods & N. Leveson (Eds.), Resilience engineering: Concepts and precepts. Hampshire: Ashgate.

Dowell, J. & Long, J. (1989). Towards a conception for an engineering discipline of human factors. Ergonomics, 32(11), 1513-1535.

Dubois, D. & Prade, H. ((1989). Handling uncertainty in expert systems – pitfalls, difficulties, remedies. In E. Hollnagel (Ed.) The reliability of expert systems. Chichester: Ellis Horwood Ltd. (64 – 99).

Feigenbaum, E. A., & Feldman, J. (1963). Computers and thought (Vol. 7). McGraw-Hill: New York.

Fitts, P. M. et al. (1951). Human engineering for an effective air-navigation and traffic-control system. Washington D. C: National Research Council

Grøtan, T. O., Storseth, E, Ro, M. H., & Skjerve, A. B. (2008). Resilience,Adaptation and Improvisation—increasing resilience by organising for successful improvisation. In 3rd symposium on resilience engineering (pp. 28-30).

Hirschhorn L. (1997). Quoted in Cook RI, Woods DD and Miller C (Eds.), (1998). A Tale of Two Stories: Contrasting Views on Patient Safety. Chicago IL: National Patient Safety Foundation, , 1998.

Hollnagel, E. (2013). Is safety a subject for science? Safety Science, 67, 21-24.

Hollnagel. E (2017) Can we ever imagine how work is done? HindSight 25 10-13

Hollnagel, E.(2009). The ETTO principle: Why things that go right sometimes go wrong. Farnham, UK: Ashgate.

Hollnagel, E. (Ed.), (2003). Handbook of cognitive task design. Mahwah, NJ: Lawrence Erlbaum Associates.

Hollnagel, E. & Woods, D. D. (1983). Cognitive systems engineering: New wine in new bottles. International Journal of Man-Machine Studies, 18, 583-600.

Hollnagel, E. (1983). The design of reliable HCI: The hunt for hidden assumptions. In J. L. Alty, D. Diaper & S. Guest (Eds.), People and computers VIII. Cambridge: Cambridge University Press.

Hollnagel, E. (1992). The art of efficient man-machine interaction: Improving the coupling between man and machine. Invited paper for CADES workshop, St. Prix, France, August 1992. (Also in: J.-M. Hoc, P. C. Cacciabue & E. Hollnagel (Eds.), Expertise and technology: Cognition & human-computer cooperation. Lawrence Erlbaum.

Hudson, P. T. W. et al. (2008). Meeting expectations: A new model for a just and fair culture. In SPE International Conference on Health, Safety, and Environment in Oil and Gas Exploration and Production. OnePetro.

Hollnagel, E. (1998) Cognitive Reliability and Error Analysis Method (CREAM). Oxford: Elsevier Science Ltd.

Kirwan, B. (1994).A practical guide to human reliability assessment. London: Taylor & Francis.

Laird, J., Newell, A. & Rosenbloom, P. (1987) SOAR: An architecture for general intelligence. Artificial Intelligence 33:1-64.

Lindsay, P. H. & Norman, D. A. (1972). Human information processing, An introduction to psychology. NewYork:Academic Press.Lerner, A. Y. (1975). Fundamentals of cybernetics. London: Plenum.

Le Coze, J. C. (2022). The 'new view' of human error. Origins, ambiguities, successes and critiques. Safety science, 154, 105853.

Leplat, J., & Hoc, J.-M. (1983). Task and activity in the psychological analysis of situations. Cahiers de Psychologie Cognitive et Cahiers de Psychologie Marseille , 3 (1), 49-63.

Meister, D. & Enderwick, T. P. (2001). Human Factors in System Design, Development, and Testing. CRC Press, Inc.

Merton, R. K. (1936). The unanticipated consequences of purposive social action. American sociological review,1(6), 894-904.

Merton, R. K., & Barber, E. (2011). The Travels and Adventures of Serendipity. Princeton University Press.

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. Psychological review, 63(2), 81.

Newell, A. (1990). Unified theory of cognition. Cambridge, MA.: Harvard University Press.

Newell, A. & Simon, H. A. (1972). Human problem solving. Englewood Cliffs, NJ: Prentice-Hall.

Nyssen, A. S., & Blavier, A. (2019). Investigating expertise, flexibility and resilience in socio-technical environments: a case study in robotic surgery. InE. Hollnagel, J. Braithwaite & R. L. Wears (Eds.) Resilient health care (pp. 97-110). CRC Press.

Papenfuss, A., & Friedrich, M. (2016, September). Head up only—a design concept to enable multiple remote tower operations. In 2016 IEEE/AIAA 35th Digital Avionics Systems Conference (DASC) (pp. 1-10). IEEE.

Reason, J. (1990). Human error. Cambridge university press.

Sarter, N.B., Woods, D.D., Billings, C.E., 1997. Automation surprises. In: Salvendy, G. (Ed.), Handbook of Human Factors and Ergonomics, 2ndEd. New York: Wiley

Shorrock. S. (2020, October 28). Proxies for work-as-done: 1. Work-as-imagined. Humanistic Systems.https://humanisticsystems.com/2020/10/28/proxies-for-work-as-done-1-work-as-imagined/

Simon, H. A. (1972). The sciences of the artificial. Cambridge, MA.: The M. I. T. Press.

Stassen, H. G. (1986). Decision demands and task requirements in work environments: What can be learned from human operator modeling. In E. Hollnagel, G. Mancini & D. D. Woods (Eds.), Intelligent decision support in process environments. Berlin: Springer Verlag.

Taylor, F. W. (1911). *The principles of scientific management*. New York: Harper.

von Neumann, J. (1958). The computer and the brain. Yale University Press

Weick, K. E. 1987. Organizational culture as a source of high reliability. California Management Review 29(2), 112-128.

Weizenbaum, J. (1976). Computer power and human reason. From judgment to calculation. San Francisco: W. H. Freeman

Wienet, N. (1988). The human use of human beings: Cybernetics and society (No. 320). Da Capo Press.

Wiener, N. (1966). God & Golem, Inc.: A comment on certain points where cybernetics impinges on religion (Vol. 42). MIT press.

Swain, A. D. & Guttman, H. E. (1983). Handbook of human reliability analysis with emphasis on nuclear power plant applications (NUREG CR-1278). Washington, DC: NRC.

Turing, A. M. (1960). Computing machinery and intelligence. Mind, LIX (236),433-460.

Mawhin, B., Cabon, P., & Buratto, F. (201 1). Integrating Aircrew Resources Variability into the Design of Future Cockpit. In 16th International Symposium on Aviation Psychology (p. 481).

Morozow, E. (2013a). The perils of perfection. The New York Times, 2 March 2013.

Motozow, E (2013b). To save everything, click here. New York: The Perseus books group

Newell, A. & Simon, H. A. (1963). GPS, A program that simulates human thought. In: E. A. Feigenbaum & J. Feldman (Eds.) Computers and thought, 279-293. New York: McGraw-Hill Book Company.

Paxton, L. J. (2007). “Faster, better, and cheaper at NASA: Lessons learned in managing and accepting risk. Acta Astronautica, 61(10), 954-963.

Strauch, B. (2017). Ironies of automation: Still unresolved after all these years. IEEE Transactions on Human-Machine Systems, 48(5), 419-433.

Taylor E W.(1919). The principles of scientific management. New York: Harper & Brothers Publishers

Taylor, F V., & Garvey, W. D. (1959). The limitations of a ‘Procrustean’ approach to the optimization of man-machine systems. Ergonomics, 2(2), 187-194.

Vicente, K. J., & Rasmussen, J. (1992). Ecological interface design: Theoretical foundations. IEEE Transactions on systems, man, and cybernetics, 22(4), 589-606.

Moehlenbrink, C., & Papenfuss, A. (2011, September). ATC-monitoring when one controller operates two airports: Research for remote tower centres. In Proceedings of the Human Factors and Ergonomics Society Annual Meeting (Vol. 55, No. 1, pp. 76-80). Sage CA: Los Angeles, CA: Sage Publications.

Ombredane, A. and Faverge, J. M. (1955). L'analyse du travail, Paris: Presses Universitaires de France.

Reason, J. T. (1990). Human error. Cambridge university press.

Roesler, A., Feil, M., Puskeiller, A., Woods, D, D. and Tinapple, D. (2001). Design is telling stories about the future.Columbus, OH: Cognitive Systems Engineering Laboratory, Ohio State University,.

Zwetsloot, G. I. J. M., Aaltonen, M., Wybo, J.-L, Saari, J. Kines, P. & Op De Beeck, R.(2013). The case for research into the zero accident vision. Safety Science, 58, 41-48.

Swain, A. D. & Guttman, H. E. (1983). Handbook of human reliability analysis with emphasis on nuclear power plant applications (NUREG CR-1278). Washington, DC: NRC.

Turing, A.M. (1960). Computing machinery and intelligence. Mind, LIX (236),433-460.

Turing, A. M. (1936). On computable numbers, with an application to the Entscheidungsproblem. J. of Math, 58(345-363), 5.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, *79*, 281-299.

Von Főrster, H., & Pőrksen, B. (2002). Understanding systems: Conversations on epistemology and ethics. Heidelberg, Germany: Carl-Auer Verlag.

Weick, K. E. 1987. Organisational culture as a source of high reliability. California Management Review 29 (2), 112-128.

Womack J, Jones D, Roos D. The machine that changed the world. New York: Simon & Schuster; 1990.

Zadeh, L. A. (2014). A note on modal logic and possibility theory. Information Sciences, 279, 908-913.

making longer: They do less than they should do, hence may render the system incapable of functioning as planned (e.g. default actions, wrong procedure, too little and too late). The fourth irony is that we try to limit or eliminate the variability of human performance that is needed for systems to function as they were intended to — or at all. This is an irony of automation, but also of human factors and HSE management in general. The ETTO-principle only referenced

The irony is that by trying to constrain human performance and introducing barriers, you actually increase variability rather than reduce or limit it! The effect is therefore the opposite of what was intended, this is more than an irony it is a paradox

Bainbridge on performance variability;

Long-term knowledge . An operator who finds out how to control the plant for himself, without explicit training, uses a set of propositions about possible process behavior, from which he generates strategies to try (e.g, Bainbridge. 1981). Similarly an operator will only be able to generate successful new strategies for unusual situations if he has an adequate knowledge of the process. There are two problems with this for machine-minding operators. One is that efficient retrieval of knowledge from long-term memory depends on frequency of use (consider any subject which you passed an examination at school and have not thought about since). The other is that this type of knowledge

Performance variability as putty

When WAI is compared to WAD, there is always a small gap or discrepancy. The two never match completely (Hollnagel, 2017). So just as when bricks are put together to build a wall or a house, something is needed to fill out and cover the gaps. The alternative would be to ensure that the bricks fit perfectly together like Lego blocks or the stones in the walls of the pyramids, Roman bridges or Japanese castles. But it is far easier to use putty to fill out the gaps and in this way disguise the differences. In the context of human-machine systems, human performance

variability plays the role of putty and helps to ensure that the wall works as it should and looks complete particularly when it has been in place for some years.

Conclusion: The irony of human factors

As this paper hopefully has demonstrated,human factors as such therefore comprise five ironies, one for each of training and design plus three for automation. The overall irony for human factors is that it considers human performance variability as a liability and, even though it is an asset without which few systems would work.

Wiener, N. (1966). God & Golem, Inc.: A comment on certain points where cybernetics impinges on religion (Vol. 42).

MIT press.

Hollnagel, E., & D Woods, D. D. (1983). Cognitive systems engineering: new wine in new bottles. International Journal of Human-Computer Studies, 5 1(2), 339-356.

WAI and WAD

Origins, Shorrock WP

Hej, nu kan jeg igen ikke bruge mit dankort. jeg ved godt at der er en maxgrænse, men det ville være fint. hvis det også fremgik af kontooversigten