

A taxonomy of interactions in socio-technical systems: A functional perspective

Tarcisio Abreu Saurin, Dr.

E-mail: saurin@ufrgs.br, DEPROT/UFRGS (Industrial Engineering and Transportation Department, Federal University of Rio Grande do Sul), Av. Osvaldo Aranha, 99, 5. Andar, Porto Alegre, RS, CEP 90035-190, Brazil.

Riccardo Patriarca, Dr.

E-mail: riccardo.patriarca@uniroma1.it, Department of Mechanical and Aerospace Engineering, Sapienza University of Rome, Via Eudossiana, 18, Rome, 00184, Italy.

Abstract

Although the modelling of interactions has long been at the core of socio-technical systems theory, and is a key for understanding resilience, there is a lack of a holistic taxonomy of interactions. This study introduces a taxonomy of interactions to be used in association with the Functional Resonance Analysis Method (FRAM). The taxonomy has nine criteria: nature of agents, output nature, levelling, waiting time, distance, degree of coupling, visibility, safety and/or security hazards, and parallel replications. For each criterion, two descriptors are proposed: what the interaction looks like; and - when applicable - the variability level of the interaction. The use of the taxonomy is presented for three systems with clearly distinct complexity characteristics: cash withdrawal from an ATM, teaching a university course, and manufacturing operations. These case studies indicate the usefulness of the taxonomy for the identification of leverage points in work system design. They also show the value of modelling the variability of the interactions

in FRAM models, in addition to the traditional modelling of the variability of the outputs of functions. Implications of the taxonomy for resilience engineering are discussed.

Keywords: interactions, socio-technical systems, complex systems, FRAM, human factors, taxonomy.

A taxonomy of interactions in socio-technical systems: A functional perspective

1. Introduction

The modelling of interactions has long been at the core of socio-technical systems theory and systems-oriented safety approaches (Clegg, 2000). In the 1940's, the studies by the Tavistock institute in coal mines concluded that the best performance arises from the harmonic interaction between the social and the technical systems (Trist and Bamforth, 1951). More recently, the concept of joint cognitive systems also relies on the notion of interactions, by assuming that the human and non-human agents in work systems form an inseparable adaptive ensemble (Hollnagel and Woods, 2005). The definition of ergonomics also highlights interactions: according to Wilson (2014) "ergonomics is the theoretical and fundamental understanding of human behaviour and performance in purposeful interacting sociotechnical systems, and the application of that understanding to design of interactions in the context of real settings". In turn, the growing interest of human factors researchers in complexity science (Walker et al., 2010) has put a spotlight on dynamic interactions, which are a defining feature of complex socio-technical systems (CSSs) (Cilliers, 1998).

Furthermore, a number of modelling approaches used in human factors, such as agent-based modelling (Baber et al., 2013), social network analysis (Houghton et al., 2006), and the Functional Resonance Analysis Method (FRAM) (Hollnagel, 2012), are essentially about the modelling of interactions. On this context, it is possibly no overstatement to say that the design of means for modelling and coping with interactions is the main concern

of current research on systems-oriented human factors approaches. Some examples of recent studies, explicitly relying on the concept of interactions can be mentioned. Bolbot et al. (2019) discuss the vulnerabilities intrinsic to tight and complex interactions in cyber-physical systems. Klockner and Toft (2018) investigated rail safety occurrences and modelled contributing factors as a network of interacting factors. Maguire (2014) discusses the impacts of the new ways of working on the interactions between users and information and communication technologies.

Regardless of the key role played by interactions, there is a lack of holistic taxonomies for modelling what they look like under different circumstances. Perrow (1984) proposed the most well-known taxonomy of interactions in CSSs, according to two axes: from linear to non-linear interactions, and from tightly to loosely-coupled. While these are core dimensions of interactions, they are hardly operationalized as metrics and may be themselves emergent outcomes of other hidden system features. For example, Perrow defines linear interactions as those in expected and familiar sequences, quite visible even if unplanned, and characterized by the proportionality between cause and effect. This definition encompasses attributes that could be assessed by their own, namely the observer's familiarity with the interaction, the visibility of the interaction, and its impacts. Furthermore, the two characteristics mentioned by Perrow are functional, rather than structural, properties of a socio-technical system. This means that they change over time (El Maraghy et al., 2014), and this might be due to the variability of their underlying contributing factors. The understanding of the said factors may be useful when having the objective of influencing the system through design.

Given this context, the research question addressed by this study is stated as follows: how should a taxonomy for interactions be defined for supporting socio-technical work system's design and analysis? The use of taxonomies in human factors is common, as it serves several purposes, such as (Olsen and Shorrock, 2010) the definition of a vocabulary for sharing information as well as support to decision-making in design by highlighting system trends, strengths, and weaknesses. Examples of such taxonomies can be cited, such as a taxonomy of slack proposed by Saurin and Werle (2017) and the human factors analysis and classification system, applied to the study of human errors in aviation (Wiegmann and Shappell, 2001).

The taxonomy of interactions proposed in this study is intended to be compatible with the FRAM, which has been the main modelling tool in resilience engineering (Hollnagel, 2012). There are two main reasons for choosing the FRAM, namely: (i) its functional emphasis, which is a key for modelling dynamic interactions; and (ii) it can in principle be applied to the modelling of any interaction type (e.g. social interactions, flow of materials, logical dependence), which implies in a broader scope when compared with other approaches. FRAM's potential in these regards has been largely confirmed in a variety of socio-technical systems, such as healthcare (Clay-Williams et al., 2015), aviation (Patriarca et al., 2017a), and maritime operations (Praetorius et al., 2016). The applicability of the taxonomy is illustrated using three systems with clearly distinct complexity characteristics: cash withdrawal from an ATM, teaching a university course, and manufacturing operations. These case studies also support a discussion of the implications of the taxonomy for the potentials of resilient systems proposed by Hollnagel (2017).

2. Background

2.1 A socio-technical perspective for the investigation of interactions

In this study, the concept of interactions is explored from a socio-technical perspective where technological, human, social and environmental components cannot be optimized individually (Trist, 1981). These components form a cooperative ensemble, and the overall system performance is mostly a function of their interactions rather than their individual properties (Hollnagel, 1998). Empirical evidence supporting this view has been gathered in a wide range of socio-technical domains (e.g. Akyuz and Celik, 2015).

In particular, studying in detail the interaction between technology and users is crucial to limit unintended consequences (Nielsen, 1990). Over years, automation acquired an increasingly central role, as proved by the large number of different models put forward for studying human-automation interaction (Parasuraman et al., 2000). In this domain, one early attempt of modelling interactions consisted of assigning tasks to machines or to humans following the MABA-MABA logic (men are better at; machines are better at) (Fitts, 1951).

Furthermore, it has been early acknowledged that the inherent complexity of socio-technical systems requires to take into account interactions at different levels of abstraction (Rasmussen, 1985). Such multi-layer structural decomposition model has been further revised through a functional deconstruction approach based on FRAM (Patriarca et al., 2017b). Consequently, a socio-technical perspective implies that the study of interactions has to acknowledge both abstract classification of roles assigned to

social and technical aspects of work, and less abstract interactions to be usable for system modelling (Baxter and Sommerville, 2011). Regarding more concrete aspects of system interactions, Mayer et al. (2014) define different ranges of interactors, either tangible such as physical buttons or knobs, or intangible such as software, graphical user interfaces, gesture or speech-control interfaces.

2.2 Interactions and FRAM

The FRAM is a method for modelling the performance of socio-technical systems, relying on the identification of variability and how it may aggregate within a work domain. It is a viable solution to explore how variability interacts leading to outcomes that are either expected or unexpected, considering interactions at different abstraction levels (Hollnagel, 2012).

The FRAM specifies interactions in terms of the relationship between the Output O of an upstream function, and any other aspect (Input I, Precondition P, Resource R, Control C, and Time T) of a downstream function. An interaction means that an output can (Hollnagel, 2012): trigger the start of another function (I); set a precondition for the start of another function, although this by itself does not start the function (P); increase or decrease as a result of carrying out the function (R); set expectations and thresholds of acceptable performance (C); set time constraints/relationships for performing the function (T). In the FRAM vocabulary, these relationships are referred to as couplings, which are hereafter interpreted as a synonym of interaction.

A core part of applying the FRAM is the assessment of the variability of the outputs of each function. Hollnagel (2012) proposes that an elaborate analysis should account for ten dimensions of variability: timing, duration, distance / length, direction, magnitude, speed, force/power/pressure, object, quantity/volume, and sequence. The simple analysis of variability, which according to Hollnagel may be sufficient in most cases, only accounts for the variability in terms of timing and precision.

3. The process development of the taxonomy

The taxonomy process development was based on contributions from literature, case studies, and feedback from a pool of experts. This process had an iterative and incremental nature (Figure 1), which in principle could continue perpetually.

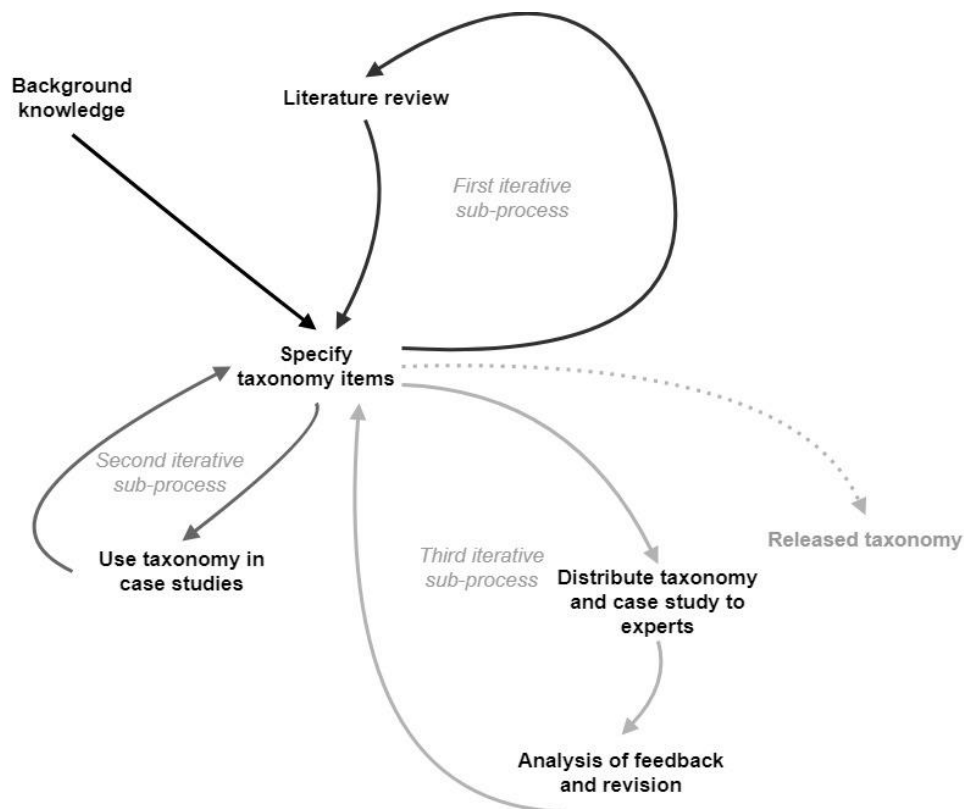


Figure 1. Iterative process for developing the taxonomy

The first iterative sub-process (literature review – items specification) was mainly related to the analysis of literature. We decided to start from Scopus database, which is the largest repository for scientific articles, and select articles which contained “functional resonance analysis method” in title or abstract or keywords. Starting from the contributions obtained from the literature search (76 documents, indexed in Scopus until 30th November 2018), a content analysis was performed to identify pieces of content in relation to interactions in the context of FRAM, systematically labelling the contents. Each paper was examined in terms of the presence or absence of potential attributes of interactions (i.e. the first set of taxonomy items).

The second iterative sub-process (case studies application by authors – items specification) was oriented to the adoption of the taxonomy to two case studies in order to test its applicability in different contexts. The case studies referred to teaching a University course, and forging operations in a manufacturing plant. To assess the reliability of the taxonomy, an inter-reliability criterion has been adopted, i.e. Cohen kappa (Cohen, 1960). Following the interpretation of Landis and Koch (1977), the values of the inter-reliability analyses confirmed an almost perfect agreement (about 0.90) for the case studies.

The third iteration sub-process (case study application by experts – items revision) was aimed to further increase the reliability of the taxonomy. For this purpose, seven international experts (four from Italy, two from Brazil, one from Australia) were invited to use the taxonomy for the same case study. The experts all have more than three years of research experience with the FRAM: one MSc student with a thesis on FRAM, two PhD students, one assistant professor, two associate professors, and one senior researcher.

Except for the MSc student, all the experts had experience both at theoretical level (as confirmed by their authored publications in the topic) and industrial level (confirmed by the projects they managed on FRAM) and joined the international FRAM community (the so called FRAMily), which make them a credible pool of validators.

In this case, a simple process was selected to minimize the background bias of experts, i.e. cash withdrawal. To remove further biases, a video recorded by one of the authors was shared among experts. Every expert was asked to apply the taxonomy to two interactions selected from the FRAM model of cash withdrawal and then answer to two questions: “Is the taxonomy item understandable?” “Is the taxonomy item helpful”, with one of the following choices: none, to a small degree, to a high degree, to a very high degree. The process was conducted through the usage of an ad hoc online spreadsheet, and, where necessary, semi-structured interviews to discuss the assigned values.

4. The taxonomy of interactions

Our operational definition of interaction, which underlies the taxonomy, is as follows: any dependence relationship between two functions in a FRAM model, which does not necessarily involve the exchange of physical or information flows. The taxonomy is comprised of nine criteria: seven original ones, one criterion (i.e. nature of agents) adapted from the FRAM, and another criterion originally proposed by the FRAM (i.e. output nature). The criteria and the descriptors of their performance levels, when applicable, are presented below.

(i) Nature of agents: according to the original FRAM proposal, the agents who perform functions can be humans, technologies or organizations (Hollnagel, 2012). Natural agents (e.g. animals, soil, atmosphere, oceans, etc.) could be another relevant type in some systems. In general, interactions involving human and organizational agents tend to be more variable than those involving only technical agents (Hollnagel, 2012). In a same function there may be a mix of the said types and these can be further sub-divided if necessary. The descriptors for this criterion are then:

- human/individual;
- human/team;
- technology/software;
- technology/hardware;
- natural agent;
- organizational agent.

As for the assessment of variability associated with this criterion, three main levels are proposed:

- Low variability: agents at both ends (i.e. upstream and downstream) usually have the same nature;
- Moderate variability: agents at both ends sometimes have the same nature;
- High variability: agents at both ends usually have a different nature.

(ii) Output nature: in the FRAM, an output is the result of the function, either an entity or a state change – these two are the descriptors of this criterion. An entity corresponds to an output that has a physical nature and is physically transformed as a result of the function. For instance, the output “medication administered”, arising from the function

<administer medication>, might be framed as an entity to the extent that the medication has a physical nature and was physically transformed as a result of being administered.

A state change corresponds to outputs of any nature (i.e. physical and non-physical) that change a non-physical characteristic (e.g. location, information content) as a result of the function. For instance, although “supplies stored” may be an output of the function <store supplies>, the supplies themselves do not physically change as a result of the function – only the location changes. As for the output variability, the two main phenotypes proposed by Hollnagel (2012) are adopted, as follows:

- Timing: on time, too late, too early, not at all;
- Precision: precise, acceptable, imprecise.

(iii) Levelling: this refers to whether the output production volume (i.e. amount of outputs) and mix (i.e. type variations of an output, such as old and young patients in a hospital, transactions types in a bank) vary over the time of the day and day of the week. Unlevelled production (e.g. a surge of trauma patients in an emergency department) tends to stress production resources in certain moments, while these may be idle most of the time (Hopp, 2018). Given its nature, the descriptor of what the criterion looks like and its variability can be merged, as follows:

- Levelled: both mix and volume do not vary substantially over time;
- Moderately levelled: either mix or volume vary substantially over time;
- Unlevelled: both mix and volume vary substantially over time.

(iv) Waiting time: this refers to the time it takes after the output is produced by the upstream function up to its actual use by a downstream function. Thus, what is measured

is the waiting time from output production to consumption. This is the main difference between the waiting time criterion and the time aspect of the FRAM functions. The traditional FRAM aspect is concerned with the time constraints of the function itself (e.g. start time, end time, time pressure) rather than taking the perspective of the waiting time. On the one hand, short waiting time tends to be desirable when the output (e.g. a patient, fresh food) properties can deteriorate in the face of long waiting. On the other hand, long waiting time may be desirable when it means a greater time window for the setup of production resources and problem-solving.

Since what counts as a short or long waiting time is context dependent, the descriptors related to this category are aimed at supporting a formalized recording of information, rather than comparisons between FRAM models. The descriptors are then as follows:

- Tight waiting time;
- Medium waiting time;
- Long waiting time.

The quantitative values of such categories depend on the system at hand, and could be (e.g.): up to 1 minute (tight), up to 1 hour (medium), up to 1 hour and more (long).

Regarding variability, three levels are proposed:

- Low variability: waiting time's values are usually the same;
- Moderate variability: waiting time's values are sometimes the same;
- High variability: waiting time's values are usually different.

(v) Distance: this criterion refers to the physical distance travelled by the output when moving from an upstream to a downstream function. As such, distance is only a relevant

criterion for interactions that involve the flow of outputs among functions. The longer the distance travelled by the output the more it is exposed to the external environment variability. Distance may be a relevant consideration even when the energy flowing is that of electronic signals (Kirilenko et al., 2017).

Similarly to waiting time, what counts as long or short distance is context-dependent. Thus, descriptors related to this category are also aimed at supporting standardized information recording rather than comparisons between FRAM models. The descriptors are:

- Short distance;
- Medium distance;
- Long distance.

Exemplar classification of distance range are proposed: up to 1 m (short); up to 100 m (medium), up to 1000 m and more (long).

In relation to variability, three levels are proposed:

- Low variability: distance's values are usually the same;
- Moderate variability: distance's values are sometimes the same;
- High variability: distance's values are usually different.

(vi) Degree of coupling: this refers to the distinction between tightly and loosely coupled interactions, which can be interpreted as two ends of a continuum. The degree of coupling is closely related to the notion of slack, since the more slack the looser the couplings. Slack is a mechanism for reducing interdependencies and minimizing the possibility of one process affecting another, and thus it makes processes loosely-coupled (Safayeni and

Purdy, 1991). From a FRAM viewpoint, there may be slack functions triggered by the output variability of upstream functions (Saurin and Werle, 2017). Similar to criterion (iii), variability is an integral part of this criterion and therefore the descriptor of what the criterion looks like and its variability can be merged, as follows:

- Tightly-coupled: there are no realistic alternative means to produce and use the upstream output. Neither waiting time nor distance are long;
- Moderately-coupled: there is at least one realistic alternative means to produce and use the upstream output. Either waiting time or distance are long;
- Loosely-coupled: there are two or more realistic alternative means to produce and use the upstream output. Waiting time and distance are long.

(vii) Visibility: this refers to the degree to which the interaction is self-explaining, without the need for verbal communication (Galsworth, 2017). The more visible the interaction, the easier tends to be its monitoring and understanding, thus reducing perceived complexity (Ramasesh and Browning, 2014). The visibility of the output of the upstream function and its status are approached from two perspectives: how visible the output and its status (e.g. late, precise) are to the agents at the downstream function, and how visible the output and its status are to other agents in the environment. An output can be “visible” in a physical sense, but still be invisible from the eyes of an untrained observer. Also, visibility can be obtained through indirect means, such as instrumentation and videos. The descriptors corresponding to visibility are:

- High visibility: the output itself and its status are visible both to the downstream agent and to agents in the environment, and there is no reliance on indirect sources of information;

- Low visibility: neither the output nor its status are visible, from the perspective of both the downstream agent and agents in the environment. There is reliance on indirect sources of information;
- Moderate visibility: any situation in which the previous two descriptors do not hold true.

Regarding variability, the proposed levels are as follows:

- Low variability: visibility levels are usually the same;
- Moderate variability: visibility levels are sometimes the same;
- High variability: visibility levels are usually different.

(viii) Safety and/or security hazards: a hazard is a “condition or object with the potential of causing injuries to personnel, damage to equipment or structures, loss of material, or reduction of ability to perform a prescribed function” (FAA, 2009). There are two dimensions for exploring this category, namely the hazardous properties of the output and the vulnerability of the output to hazards in the environment. There may be either safety implications, when the hazards are unintentionally released and no harm is desired (e.g. occupational accidents), or security implications, when there is an intention to release the hazard and cause harm (e.g. terrorism). This criterion does not account for emergent hazards arising from several interdependent interactions. The descriptors are presented below, and separate assessments should be carried out for safety and security hazards.

- No safety (and/or security) hazards: the output has no relevant hazardous properties and the environment does not pose any significant hazards to the output;

- Either the output is hazardous or the environment poses hazards to the output (safety and/or security);
- The output has hazardous properties and the environment poses significant hazards to the output (safety and/or security).

The variability of the safety/security hazards is mostly linked to where the output and its environment are positioned in a continuum ranging from a technical to a socio-technical system. The closer to a purely technical system (e.g. a product with toxic properties), the less variable the hazards tends to be. As such, the proposed variability levels are as follows:

- High variability: hazard is dynamic, changing over time and under conditions that cannot be easily anticipated;
- Moderate variability: hazard is dynamic, changing over time and under conditions that can be easily anticipated (e.g. hour of the day, day of the week, location, weather);
- Low variability: hazard is static, not changing over time.

(ix) Parallel replications: the traditional FRAM models do not make it clear how many replications a same function has in a given moment in time. For instance, there may be a generic function <administer medications to patients> performed by a generic nurse. However, in reality, there may be several nurses doing the same function in parallel. As such, we propose that, when replication occurs, an estimate of the number of parallel replications should be acknowledged. Furthermore, information on the maximum number of possible parallel replications sheds light on the overall capacity of the system, which can be checked against demand.

More replications create more opportunities for unintended and non-linear interactions. These may be beneficial if there is a possibility of relocating resources if necessary, and thus loosening couplings that otherwise would be tighter – e.g. nurse A who is administering medications to a patient may provide advice on how to administer medications to nurse B who is in the same room caring another patient. The descriptors for this criterion are as follows:

- No parallel replications;
- Medium number of parallel replications;
- Large number of parallel replications.

These descriptors are context-dependent as well, but exemplar values could be: medium (up to ten), large (up to hundreds, and more).

Variability levels are applicable to this criterion, and the levels as follows are proposed:

- Low variability: parallel replication values are usually the same;
- Moderate variability: parallel replication values are sometimes the same;
- High variability: parallel replication values are usually different;

5. The process of applying the taxonomy in the case studies

5.1 Selection of case studies and steps for applying the taxonomy

The steps for applying the taxonomy were the same in all cases, as follows:

Step 1: the development of a FRAM model and a corresponding instantiation of this model, either involving a past event, the present everyday work, or a future scenario;

Step 2: since a FRAM model may have dozens of interactions, it is not practical to apply the taxonomy for all of them. Thus, some interactions should be prioritized. Two prioritization criteria are proposed, namely: interactions involving functions that directly produce the main output of the whole system (e.g. administering classes in the teaching case study); and interactions that have a larger number of upstream and downstream couplings – i.e. functions with higher in-degree and out-degree values;

Step 3: application of the taxonomy for the selected interactions;

Step 4: the proposition of recommendations for influencing the interactions in the desired direction, if necessary. First, the possibility of eliminating interactions should be considered, since this can make the system less vulnerable to unexpected and undesired interactions. Second, it is necessary to appreciate whether the removal of any interaction does not imply in creating compensating interactions, which may bring up their own even worse risks. Third, if the interaction cannot be eliminated, it should be verified whether it is necessary and possible to influence it by design, using the results of applying the taxonomy as a source of improvement opportunities identification.

5.2 Data collection and analysis

The taxonomy was tested in three case studies, which set a basis for its evaluation. The cases represent markedly different systems, thus allowing for the investigation of the applicability of the taxonomy to different contexts. The cases involved: cash withdrawal from an ATM; teaching a University course; and forging operations in a manufacturing plant.

The cash withdrawal is the case study used in the third iterative sub-process (cf. §3.2.c) whose main source of data was the filming of a withdrawal carried out by one of the authors of this paper. The ATM was located within the private premises of a bank branch, and the film was recorded on a weekend, when there was no one else in the facility. This was useful to avoid interruptions and to model a simple situation. Based on watching the film multiple times (about 3 minutes), the usual steps for developing a FRAM model were followed (Hollnagel, 2012). Given the routine nature of this activity, the identification of functions and variability sources was fairly straightforward.

The teaching case refers to one of the two case studies for taxonomy application by the authors of this paper (cf. §3.2.b). It was based on the experience of the authors who have both administered graduate and undergraduate courses for several years. The FRAM model corresponds to an everyday class at the undergraduate level, in the context of the institution of one of the authors. There was no formal data collection in this case, given the deep insider's experience of the researchers. This case study has been selected since it may represent an easy exemplar validation of the taxonomy.

The third example refers to forging operation in a manufacturing plant (second of the two case studies for taxonomy application by the authors of this paper). Data in this case were gathered by means of documents, focus groups, semi-structured interviews, and observations. This example was chosen because the authors have both experience in industrial operations, with one of the authors participating in the original model development, documented in Gattola et al. (2018).

For the three case studies, some interactions were selected for applying the taxonomy based on the previously mentioned criteria and application steps. In the cash withdrawal case study, in case of disagreement, it is reported the most frequent item selected in the taxonomy by the experts (cf. §3.2.c).

6. Results

6.1 The cash withdrawal case

Figure 2 presents the FRAM instantiation for the cash withdrawal case. The function on the top, <monitor transaction>, is performed by the software that controls the operation. Each function carried out by the customer, such as <insert card>, sends an electronic signal that is an input for <monitor transaction>. Then, the software releases the next screen and sets a time limit for performing the next function. Therefore, the output of <monitor transaction> is coupled with the time aspect of other functions.




Figure 2. FRAM instantiation of the cash withdrawal case study. Notes: (i) functions in red are those selected for analysis; (ii) waves inside the hexagons indicate variability in the function's output.

Considering that the interactions shown in Figure 2 are similar, only two were selected for applying the taxonomy (Table 1): I-1 <insert card> - <monitor transaction>, and I-2 <type PIN> - <retrieve cash>.

Table 1. Application of the taxonomy to the ATM cash withdrawal case study. Note: the results according to the descriptors previously presented are in *Italics*.

Taxonomy category / Interaction	I-1: insert card – monitor transaction	I-2: type PIN – retrieve cash
Nature of the agents	<i>Human/Individual – Technology/Software</i> <i>Moderate variability:</i> sometimes there are different ATM models in the same branch, and there are demographic variations in the customers (old vs. young, literate vs. illiterate, blind vs. non-blind)	<i>Human/Individual – Human/Individual</i> <i>Low variability:</i> the same person is doing two consecutive functions, moderated by the function <monitor transaction>
Output nature	<i>State change</i> <i>Acceptable precision:</i> sometimes the ATM does not read the card at the first time it is inserted	<i>State change</i> <i>Imprecise:</i> slips and memory lapses when typing the PIN are common
Levelling	<i>Moderately levelled:</i> demand for the ATM increases around noon, while the mix probably does <u>not change</u> overtime	<i>Moderately levelled</i>
Waiting time	<i>Tight waiting time:</i> there is virtually no significant delay after inserting the card and the detection of this action by the computer <i>Low variability</i>	<i>Tight waiting time. Low variability.</i> After typing the PIN, a few seconds pass up to the release of cash. If not removed briefly, the cash is pulled back into the ATM.
Distance	<i>Short distance. Low variability.</i>	<i>Short distance. Low variability.</i>
Degree of coupling	<i>Tightly-coupled:</i> there is no alternative means of producing and using this output after starting the upstream function	<i>Tightly-coupled</i>
Visibility	<i>High visibility:</i> from the perspective of both the environment and downstream agent (i.e. computer), visibility is high, since the computer promptly detects the card insertion <i>Low variability</i>	<i>Moderate visibility:</i> although the upstream and downstream agents are the same, the PIN appears on the screen only as **** <i>Low variability</i>
Safety and/or security hazards	<i>No safety hazards</i> <i>No security hazards</i> <i>Low variability</i>	<i>No safety hazards</i> <i>Either the output or its environment has security hazards:</i> If a large amount of cash is removed at once, this may call the attention of people around. However, this is unlikely given the daily withdrawal limits set by the bank. <i>High variability:</i> the location of the ATM (e.g. region within a given city) has an influence on the security hazards
Parallel replications	<i>No parallel replications:</i> in the specific instantiated scenario, there were no other people using the existing three neighbouring ATM <i>Moderate variability:</i> the number of replications depends on the size of the branch	<i>No parallel replications</i> <i>Moderate variability</i>

The main lessons learned from this case study are as follows: (i) redundant ATM on standby offers an alternative for transforming a tightly-coupled system of interactions into

a loosely-coupled one – e.g. using a neighbouring ATM if there is a technical failure in any of them; (ii) variability in terms of nature of agents may be a drawback, in face of an inflexible software/hardware; and (iii) need for visibility is contingent. On the one hand, the lack of visibility of the PIN made sense in I-2. On the other hand, it could be beneficial if visibility was given to the time available for the consumption of the output by the downstream function –e.g. the output of <type amount> will vanish if not consumed within a certain time limit set by the ATM software.

Improvement opportunities arising from using the taxonomy may be highlighted, namely: (i) to develop a more flexible software/hardware, which can be adaptable to different profiles of users – e.g. touch screens that allow for enlarging the characters, use of icons as a support for non-native speakers and illiterate people; and (ii) to give visibility for the time available for performing the next function – e.g. by posting a countdown on the screen. Besides, some interactions can be eliminated due to the introduction of new human-computer interfaces in the near future – e.g. the client could be automatically identified based on his/her iris, eliminating the need for cards. Elimination of the keyboard is less likely to be useful, since it provides privacy that would be compromised if voice commands were used in public spaces.

6.2 The teaching case

Figure 3 presents the FRAM instantiation for a scenario of teaching a course at undergraduate level. Classes typically occur once or twice a week, for groups of 30 students. According to the defined boundaries, the starting function is <go to the University - professor>. Once arriving at the building entrance hall, which is located on

the ground floor (there are 7 floors), the function <collect classroom keys> is performed. The keys of all classrooms are stored in the reception desk, in which a security guard makes written records of who collects each key and when. The output of the collect keys function is often late because the keys are not available. This usually occurs either because the keys were not returned to the security guard by the person who was in the classroom, or because there is still someone in there. This triggers the need for looking for the keys elsewhere and may delay the downstream functions. After having access to the classroom, the function <switch on computer and other equipment – e.g. slide projector> is conducted. As a precondition for this function, the output of <maintain computers and other equipment> should be precise and on-time. There are two main inputs for the core adding-value function <administer class>: the upload of slides on the computers and the presence of a minimum number of students in class – the start can be delayed a few minutes since many students can be late. The downstream boundary of the model is the function <apply knowledge>, which is performed by the students and is influenced by the variability of all upstream functions, besides a number of contextual factors not encompassed by the model.

Table 2 presents what three interactions look like in light of the taxonomy. The selected interactions (I) are: I-1 <give keys back to security> - <collect classroom-keys>, I-2 <administer class> - <apply knowledge>, and I-3 <maintain computers and other equipment> - <switch on computers and other equipment>. These interactions involve the main human agents participating in the system (professor, students, IT staff, and security), as well as technological artefacts.




Figure 3. FRAM instantiation of the teaching case study.

Table 2. Application of the taxonomy to the teaching case study.

Taxonomy category / Interaction	I-1: return keys – collect keys	I-2: administer class – apply knowledge	I-3: maintain computer and other equipment – switch on computer and other equipment
Nature of the agents	<i>Human/Individual – Human/Individual</i> <i>Moderate variability:</i> the person who gives back the keys may not be the same who collected the keys.	<i>Human/Individual – Human/Individual</i> <i>Moderate variability:</i> sometimes knowledge is applied by students as part of teamwork	<i>Human/Team – Human/Individual</i> <i>Low variability</i>
Output nature	<i>State change</i> <i>Imprecise and not at all:</i> sometimes the person in charge of the keys forgets to give these back	<i>State change</i> <i>Acceptable precision:</i> the quality of the lecture depends on a number of factors – e.g. fatigue	<i>Entity</i> <i>Too late and acceptable precision:</i> it is often delayed, problems not definitely solved
Levelling	<i>Unlevelled</i>	<i>Levelled:</i> the output of <administer class> is stable in terms of mix and volume. However, the use of this as an input for <apply knowledge> is irregular. There may be accumulation of knowledge waiting to be used.	<i>Moderately levelled:</i> frequency of corrective maintenance varies over time. Mix is more stable.
Waiting time	<i>Long</i> <i>High variability:</i> the waiting time varies from a few minutes to several hours, depending on the schedule of classroom occupation	<i>Long</i> <i>High variability:</i> the waiting time between the output production and its use can vary from days to months or even years.	<i>Medium</i> <i>High variability:</i> it can vary from minutes to days, since the corrective maintenance activities occur on demand, and there may be other priorities
Distance	<i>Short</i> <i>Low variability</i>	<i>Long</i> <i>High variability</i>	<i>Short</i> <i>Low variability</i>

Degree of coupling	<i>Moderately coupled:</i> there is a spare key set in each department reception area.	<i>Loosely-coupled:</i> there are many possible ways of using the output of <administer class>, and there is a <u>significant slack in terms of time</u>	<i>Loosely-coupled:</i> there are several other desktops in the classrooms, which could be a replacement
Visibility	<i>Low visibility:</i> there is a low visibility for the downstream agent, since the person who collects the key does not know in advance whether or not the keys were returned. High visibility for agents in the environment. <i>Low variability</i>	<i>High visibility:</i> there is a high visibility for the downstream agents, since the students attend the class. <i>Moderate variability:</i> although the class is visible in a physical sense, there may be wide variations regarding students perceptions	<i>Low visibility:</i> the downstream agent (professor) has no visual cues on the maintenance status of computers and equipment. <i>Low variability</i>
Safety and/or security hazards	<i>No safety hazards</i> <i>No security hazards</i> <i>Low variability</i>	<i>No safety hazards</i> <i>No security hazards</i> <i>Low variability</i>	<i>Either the output or its environment have safety hazards:</i> inadequate maintenance of electrical equipment can pose safety hazards <i>Moderate variability:</i> it depends on the age of the equipment and the nature of the faulty maintenance <i>No security hazards</i>
Parallel replications	<i>No parallel replications</i> <i>Moderate variability</i>	<i>No parallel replications:</i> although the same course may be given by the same or another professor (to another class), it does not occur in parallel <i>Low variability</i>	<i>Medium number of parallel replications:</i> maintenance activities in parallel are common, given that there are several maintenance staff <i>Moderate variability:</i> it depends on the variation in demand for maintenance

Some insights from Table 2 may be highlighted, namely: (i) visibility was high only when the downstream agent was physically present in the same environment where the upstream output was produced (i.e. in the case of I-2); and (ii) replications that occur in different moments in time (e.g. the same class to a similar group of students, at another time) may also offer an opportunity for loosely-couplings as well as for learning from experience.

Re-design recommendations start by considering the elimination of unnecessary complexity. For instance, I-1 (return keys – collect keys) could be eliminated if classrooms were kept permanently unlocked during business hours. While this could create security threats, as well as new functions (e.g. an administrative employee opening

the rooms at the beginning of the day and closing them at the end of the day), these would have a low frequency, and probably a complexity reduction net effect. Furthermore, the visibility of some interactions can be enhanced – e.g. by posting a schedule of planned versus actual preventive maintenance on public display.

6.3 The manufacturing case

This case study is inspired by a previous research conducted in a metalworking company producing power-tools accessories (Gattola et al., 2018). The plant produces a wide variety of products, which can be grouped into three main families, depending on the dimension. The main production process is divided into six phases: turning, milling, forging, tempering, sandblasting and packaging.

The focus of this case study is on the forging operation, which is the most critical part of the production process. At the beginning of every work-shift, the operator performs a conformity check following the so-called 6S check, which includes routine functions such as <control oil level>, <control cooling water filter>, <control collective and individual barriers>, <clean machine floor>, and <clean machine panels>. Production can only start after the forging machine is properly setup. The operator's functions are then <turn on the machine> and after positioning the raw material in the belt of the machine, she <move and heat the raw material>, and <forge the semi-finished product> by means of a standard load automatically provided by the machine. In case of a technical problem encountered in the setup procedure, the operator in charge of the process shall contact maintenance technicians. This last process part represents the core of the proposed instantiation (see Figure 4).




Figure 4. FRAM instantiation of the manufacturing case study.

Table 3. Application of the taxonomy to the manufacturing case study.

Taxonomy category / Interaction	I-1: Control oil level – Contact maintenance team	I-2: Execute maintenance – Turn on the machine
Nature of the agents	<p><i>Human/Individual – Human/Individual. Low variability.</i></p> <p>The operator in charge of the 6S activity (control oil level) is responsible for both controlling oil levels and contacting the maintenance technician, if the routine check is not positive.</p>	<p><i>Human/Team – Human/Individual. Low variability.</i></p> <p>The maintenance team executes maintenance. The operator is responsible for the downstream action. The process information in this interaction refers to the flow from team’s actions to individual’s actions.</p>
Output nature	<p><i>State change.</i> The output oil level controlled is a state change because the oil is not changing as a result of the function, but it rather has a different dimension.</p> <p>The output may be of <i>acceptable precision</i>, since the oil level is visually measured in terms of centimetres of oil on the end of the dip stick inserted. It may also be performed <i>too early</i>, when the machine is not yet cold.</p>	<p><i>Entity.</i> The output machine maintained has a physical nature which adds a key value step to the process of maintenance management.</p> <p>The output may be <i>imprecise</i> with respect to maintenance procedures. It may be <i>too late</i> depending on the delay from the original request.</p>
Levelling	<p><i>Levelled</i></p> <p>The number of maintenance requests is reasonably low, and basically the same over time, as confirmed by the fairly constant time between maintenance reported in the historic data.</p>	<p><i>Unlevelled</i></p> <p>Both mix (type of maintenance intervention) and volume (number of man-hours required for the intervention) vary for each time the interaction is activated.</p>
Waiting time	<p><i>Tight waiting time. Moderate variability.</i></p> <p>The interaction is generally performed in a short time interval (few minutes) to allow a prompt maintenance intervention. There is moderate variability, since in case of strict production plans it could be required to act</p>	<p><i>Medium waiting time. High variability.</i></p> <p>The intervention are generally performed within 1 hour (medium waiting time), but there is high variability due to the variability of required maintenance actions, which usually is assessed only when the team reach the machine.</p>

	promptly, or vice versa, deferring it to allow the operator accommodating other priority requests.	
Distance	<i>Medium distance. Low variability.</i> The operator has to call the maintenance from the central office, which is located less than 50 meters away from the machine.	<i>Short distance. Low variability.</i> The distance is short and not relevant, since the maintenance action is conducted close to the machine.
Degree of coupling	<i>Loosely-coupled.</i> There are two alternative means of performing the interaction, i.e. calling the maintenance technicians through the operator's personal phone or contact them by e-mail.	<i>Moderately-coupled.</i> Generally there is at least one alternative way to perform the interaction, depending on the components. The coupling level of this interaction can be represented, as a proxy measure, by the number of maintenance procedures to perform the task, which can be stricter (requiring certified technician, or specific tools) for some type of intervention.
Visibility	<i>Moderate visibility. High variability.</i> From the perspective of the downstream agent (i.e. maintenance technicians), visibility is moderate, since the operator does not necessarily refer the issues he is facing, due to her local understanding of the situation. It is highly variable, because it generally depends on the experience of the operator, and on the time available to properly assess the scenario.	<i>Moderate visibility. Moderate variability.</i> The downstream agent generally has indirect sources of information. If the intervention is conducted during her work shift, she acknowledges the completion of the maintenance activity on person, otherwise the operator has to check the daily activity report in the main office. Variability emerges from such possible scenarios.
Safety or security hazards	<i>No safety (and security) hazards. Low variability.</i> The action is pretty straightforward with very limited potential for any type of risks for the operator. There is no security implication.	<i>The output has hazardous properties and the environment poses significant hazards to the output. High variability.</i> <i>No security hazards.</i> Even if there is no security implications at this level, there could be safety issues, since an imprecisely executed maintenance may jeopardise the operator's safety. It could be highly variable since there could be several combinations of events leading to critical consequences.
Parallel replications	<i>Low number of parallel replications. Moderate variability.</i> The same interaction is normally performed one time for work shift, but can be performed multiple times if different production plans require it.	<i>Medium number of parallel replications. Moderate variability.</i> The same interaction is performed multiple times for the same machine, with the potential for multiple interventions, by different teams for different types of faults to be diagnosed and managed.

From the usage of the taxonomy, it is possible to define some recommendations, on both technical and management aspects. For example, one criticality that emerges from both I-1 and I-2 refers to waiting time between making contact with maintenance team and the subsequent maintenance execution. In this case, referring to I-1 <control oil level> it is requested that the operator checks manually at the beginning of her work shift the machine oil level, and make contact only in case the level is not satisfactory. A recommendation in this case would suggest inviting the operator in taking note of the assessed oil level, as well as of the other 6S checks. The reported data would be useful to

feed an algorithm which combines the production plan and the reported measures to better schedule the maintenance interventions, loosening the time pressure for downstream interactions.

As a consequence, this process change would also imply benefits with respect to the visibility and on the levelling of the interactions, allowing the maintenance team to have more formal data for setting up properly the intervention. Such change would imply the shift towards a dynamic condition-based maintenance, relying on prognostic models. Following the analysis of the taxonomy's items safety/security hazards, it is also recommended to promote toolbox meetings among maintenance teams in order to discuss the potential safety and production criticalities emerging from imprecise execution.

7. Discussion

7.1 Taxonomy assessment

The feedback from the seven experts indicated, at least for those who already have experience with traditional FRAM, that the taxonomy items are acceptably easy to be interpreted and they do provide a helpful approach to deal with the complexity of socio-technical systems (Figure 5). It can be inferred from this that, while the taxonomy increases the complexity of the FRAM analysis, it adds value to cope with the complexity of the representation.

A useful feedback from two experts refers to the “safety hazard” and “security hazard” criteria. The experts recommended the use of this criteria to be accompanied by a punctual

definition of what safe/unsafe (secure/unsecure) means in the specific case study under examination, rather than simply use the related label. Furthermore, one expert pointed out that the criteria “Levelling” and “Parallel Replications” may be not clearly understandable (cf. Figure 5). For such criteria, the descriptors were furtherly refined in the third iterative sub-process, in order to convey them in sufficiently abstract terms, so as they could be applicable across domains, and at the same time being precise enough for reliable assessments.

The assessment also pointed out that analysing some features of an interaction may be unnecessary under certain circumstances. This comment emerges from Figure 5, where some experts assessed the criteria “Nature of agents variability”, “Safety (Security) hazard value/variability”, and “Parallel Replications” to be helpful at a small degree. For some systems such aspects would not be relevant, indicating that the taxonomy is intended to be flexible and better accommodate the analysts’ needs for the case at hand. The taxonomy may be used partially and in an ad-hoc manner, to the extent that it has utility for the analyst.



Figure 5. Feedback by the seven experts involved in the study.

7.2 Implications for resilience engineering

In principle, the taxonomy can contribute to the operationalization of the four potentials of resilient systems defined by Hollnagel (2017), as discussed below:

(i) The potential to *respond* implies knowing what to do, responding to regular and irregular changes and opportunities (Hollnagel, 2017). In this respect, the system re-design recommendations discussed in the case studies illustrate how the taxonomy application can give rise to responses to either undesired or unsatisfactory characteristics of the interactions;

(ii) The potential to *monitor* implies knowing what to look for, monitoring what could seriously affect performance in the near term, positively or negatively (Hollnagel, 2017). The taxonomy criteria may give rise to some metrics worth monitoring. For instance, the monitoring of the interaction levelling implies in the need for monitoring the rate and mix of output production at the upstream and consumption/use at the downstream function. This data might be useful for the re-design of the production resources – e.g. increasing capacity by adding more ATM;

(iii) The potential to *learn* implies knowing what has happened, acquiring the right lessons from the right experience (Hollnagel, 2017). The FRAM can be interpreted as a learning platform (Clay-Williams et al., 2015), applicable both for modelling past events (e.g. accidents) as well as for a risk analysis, looking into future scenarios. Both situations offer learning opportunities, which can be enriched by a structured recording of the taxonomy application;

(iv) The potential to *anticipate* implies knowing what to expect, preparing for developments further into the future, such as disruptions, constraints or opportunities (Hollnagel, 2017). This potential may benefit from the use of the taxonomy jointly with FRAM models focused on risk analysis, as well as from applying the variability descriptors. The variability of the interaction could be assessed considering longer time

horizons into the future – e.g. which would the expected variability of the interaction within one year into the future?

8. Conclusions

The proposed taxonomy contributes to a deep understanding of the functional interactions in socio-technical systems. The two dimensions of the taxonomy descriptors (i.e. output characteristics from the viewpoint of the downstream function, and variability of the interaction) support the development of a structured database for recording the results of the analysis.

The emphasis on describing the interactions and their variability (instead of being limited to the variability of the outputs) is a distinctive taxonomy's feature in relation to the original FRAM. For example, the variability according to the “waiting time” criterion depends on the match between the rate of output production at upstream and the rate of the same output use by the downstream function. Similarly, the description of the criterion “distance”, and its corresponding variability, depends on the interaction between upstream and downstream function – i.e. a same upstream output may have different distance description and variability from the perspective of different downstream functions. As such, the taxonomy is a complementary analytical tool to the FRAM, by shedding light on system aspects that are implicit in FRAM models. **Indeed, the taxonomy forces the FRAM model developers to make it explicit their assumptions on the nature of the interactions in the scenario considered. This can contribute to the development of more realistic models.**

The usefulness of the taxonomy was demonstrated through the analysis of three case studies in which work system re-design opportunities were identified. These opportunities may be logically connected to the four potentials of resilient systems, thus making a link between the taxonomy and re-design actions consistent with resilience engineering. Furthermore, the case studies suggest that all combinations between the taxonomy's categories can be possible, which provides empirical evidence of the need for uncovering what is beneath the linear/non-linear versus tight/loose couplings taxonomy. The variability descriptor is also a recognition of the need for evaluating the interactions over time, instead of static snapshots.

However, any taxonomy is a social construct and it is not definitive (Parasuraman et al., 2008). As such, new conceptual and technological developments may imply in the need for revising the taxonomy criteria and its descriptors. Another limitation refers to the high number of interactions that may exist in FRAM models, which does not make it practical the full application of the taxonomy for all interactions. In order to overcome this limitation, two possibilities are the prioritization of some interactions for full taxonomy application and a partial application of the taxonomy. Furthermore, the taxonomy reliance on the FRAM can be seen as a limitation to the extent it may need to be adapted in order to be compatible with other modelling tools. Lastly, although the taxonomy was tested in different scenarios, its full generalizability depends on its application to other contexts.

Some opportunities for future studies can be mentioned, namely: (i) to develop a metric to evaluate the complexity of FRAM models, based on the assignment of scores to the descriptors; (ii) to apply the taxonomy to a wide range of systems, in order to identify patterns that could set a basis for standardized taxonomies of socio-technical systems;

and (iii) to adapt the taxonomy to other system modelling approaches, such as causal-loop-diagrams and social network analysis.

References

Akyuz, E., Celik, M., 2015. Application of CREAM human reliability model to cargo loading process of LPG tankers. *J. Loss Prev. Process Ind.* 34, 39–48.

<https://doi.org/10.1016/j.jlp.2015.01.019>

Baber, C., Stanton, N. A., Atkinson, J., McMaster, R., & Houghton, R. J. 2013. Using social network analysis and agent-based modelling to explore information flow using common operational pictures for maritime search and rescue operations. *Ergonomics*, 56(6), 889-905.

Baxter, G., Sommerville, I., 2011. Socio-technical systems: From design methods to systems engineering. *Interact. Comput.* 23, 4–17.

<https://doi.org/10.1016/j.intcom.2010.07.003>

Bolbot, V., Theotokatos, G., Bujorianu, L., Boulougouris, E., Vassalos, D. 2019.

Vulnerabilities and safety assurance methods in Cyber-Physical Systems: A comprehensive review. *Reliability Engineering and System Safety*, 182, 179-193.

Cilliers, P., 1998. *Complexity and Postmodernism: Understanding Complex Systems*. Routledge, London.

Clay-Williams, R., Hounsgaard, J., Hollnagel, E., 2015. Where the rubber meets the road: using FRAM to align work-as-imagined with work-as-done when

implementing clinical guidelines. *Implement. Sci.* 10, 125.

<https://doi.org/10.1186/s13012-015-0317-y>

Clegg, C., 2000. Sociotechnical principles for system design. *Applied Ergonomics*, 31, 463-477.

Cohen, J., 1960. A Coefficient of Agreement for Nominal Scales. *Educ. Psychol. Meas.*

<https://doi.org/10.1177/001316446002000104>

ElMaraghy, H.; AlGeddawy, T.; Samy, S.; Espinoza, V. 2014. A model for assessing the layout structural complexity of manufacturing systems. *Journal of Manufacturing Systems*, 33, 51-64.

FAA, 2009. Risk management handbook 8083, Flight Standards Service, 2009.

Fitts, P.M., 1951. Human engineering for an effective function allocation. National Research Council, Washington, DC.

Galsworth, G. D. 2017. Visual workplace: visual thinking. CRC Press.

Gattola, V., Patriarca, R., Tomasi, G., Tronci, M., 2018. Functional resonance in industrial operations: A case study in a manufacturing plant. *IFAC-PapersOnLine* 51, 927–932. <https://doi.org/10.1016/j.ifacol.2018.08.489>

Hollnagel, E., 2017. Safety-II in practice: developing the resilience potentials. Taylor & Francis.

Hollnagel, E., 2012. FRAM: The Functional Resonance Analysis Method - Modelling Complex Socio-technical Systems. Ashgate.

Hollnagel, E., 1998. The Cognitive reliability and Error Analysis Method. Elsevier, Oxford, UK.

Hollnagel, E., Woods, D., 2005. Joint Cognitive Systems: Foundations of Cognitive Systems Engineering. Taylor & Francis, Boca Raton.

Hopp, W. J. 2018. Positive lean: merging the science of efficiency with the psychology of work. *International Journal of Production Research*, 56(1-2), 398-413.

Houghton, R. J., Baber, C., McMaster, R., Stanton, N. A., Salmon, P., Stewart, R., & Walker, G. 2006. Command and control in emergency services operations: a social network analysis. *Ergonomics*, 49(12-13), 1204-1225.

Kirilenko, A., Kyle, A. S., Samadi, M., & Tuzun, T. 2017. The Flash Crash: High-frequency trading in an electronic market. *The Journal of Finance*, 72(3), 967-998.

Klockner, K., & Toft, Y. 2018. Railway accidents and incidents: Complex socio-technical system accident modelling comes of age. *Safety science*, 110, 59-66.

Landis, J.R., Koch, G.G., 1977. The Measurement of Observer Agreement for Categorical Data. *Biometrics*. <https://doi.org/10.2307/2529310>

Maguire, M., 2014. Socio-technical systems and interaction design - 21st century relevance. *Appl. Ergon.* 45, 162–170. <https://doi.org/10.1016/j.apergo.2013.05.011>

Mayer, S., Tschofen, A., Dey, A.K., Mattern, F., 2014. User interfaces for smart things - A generative approach with semantic interaction descriptions. *ACM Trans. Comput. Interact.* 21, 1–25. <https://doi.org/10.1145/2584670>

Nielsen, J., 1990. A meta-model for interacting with computers. *Interact. Comput.* 2, 147–160. <https://doi.org/90/010147-14>

Olsen, N. S., & Shorrock, S. T. 2010. Evaluation of the HFACS-ADF safety classification system: inter-coder consensus and intra-coder consistency. *Accident Analysis & Prevention*, 42(2), 437-444.

Parasuraman, R., Sheridan, T.B., Wickens, C.D., 2008. Situation awareness, mental workload, and trust in automation: Viable, empirically supported cognitive engineering constructs. *J. Cogn. Eng. Decis. Mak.* 2, 140–160.

Parasuraman, R., Sheridan, T.B., Wickens, C.D., 2000. A Model for Types and Levels of Human Interaction. *IEEE Trans. Syst. Man Cybern. Part A Syst. Humans* 30, 286–297.

Patriarca, R., Di Gravio, G., Costantino, F., 2017a. A Monte Carlo evolution of the Functional Resonance Analysis Method (FRAM) to assess performance variability in complex systems. *Saf. Sci.* 91, 49–60. <https://doi.org/10.1016/j.ssci.2016.07.016>

Patriarca, R., Bergström, J., Di Gravio, G., 2017b. Defining the functional resonance analysis space: Combining Abstraction Hierarchy and FRAM. *Reliab. Eng. Syst. Saf.* 165. <https://doi.org/10.1016/j.ress.2017.03.032>

Perrow, C., 1984. *Normal Accidents - Living with High-risk Technologies*. Harper Collins Publishers.

Praetorius, G., Graziano, A., Schröder-Hinrichs, J.-U., Baldauf, M., 2016. FRAM in FSA—Introducing a function-based approach to the formal safety assessment framework, *Advances in Intelligent Systems and Computing*. https://doi.org/10.1007/978-3-319-41682-3_34

Ramasesh R V., Browning TR., 2014. A conceptual framework for tackling knowable unknown unknowns in project management. *J. Oper. Manag.* 32:190–204. [doi:10.1016/j.jom.2014.03.003](https://doi.org/10.1016/j.jom.2014.03.003).

Rasmussen, J., 1985. The role of hierarchical knowledge representation in decisionmaking and system management. *IEEE Trans. Syst. Man. Cybern. SMC-* 15, 234–243. <https://doi.org/10.1109/TSMC.1985.6313353>.

Safayeni, F., Purdy, L., 1991. A behavioral case study of just-in-time implementation. *Journal of Operations Management*, 10 (2), 213-228.

Saurin, T. A.; Werle, N. J. B. 2017. A framework for the analysis of slack in socio-technical systems. *Reliability Engineering & System Safety*, 167, 439-451.

Trist, E.L., 1981. *The evolution of socio-technical systems: A conceptual framework and action research program*. Ontario, Canada.

Trist, E.L., Bamforth, K.W., 1951. Some Social and Psychological Consequences of the Longwall Method of Coal-Getting: An Examination of the Psychological Situation and Defences of a Work Group in Relation to the Social Structure and Technological Content of the Work System. *Hum. Relations* 43–38, 3–38.
doi:10.1177/001872675100400101

Walker, G.H., Stanton, N.A., Salmon, P.M., Jenkins, D.P., Rafferty, L., 2010. Translating concepts of complexity to the field of ergonomics. *Ergonomics* 53 (10), 1175-1186.

Wiegmann, D. A., & Shappell, S. A. 2001. Human error analysis of commercial aviation accidents: Application of the Human Factors Analysis and Classification System (HFACS). *Aviation, space, and environmental medicine*, 72(11), 1006-1016.

Wilson, J., 2014. Fundamentals of systems ergonomics/human factors. *Applied Ergonomics* 45, 5-13.