Diagnostic Quality Problem Solving: A Conceptual Framework and Six Strategies

JEROEN DE MAST
INSTITUTE FOR BUSINESS AND INDUSTRIAL STATISTICS OF THE UNIVERSITY OF AMSTERDAM

Diagnostic problem solving, which is the task of discovering causal explanations for unwanted effects, is an important element of problem solving. This paper contributes a conceptual framework for the generic process of diagnosis in quality problem solving by identifying its activities and how they are related. It then presents six strategies that structure the diagnostic process by suggesting a certain sequence of actions and techniques. The paper analyzes when each of these strategies is likely to be effective and how it may help in making the diagnostic process more efficient. Finally, the paper proposes and motivates a generic sequence of stages in diagnosing quality problems.

The framework offers a scientific basis for studying and evaluating problem-solving methodologies such as Six Sigma's DMAIC model, Kepner and Tregoe's problem analysis method, and Shainin's system. For the practitioner, the framework clarifies the rationale for many problem-solving techniques offered in courses and textbooks. The paper also offers indications and contra-indications when techniques are promising, and demonstrates how they fit together in a coherent strategy.

Key words: diagnosis, DMAIC, problem solving, root-cause analysis

INTRODUCTION

Problem solving in the face of quality, reliability, and performance problems has been and still is an important task in operations management (for example, Balakrishnan et al. 1995; Ho and Sculli 1997; MacDuffie 1997). In recent years, large numbers of professionals have been trained in advanced problem-solving methodologies, and the Six Sigma phenomenon in particular has spurred a flood of courses and textbooks on problem solving following its define-measure-analyze-improve-control (DMAIC) model (Chakravorty 2009; de Mast and Lokkerbol 2012). The 1980s and 1990s saw the emergence of problem-solving approaches such as the Shainin System (Shainin 1993; Steiner, MacKay and Ramberg 2008) and Kepner and Tregoe's (1997) problem analysis method, still widely taught and applied in industry. Simpler models, such as the plan-do-check-act (PDCA) and eight disciplines (8D) models, are also generally applied, as well as techniques such as root cause analysis, brainstorming, 5 Whys, and the cause-and-effect diagram.

The discovery of the causes of a problem is called diagnosis, and it is an essential element of problem solving. This paper contributes a conceptual framework for the diagnosis process in quality problem solving. It then presents six strategies that structure the diagnostic process using a certain sequence of actions and techniques. The paper analyzes when each of these strategies is likely to be effective and how it may help in making the diagnostic process more efficient. The focus is on the diagnosis of nonroutine problems, that is, problems that are novel to the problem solver.

The motivation for the author's study is the observation that the scientific underpinning for many of the earlier mentioned methods and techniques is weak. Further, practitioners would like to identify the causes of quality problems with as little effort and as fast as possible. To this end, they need a strategy that specifies when and how techniques should be applied, what data to gather, and in general, what steps to take. Typical books aimed at practitioners, including most accounts of Six Sigma's DMAIC method (such as Breyfogle 2003; Pyzdek 2003), describe a large number of techniques and methods for finding the causes of problems, but without much coherence or structure. The little strategic support that is offered for controlling the efficiency of the diagnostic process often lacks operationality or is tenuous (de Mast and Lokkerbol 2012). The systems promoted by Shainin (Steiner, MacKay, and Ramberg 2008) and Kepner and Tregoe (1997) integrate techniques into diagnostic strategies, but do not provide a scientific understanding of the principles and rationales of problem diagnosis. Moreover, these methods are rather strongly tied to a single diagnostic strategy (de Mast 2011), referred to as branch-and-prune in this paper.

The conceptual framework for problem diagnosis provides a strong scientific basis, allowing quality engineers to further develop problem-solving methods. Also, it facilitates a better understanding and critical evaluation of existing methods. For practitioners, the analysis offers soundly based support for choosing between strategies and techniques.

In the next section, the author presents his conceptual model of the generic diagnostic process. In the subsequent section, he presents six strategies for diagnosis and proposes a heuristic guideline for choosing among them in practice. In the concluding section he discusses some ramifications of the framework for methodological support for problem solving on the shop floor, and suggests directions for future research.

CONCEPTUAL MODEL OF THE DIAGNOSTIC PROCESS

Definition of the Subject of Study

The task of discovering a causal explanation for unwanted effects is called diagnostic problem solving, or diagnosis (Smith 1988). It is one of the core tasks in problem solving (Smith 1988), and is often the basis for the subsequent design of a solution. The process is called the "Diagnostic Journey" (Juran 1998), and it is the function of the analyze phase in Six Sigma's DMAIC method (de Mast and Lokkerbol 2012).There is substantial scientific literature on diagnosis in fields such as troubleshooting of devices (for example, Morris and Rouse 1985; Davis and Hamscher 1988), medical diagnosis (Pople 1982; Norman 2005), artificial intelligence (A.I.) (Chittaro and Ranon 2004; Torasso and Torta 2005), and A.I. in medical diagnosis (Keravnou and Washbrook 1989; Lucas 1997). The quality and industrial engineering fields have produced many accounts in the practitioners' literature, as well as an occasional discussion in the scientific literature (for example, Wagner 1993; Smith 1998; de Mast 2011). The stance taken in this paper is that one can learn a lot from the scientific advances in other fields. The author offer a synthesis of strands and fragments of research in a variety of disciplines, and translates them into a coherent whole for quality engineering. Characteristics of the diagnostic process vary, however, from one context to another, and first he develops a solid description of the sort of diagnosis in quality problem solving that he aims to study.

For the diagnosis of routine problems, a list of known fault types is given to the diagnostician a priori, and diagnosis boils down to selecting from this list the most likely cause. This type of diagnosis is a classification problem of inferring a predefined fault type from observed symptoms. Such tasks have been well studied in A.I., troubleshooting, and medicine (for example, Custers, Regehr, and

Norman 1996; Venkatasubramaniana et al. 2003a; and Torasso and Torta 2005). Solving *nonroutine* problems, to the contrary, often involves the discovery of altogether new fault mechanisms, and the set of potential fault types is not given a priori. This paper, therefore, studies diagnosis as a discovery process, rather than as a classification task.

The author does not assume that all relevant data and domain knowledge are given to the diagnostician a priori. He studies the entire iterative process including diagnostic reasoning (how to get from data to a diagnosis) and data and knowledge acquisition (what data to gather). The intended result of diagnosis in quality problem solving is a causal explanation of problematic behavior that is useful for the subsequent design of a solution. This differs from troubleshooting or identifying the malfunctioning component or connection in a device, without a full explanation of what has happened and what is wrong with the component or connection in question (Wagner 1993).

Diagnosis as a Search Process Through a Problem Space

The A.I. and troubleshooting fields have produced a limited number of problem diagnosis conceptualizations, such as Rasmussen (1981) and Clancey (1988), but none of them are directly applicable to the diagnosis of nonroutine quality problems as described previously. The author's model uses the highly influential concepts of Newell and Simon (1972), who characterize problem-solving processes as a search through a state space driven by operators. From Keravnou and Johnson (1989) he borrows the idea to characterize the states in the diagnostic search in terms of hypotheses and their status. Also from Keravnou and Johnson (1989) he takes the idea that operators are actions that modify the status of hypotheses, thus progressing the search from one state to the next.

The proposed model is summarized in Figure 1. The process of diagnostic problem solving is conceptualized as a search through the field of potential causal explanations, until an explanation is found that is both accepted as true and sufficiently specific. These candidate explanations in the diagnostician's mind are called hypotheses. It is a near universal finding that expert diagnosticians use working hypotheses to bring focus to data gathering and inquiry (for example, Pople 1982; Boreham 1986). During the search, the diagnostician generates, rejects, and accepts hypotheses on the basis of the available domain knowledge and observations. At any moment, there will be a particular set of hypotheses under consideration, and sets of hypotheses already rejected or accepted. These hypotheses and their status (accepted, rejected, or under consideration) define a state in the search process.

The search progresses from one state to the next as the diagnostician performs three types of actions: hypothesis generation, testing, and evaluation. These three operators represent the basic forms of diagnostic inference: abductive, deductive, and inductive (Keravnou and Johnson 1989), which are also prominent in philosophies of discovery (see for example Niiniluoto 1999).

Figure 1 Conceptual model of the diagnostic process.

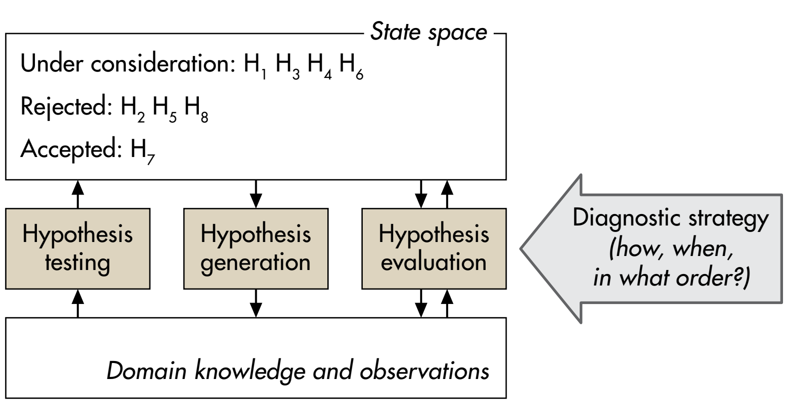

1. *Hypothesis generation:* On the basis of domain knowledge and observations, the diagnostician invents a new candidate explanation and adds it to the collection of hypotheses under consideration. Generated hypotheses may range from general and broad causal directions (for example, "The problem is caused in the soldering process") to specific and detailed causal explanations ("The problem is a short circuit, created by contaminations with salts, deposited by a newly introduced soldering flux."). Often, a number of complementary hypotheses are generated, such as the cause is in subsystem *A, B,* or *C.* (Keravnou and Johnson 1989).
2. *Hypothesis testing:* Given the hypotheses under consideration, the diagnostician determines what observations or knowledge are needed in order to evaluate a particular hypothesis. When necessary, the diagnostician does new observations or tests and adds the results to the body of domain knowledge and observations.
3. *Hypothesis evaluation:* Given the hypotheses under consideration, and given the available domain knowledge and observations, the diagnostician decides to change the status of a hypothesis from under consideration to accepted or rejected. Accepted and rejected hypotheses are added to the domain knowledge as findings.

The search does not necessarily end when a certain hypothesis is accepted as true, as this hypothesis may be too general and in need of further elaboration. Rather, the goal state of the diagnostic search is reached upon the acceptance of a hypothesis that offers a sufficiently specific explanation of the malfunction to enable the design of a solution. Thus, after accepting a broad hypothesis ("The problem's cause is in the soldering process"), the diagnostician may continue the search by generating, testing, and evaluating more specific elaborations of this initial explanation ("The cause is in the soldering flux," "The alloy's temperature is too low," and so on).

The hypothesis generation and evaluation operators are driven by domain knowledge and observations. Conversely, hypothesis testing directs the collection of additional knowledge and observations on the basis of the hypotheses under consideration. Domain knowledge is knowledge about the product or process under study: how it works, what parts it is composed of, how it has been used, what it behaves like normally, and so on. For example, the SIPOC model and flowchart technique (Pyzdek 2003) are often used for laying down domain knowledge about a production process. A particular form of domain knowledge concerns problems that have troubled the system or comparable systems in the past (fault knowledge). Table 1 lists typical elements of domain knowledge and observations used for diagnosis. The table was composed by the author, based on the types of domain knowledge used in A.I. systems for diagnosis (Clancey 1988; Keravnou and Washbrook 1989; Lucas 1997; Chittaro and Ranon 2004), and translated to the context of quality problem solving.

The domain knowledge and observations needed for diagnosis are typically not given a priori in a complete and consistent form, but are collected in interaction with the generation and evaluation of hypotheses. Most diagnostic searches start, however, with a first round of domain knowledge and observations gathering. This initial round is sometimes called cue acquisition or the study of symptoms (Gryna 1988). Where data and knowledge acquisition later in the diagnostic process will be greatly guided by working hypotheses under consideration and the diagnostic strategy in use, the initial cue acquisition tends to be quite general.

Practical Example: Electrical Instabilities

Since the framework the author develops is fairly abstract, he describes a real-life example, which is presented in more detail in de Mast (2011). The example serves no other purpose than illustration, and it will be revisited throughout the remainder of the paper. The case took place at a manufacturer of electrical devices. Quite suddenly, from week 29

Table 1 Forms of domain knowledge and observations used in diagnosis.

Physical structure A model of the system's anatomy in the form of a decomposition of the product or process into subsystems, components, stations, and parts.
Functional structure A model of how the product or process works, by specifying the function of components (in terms of input and intended output) and their linkages.
Operations context Especially for production processes, knowledge about how and when the process has been deployed as a resource in the production schedule (for example, for what type of products, in what batch sizes, in which shifts?).
Normal behavior Knowledge about normal states and normal behavior of the product or process.
General knowledge Scientific and professional knowledge of physics, electronics, chemistry, and other relevant fields.
Fault knowledge Knowledge about problems that have troubled the product or process in the past. Fault knowledge can be in a raw form (recollections of past problematic episodes in people's memory, or as descriptions in logs or on the Internet). Often it is in a compiled form, such as fault dictionaries (a list of known fault types with typical symptoms) or taxonomies (tree-type classifications of known fault types, providing operational definitions that guide diagnosis).
Observations Measurement data collected from the process or products, including data from experiments, and qualitative and less-structured observations in the form of findings and anecdotes.

in 2008 onward, a problem emerged in the assembly of a type of products the author refers to as *TA*. Interestingly, a product named *TB* that was nearly identical in design and production process, was unaffected. The problem was described as an electrical instability, and it destroyed about 12 percent of the *TA*s when connected to a power supply.

The main suspect was the connector linking the *TA*s to the power supply. In brainstorming sessions the engineers and operators generated possible causes, mainly related to the connector, and they also did some close examinations of the damage's appearance. After three months, these investigations still had not identified the problem's true cause. Eventually, the engineers designed a swapping test, in which the power supply and connector used for the *TA*s were adjusted such that they could be used for the *TB*s. The test demonstrated that the power supply and connector gave no problems when used for *TB*s. The team inferred that the problem's cause must be in the products themselves, including their production history, and not in the connector or the power supply.

These findings completely shifted the focus of the investigations. The team studied the production processes of the *TA*s and *TB*s more closely, and identified the soldering process as the main difference between them. Meanwhile, a team member did a literature search for known issues with this sort of devices, and found a list of four suspects, including contamination with salts. Cleaning the surfaces of some *TA*s and analyzing the residues demonstrated the presence of sodium chloride ("table salt"). Closer examination revealed that different soldering fluxes were used for the *TA*s and *TB*s, and that the flux used for the *TA*s contained sodium chloride. This flux had been introduced in the process recently. An engineer measured the conductivity of a surface contaminated with this flux, and found it sufficient for a short circuit given the high voltages the products used. The use of the soldering flux containing sodium chloride was then discontinued, and four weeks later not a single new electrical instability had occurred.

The case is a good example of the sort of diagnostic problem solving that is the focus of this paper. Diagnosis resulted in an accepted causal explanation: the introduction of a new soldering flux had resulted in contaminations of the products' surfaces with salt, which in turn created a short circuit. This destroyed the products when connected to a high-voltage power supply. The diagnostic search was an iterative process of hypothesis generation, testing, and evaluation, alternated with data and knowledge acquisition. The search resulted in the discovery of a fault mechanism that was novel to the people involved.

Early in the search, hypotheses focused on the connection of the products to the power supply. The turning point in the search was the generation and systematic testing of the three complementary hypotheses that the cause would be related to the product itself, to the power supply, or to the connector (the swapping test). Upon rejection of the latter two, the remainder of the project focused on refining the hypothesis that the cause is related to the product itself. A literature study, the observation that the soldering process is the main difference between the *TA*s and *TB*s, and some further investigations yielded the conclusion that was ultimately accepted as the problem's diagnosis.

SIX DIAGNOSTIC STRATEGIES

A Study of Diagnostic Strategies

The process of diagnosis progresses by hypothesis generation, testing, and evaluation (see Figure 1). A diagnostic strategy is a structure that prescribes how to do these actions in a certain order and in certain ways, for example by suggesting what sort of hypotheses to generate and in what order to test them. The function of a strategy is to make the diagnostic process more efficient (that is, it aims to reduce the expected amount of effort).

In the subsequent sections the author presents six diagnostic strategies. They were identified by studying approaches recommended in the practitioners' literature, from which he studied a large sample (references in the subsequent sections are restricted to Gryna 1988; Breyfogle 2003; Pyzdek 2003; George et al. 2004; Gitlow and Levine 2004). Further, he studied a large number of case-study descriptions discussed in journals such as *Quality Engineering* and the *Journal for Quality Technology,* as well as in practitioners' books such as Bhote (1991). A last source of information was the before-mentioned scientific literature in A.I. and medical diagnosis. For each of the identified approaches, the strategic idea was pinpointed, and where possible related to principles in the problem-solving literature. Further, for each of the identified strategies, the types of domain knowledge or observations needed as input were analyzed, and what strategic value it can bring in the diagnostic process. This resulted in a characterization of the situations where each strategy is promising, thus providing support to practitioners for choosing among alternative approaches. The author has also associated well-known quality problem-solving techniques to the discussed strategies, thus clarifying their rationale and applicability to practitioners.

As a basis for comparison, the author first describes the least-efficient approach for diagnosis that is realistic: a blind trial-and-error search. Suppose the diagnostician is only given a list of potential fault types but no further knowledge, nor is it possible to acquire further knowledge (except for trying out candidate causes). The only option in that case is to randomly try fault types until finding the one that fixes the observed problem. If even a list of potential fault types is absent, the strategy boils down to randomly inventing causal explanations and trying them out. Such blind trial-and-error approaches represent a limiting case in terms of diagnostic efficiency. In the case of a finite number n of potential causes, the expected number of trials for identifying the true one is $(n+1)/2$. In the remainder of this section, the author describes six strategies that aim to make the search process more efficient.

Lucky Guess Strategy

In a lucky guess strategy, the diagnostician thinks he or she recognizes the symptoms of a known problem, and first tests this explanation before embarking on more systematic diagnostic efforts. This sort of reasoning, where a cause is conjectured based on experiential association ("it has caused similar problems in the past"), is named *shallow* or *nonanalytic reasoning* in the diagnostic literature (Milne 1987; Keravnou and Washbrook 1989; Eva 2005). It is driven by a diagnostician's recollection of earlier experiences with similar problems (fault knowledge).

If the initial guess is right, a lucky guess strategy is the most efficient approach possible, and especially for routine problems, experts' recognition of familiar symptoms may obviate more elaborate searches. The literature on medical diagnosis asserts that expert physicians make most diagnoses by shallow reasoning, resorting to deeper reasoning only in novel or atypical cases (Elstein and Schwarz 2002). But lucky guess approaches are also fallible and risky; if the first guess is a dead end, the strategy may bog down the diagnostician to the wrong part of the search space, or the search may quickly degenerate into blind trial-and-error. According to Wagner (1993), in particular less-experienced problem solvers tend to fixate early in the search on a specific explanation, and study it in full detail, instead of switching to a more systematic search strategy.

Lucky guesses are unlikely to be correct for novel fault types, as they are driven by past experience with similar problems. Indications for attempting a lucky guess include:

- The problem appears to be routine.
- The immediate evidence suggesting the lucky guess explanation is very strong.
- The cost and effort in testing the guess are relatively minor.

Symptomatic Search Strategy

Most computer problems can be solved by entering an error message verbatim in an Internet search engine, which usually produces an overview of the likely causes. In a symptomatic search the diagnostician uses a set of observations representing the problematic behavior as a search template to find a matching set in a library of known symptoms and their likely causes (Venkatasubramanian, Rengaswamy, and Kavuri 2003b). These libraries may consist of compiled fault knowledge (such as fault dictionaries and taxonomies), but the diagnostician may also search in knowledge stores of a less-structured nature, such as a general search engine on the Internet. Alternatively, the diagnostician may consult the expert literature, or discuss the problem with colleagues or experts in a group meeting, hoping they recognize the symptoms as indicative for a known fault type.

The name *symptomatic search* was taken from the literature on troubleshooting (Rasmussen 1981; Rouse 1983). In the A.I. literature, diagnosing a problem by matching observed characteristics to those of cases stored in a library, is called case-based reasoning (for example, Kolodner, Simpson, and Sycara 1985; Portinale et al. 1994).

Symptomatic searches may give a short-cut, and rather than discovering the causes of a problem oneself, the diagnostician reuses knowledge gathered in earlier problematic episodes. As with lucky guesses, symptomatic searches are driven by fault knowledge, as well as by a description of symptoms. For that reason, they are unlikely to be effective for novel fault types. Indications for attempting a symptomatic search include:

- Rich fault knowledge is available.
- The problem is unlikely to be novel.
- The symptoms are salient and specific.

Branch-and-Prune Strategy

The next three strategies aim to reduce the extensiveness of the search space by ruling out entire classes of causal directions ("pruning"). De Mast (2011) offers an earlier discussion. The first one, the branch-and-prune strategy, seeks to reduce the extensiveness of the search space by first splitting

Figure 2 Branching and pruning in diagnosing electrical instabilities.

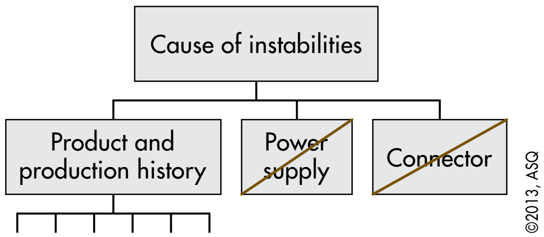

the space into broad but complementary hypotheses (the "branch" step). This coarse partitioning is done on the basis of the structure of the process or product under study, or on the basis of generic structures such as time and space. Next, tests or observations drive the elimination of entire branches ("prune"), and the diagnostician studies only the retained branch(es) in more detail.

The breakthrough in the electrical instabilities example (de Mast 2011) came from the swapping test, where the connector and power supply of the *TA*s were tested on the *TB*s. The test branched the search space into three complementary and broad hypotheses, namely, that the cause is either in the products themselves, in the power supply, or in the connector (see Figure 2). Based on the test's results, the latter two branches were pruned, and the power supply and connector were discarded from the remaining search. Instead, the team elaborated the causes-related-to-the-product branch in more detail.

Branch-and-prune strategies treat the space of potential causal explanations as a hierarchical tree structure, with broad and general hypotheses in the higher levels, and more specific refinements of these as branches in the lower levels. Hypotheses in a single layer should be complementary (for example, the cause is either in the product, in the power supply, or in the connector). The diagnostician works top-down, first pruning most of the branches in a level, and only then elaborating the retained branches into more detailed and specific hypotheses. The idea to use a hierarchy of hypotheses on various degrees of abstractness is known in A.I. as hierarchical model-based diagnosis (Mozetič 1991; Chittaro and Ranon 2004). A strategy of eliminating broad classes of causes, followed by zooming-in on the retained classes, is promoted in the practitioners' literature by Shainin (1993).

The branching step is based on system structure, such as a product or process's physical or functional structure. Also, generic structures such as time and space may serve as a basis for branching the search tree (de Mast 2011). The pruning step is driven by observations collected in a stratified sample, with the strata determined by the branches under study.

A rigorous proof of the general efficiency of branch-and-prune strategies is difficult. In simplified and stylized situations, however, the efficiency is easy to demonstrate. Suppose a problem has n possible causes. In blind trial-and-error, the mean number of causes to test before the right one is found is $(n + 1)/2$. Suppose the system can be decomposed into m subsystems of k components each ($km = n$), and suppose it is possible to confirm or reject the relevance of each subsystem without knowing the details and specifics within the subsystem. Suppose further that there are no interdependencies between the causes in one subsystem and another. The expected number of tests for pinpointing the problem's cause is now: $m + (k + 1)/2$ (assuming one needs to test all m subsystems to identify the relevant one, and assuming one applies blind trial-and-error to single out the cause from the k candidates within the selected subsystem). If, for example, $n = 1{,}000$, and $m = 10$, $k = 100$, this amounts to 60.5 instead of 500.5 expected trials. Applying multiple branching and pruning steps consecutively improves the efficiency even further. Note that the gain in efficiency does not only concern the hypothesis testing effort, but also the effort of acquiring domain knowledge about a specific subsystem. Besides this stylized demonstration of the efficiency of branch-and-prune strategies, it is also an empirical finding that efficient professional problem solvers tend to apply such top-down refinement strategies (for example, Kassirer and Gorry 1978; Smith et al. 1986; Boreham 1986; Schaafstal, Schraagen, and Van Berlo 2000).

Branch-and-prune strategies are promising if the system under study (and thus the search space) is complex and extensive, and if it suggests a strong structure for branching the search space into non-interdependent classes. Next, the author discusses some well-known techniques and variants that he claims are based on branch-and-prune tactics.

Sequential structure and bisection

Bisection (Morris and Rouse 1985), also known as the half-split strategy or process dissection (Gryna 1988), is a branch-and-prune strategy based on sequential structure, such as the sequence of steps in a production process. The idea is to observe whether the problem is manifest halfway in the sequence or process, thus establishing whether the cause acts in the first or second half. Next, one studies the relevant half in more detail, possibly by applying bisection again to the selected half, and discarding the other half.

Branch-and-prune on the basis of physical structure: Component swapping

Branches could be defined on the basis of a decomposition of a process, machine, or product into its physical subsystems, components, and parts. Components search, or component swapping (Bhote 1991), is a practical technique for branching and pruning on the basis of such a physical decomposition. The team in the electrical instabilities example (de Mast 2011) did half a component swapping test by swapping the power supply and connector of the problematic *TA*s with those of the unaffected *TB*s. The observation that the connector of the *TA*s gave no problems when used for the *TB*s ruled out causes associated to the power supply and the connector, and established that the cause must be related to the product itself. Note that in a full component swapping test, the team would also have used the connector and power supply of the *TB*s for the *TA*s, and would probably have observed that this combination did reproduce the problem.

Multi-vari studies and charts

This approach has been proposed in particular for problems related to excessive variation (Gryna 1988). Typical multi-vari studies branch the search tree in four classes: causes that vary over time, between-production streams, within-production streams but between products, and causes that vary within products. The diagnostician collects a stratified sample of data following the same structure. The results are often presented graphically in a multi-vari chart (de Mast et al. 2001). The subsequent study focuses on the class that appears to contain the dominant source of variation, ignoring the other classes.

Simpler multi-vari studies only compare variation across production streams: Does the problem manifest itself more in some streams than in others? Is the main variation within streams or across streams? They are driven by simple group comparison techniques such as the analysis of means (ANOM), boxplots, or the analysis of variance (ANOVA).

Time series plots allow branching and pruning based on patterns over time. In the electrical instabilities example (de Mast 2011), a time series plot would have shown that the problem emerged quite suddenly in week 29, and remained at a constant level in the weeks thereafter. Explanations not involving a change in week 28 or 29, therefore, could be pruned from the search space, homing in on events and changes occurring around week 29.

4W2H, defect check sheets, concentration diagram

A rather informal way of applying branch-and-prune tactics is to acquire cues guided by the generic questions of who, where, when, what, how, and how much (known under the acronym 4W2H). These questions have the diagnostician probe the search space in terms of spatial or physical structure (where), temporal structure (when) and functional structure (what), exploring whether the search space can be pruned and the search focused. Also typical formats for cue acquisition, such as the concentration diagram, defect location check sheet, defect type check sheet, and process check sheet (Pyzdek 2003, 274-276), structure findings in physical, spatial, functional, and other structures, thus facilitating a branch-and-prune strategy.

Branch-and-prune on the basis of functional structure

Observing that some subfunctions of a process or product are malfunctioning, while others function normally, allows the diagnostician to focus on the components of the product or process related to the faulty subfunctions, and ignore the rest (Rasmussen 1981; Wagner 1993). Suppose one debugs a computer program consisting of a multitude of subroutines, and suppose these subroutines are associated to functions and subfunctions of the program. Observing or testing which of the program's functions are normal, and which are invalid, focuses inquiry to the relevant subroutines.

Branch-and-prune on the basis of operations context

Operations context refers to structures induced by the production schedule: different production streams, product types, shifts, and other structures defining strata. As in multi-vari studies, the diagnostician establishes whether the problem is present in some, but not in other strata. For example, upon observing that the electrical instabilities (de Mast 2011) affected the *TA*s but not the *TB*s, the team inferred that the cause must be something that distinguishes the design or the production process of both products. This focused attention on the soldering process, which was the main difference.

Pruning Following a Proximate Causes Strategy

In a proximate causes strategy, one achieves a more focused problem description by moving upstream in the chain of cause and effect. Such a strategy starts with a close examination of symptoms, possibly involving a disassembly of malfunctioning parts or products (Gryna 1988, refers to such examinations as *autopsies*). On the basis of these examinations, the diagnostician seeks to reconstruct the immediate causes of the problem, thus moving the problem description step by step in anti-causal direction. This may either pinpoint the root cause of the problem, or focus the diagnostic effort, as the identification of proximate causes (and thus the nature of the causal mechanism) may greatly prune the search space. Note that the well-known 5 Whys technique (ask "Why?" five times) facilitates a proximate causes strategy.

In the electrical instabilities example (de Mast 2011), the team did a close inspection of some destroyed products, looking whether the damage's position and appearance would provide clues as to the immediate cause; these studies were, however, fruitless. Later in the project, another occurrence of a proximate causes strategy was more successful; the detection of sodium chloride on the products' surfaces allowed the team to reason back from the electrical instabilities to their immediate causes, producing the cause-and-effect chain: salt residue → short circuit → electrical instability. As a result, the problem description was recast from "What causes the electrical instabilities?" to "Where does the salt residue come from?" Consequently, the search space could be pruned by discarding from further consideration all elements of the process that could not plausibly be expected to leave a salt residue.

The proximate causes strategy is known in the general problem-solving literature as forward search (Norman et al. 2000), a name, somewhat confusing in the context of diagnosis, reflecting that one reasons forward from givens toward a goal-state. The strategy may be promising early in the diagnostic process, when the search space is wide and complex, and where the identification of proximate causes may yield a better-focused problem definition. Indications for attempting a proximate causes strategy include:

- The current problem description is unspecific and unfocused.
- Symptoms give clear cues as to the sort of causal mechanism that produced them.

Pruning on the Basis of Syndromes

In a proximate causes strategy the search space is pruned by reasoning backward from the problem to its immediate causes. In a branch-and-prune strategy, the search tree is pruned on the basis of system structure. A third option is pruning on the basis of syndromes. In this approach, the diagnostician observes a series of occurrences of the problem and tries to identify patterns in concomitant symptoms (the *syndrome*). Such patterns may reveal a characteristic of the causal mechanism that helps in ruling out options. The pattern is typically contrasted to normal or unproblematic behavior.

De Mast and Trip (2007) discuss an example concerning eccentricity of pins on cell-phone components. A histogram of 125 eccentricity measurements brought to light that these values had a bimodal distribution, and thus, that two homogeneous populations could be discerned. This taught the engineer an important characteristic of the eccentricity's cause: that it must be a phenomenon with two clearly discernable states. As it turned out, there were two molds in the process, and one of them was worn out. In the same paper, de Mast and Trip (2007) describe a problem concerning excessive variation in a cutting process. After much but fruitless detective work, the breakthrough in understanding the causes of the variation came from a time series plot revealing that the deviations have a cyclical pattern. Time series analysis indicated that the cycles had a period of 40 or 80 products. The number of products on one loop of the conveyor belt was 80, thus focusing the attention to properties of the belt. Subsequent studies demonstrated that the belt's flexibility was a major cause of the variability in the dimensions of cut products.

If large sets of production data are available, in which many variables have been recorded, multivariate statistical techniques such as principal components analysis (PCA), cluster analysis, and partial least squares (PLS) may be used to identify syndromes; Garcia-Muñoz et al. (2003) describe an interesting example. In a syndrome-based strategy, one focuses on establishing the pattern of symptoms rather than finding the cause. Therefore, this approach is a primer for other strategies. Since one cannot identify a pattern from a single or just a few samples, syndrome-driven strategies require larger numbers of observations and can be quite laborious. This in itself may disqualify the strategy in the case of a problem that occurs relatively rarely.

Pairwise comparison

Pairwise comparison is a well-known technique that involves the systematic comparison of the problematic to the unproblematic by comparing the best-of-the-best (BOB) products to the worst-of-the-worst (WOW) products (Steiner, MacKay, and Ramberg 2008). A similar approach is Kepner and Tregoe's (1997) "*is* versus *is not* analysis," which has the problem solver identify what distinguishes objects, behavior, locations, and situations where the problem is from those where it could be but is not.

Funneling Strategy

Funneling strategies seek to test an enumerable list of specific candidate causes in an efficient manner. By enumerating elements of the process or machines as possible causes, or by listing obvious suspects, one generates a set of specific and detailed hypotheses. Based on this list, the diagnostician designs an efficient testing strategy, for example, involving a statistically designed experiment, Shainin's variable search procedure (Dasgupta, Adiga, and Wu 2011), or testing candidates one by one after sorting them by plausibility.

Funneling strategies are called differential diagnosis in medicine (Eva 2004), and are often prescribed in accounts of Six Sigma's DMAIC method (for example, George et al. 2004, 12-13; Gitlow and Levine 2004, 146ff.), typically using a brainstorming session to identify candidate causes, and a statistically designed experiment to test them efficiently. They work with rather specific hypotheses, and are efficient for trying out candidate causes in a compact area of the search space. There are dangers when a funneling strategy is tried before the search space has been sufficiently narrowed down to the right area. The first danger is that the search space is so extensive that potential causes are not enumerable or otherwise multitudinous, and testing all of them, even with an efficient experimental design, is too laborious. Mooren, de Mast, and Does (2012) present a case about premature wear-out of drills that illustrates this problem. The other danger is that the diagnostician, faced with an extensive search space, only raises candidate causes in a narrow area of the space, and thus may get bogged down to the wrong part of the search space.

In the electrical instabilities example (de Mast 2011), the team first attempted a funneling strategy, generating possible explanations in brainstorming sessions. These early attempts focused on the connection of the product to the power supply. The team persevered rather long in exploring this part of the search space (about three months), and in the end, it turned out to be a dead end. It was an external consultant who realized that the efforts lacked a systematically established focus. This focus was achieved by three pruning maneuvers. First, a branch-and-prune step, based on the physical decomposition of the system into three subsystems, focused the search on the product and its production history (the swapping test). Next, proximate causes tactics focused the search on process elements that could be expected to leave a salt residue. Finally, a branch-and-prune step based on the operations context focused the search on the difference in the production processes of the *TA*s and *TB*s: the soldering process. These tactics gave sufficient focus to identify the soldering flux as the culprit.

Thus, the most important indication for attempting a funneling strategy is:

- The search space is compact and focused.

Group meetings

Group meetings are popular in shop-floor problem solving to drive a funneling strategy, with a cross-disciplinary group generating a varied set of hypotheses for further testing. The identification of candidate causes is often guided by categories such as materials, machine, method, personnel, measurement, and environment (Breyfogle 2003). Such systems of standard causes are often helpful for enlarging the scope of the causal field under consideration. They allow the group to consider to which categories candidates belong, and which categories aren't being considered.

If the set of candidates is large, one needs a way of sequencing them for testing and evaluating. The affinity diagram and multi-voting technique (George et al. 2004) are often suggested for this purpose.

DISCUSSION AND CONCLUSIONS

A Generic Sequence of Strategies

To declare one of these six strategies universally best would miss an important point, as their effectiveness and efficiency strongly depend on the situation. Rather, diagnosticians are best advised to think strategically, but not to follow any particular strategy rigidly. In each stage of the diagnostic search the diagnostician should reassess the situation and its tactical consequences, and be opportunistic in switching from one strategy to another. In the electrical instabilities example, the diagnosis was established by a combination of branch-and-prune tactics (for example, the swapping test), proximate causes tactics ("Where does the salt residue come from?"), and a symptomatic search ("What are known issues with this sort of device?"). Opportunistic switching between strategies or perspectives is advised or noted by Boreham (1986), Davis and Hamscher (1988), Keravnou and Washbrook (1989), and Norman (2005).

Table 2 briefly summarizes the strategies described in the previous section. The table also suggests a certain order of these strategies that the author proposes is a rational sequence of diagnostic efforts. Prescribing a certain order rigidly would fail to appreciate the uniqueness of problems. However, the previous section has shown that each strategy has certain prerequisites and a certain strategic value, and this allows a characterization of the sort of situations where each strategy may be useful.

Table 2 Six strategies in a rational order.

1. Lucky guess strategy The diagnostician recognizes the symptoms of a known problem.	Known problem?
2. Symptomatic search strategy Symptoms are used as a query in a search through a knowledge store of known problems.	
3. Proximate causes strategy A more focused problem description is achieved by reasoning backward from the problem to its immediate causes. *5 Whys, autopsy*	Achieve focus on the relevant part of the search space by pruning
4. Branch-and-prune strategy The search space is split into high-level classes ("branch"); irrelevant classes are discarded from the search ("prune"), and the retained branches are elaborated in more detail. *Bisection (half-split strategy), component swapping, multi-vari study, 4W2H*	
5. Syndrome-driven pruning strategy The search space is pruned by identifying characteristics of the causal mechanism from patterns in observed symptoms. *Pairwise comparison*	
6. Funneling strategy An enumerable list of specific hypotheses is tested in an efficient manner. *Group meetings and designed experiments*	Efficient testing of detailed hypotheses

He thinks the sequence proposed in Table 2 can be motivated on the basis of this analysis at least as a heuristic advice.

The rationale for the given sequence is as follows. Confronted with a problem, the diagnostician first figures out whether he or she is facing a known problem (lucky guess or symptomatic search strategy). If the symptoms are recognized as being related to a known problem, reusing knowledge from earlier episodes may obviate more elaborate problem-solving efforts. If symptoms do not give a match in symptomatic searches, the diagnostician must discover the causal mechanism him or herself. Especially if the search space is large and complex, the second stage is about focusing the search by pruning the search space. Proximate causes tactics achieve this by working backward from the symptoms to their immediate causes, thus giving a more focused problem definition. Branch-and-prune strategies exploit system structure and observations to focus the search. And syndrome-driven pruning strategies apply pruning on the basis of characteristics of the causal mechanism revealed by patterns in the symptoms.

Once the search space has been pruned to a manageable size, a funneling strategy has the diagnostician enumerate detailed candidate causes in the focal area of the search space, and design an efficient sequence of tests for singling out the true cause.

Conclusions and Suggestions for Future Research

As early as the 1950s, Joseph Juran described the diagnostic journey as part of the Universal Sequence for Breakthrough, and offered a number of techniques for diagnosing quality problems (Juran 1998). Despite a multitude of books on problem solving in the practitioners' literature, the academic literature about quality management has largely not followed up this early work and has not systematically studied diagnostic problem solving. As a basis for such research, the author offers a conceptual framework that is firmly grounded in the literature of fields such as troubleshooting, medical diagnosis, and A.I., but modified to the sort of problem solving common in quality engineering. He shows that the process of problem diagnosis can be understood as a search through a space of potential explanations (hypotheses). The search progresses toward a goal state by activities that generate, test, or evaluate hypotheses. Diagnostic strategies structure these activities in a way designed to make the process efficient.

The proposed framework offers a structure that researchers can use as a basis to compare and evaluate various methodologies such as Six Sigma's DMAIC method, the system proposed by Shainin, and Kepner and Tregoe's problem analysis procedure. Such methodologies are the basis for courses in quality improvement taught to large numbers of professionals, and this fact underlines the importance of a systematic and rigorous scientific appraisal of their merits.

Support given in practitioners' books for identifying the causes of quality problems tends to lack structure (with the aforementioned work by Shainin and Kepner and Tregoe as notable exceptions). Popular books such as Breyfogle (2003) and Pyzdek (2003) describe isolated techniques, but do not integrate them into a diagnostic strategy or logical sequence. In recent years, the statistics community in the quality field has become aware of the dangers of studying techniques merely in isolation, and has emphasized the importance of integrating techniques into strategies (for example, Anderson-Cook et al. 2012; Steiner and MacKay 2013). The proposed framework offers a basis for taking this initiative beyond purely statistical techniques for problem solving. The linkage of well-known problem-solving techniques to the generic strategies in which they are useful, clarifies their rationale. It indicates to practitioners in which situations these techniques are promising, and it explains their function in the diagnostic process.

The presented overview of six strategies allows researchers to reveal which strategic ideas are underrepresented in existing quality problem-solving methodologies. Current methodologies tend to be limited to a single or only a few strategies, such as the funneling strategies promoted in many accounts of Six Sigma (de Mast and Lokkerbol 2012) and the branch-and-prune strategy in Shainin's method (de Mast 2011). The author's framework offers six generic strategies, thereby demonstrating that the range of approaches for problem diagnosis is much wider than recognized in any account in the quality field known to the author. This wide range of alternatives is relevant. Table 2 shows that the viability of each of these strategies depends on the stage in the diagnostic process. Also, he has specified indications for the viability of each strategy, demonstrating that different situations make different strategies promising. Most salient to the author is the almost sole reliance on funneling strategies in most accounts of Six Sigma's DMAIC method. As discussed earlier, a funneling strategy is typically inefficient without first achieving focus to the relevant part of the search space. A critical review of popular accounts of the DMAIC method (de Mast and Lokkerbol 2012), such as Breyfogle (2003), Pyzdek (2003), Gitlow and Levine (2004), and George et al. (2004), reveals that pruning strategies such as branch-and-prune are mentioned cursorily at best. Hopp and Spearman (2008) criticize DMAIC for failing to offer a provision for using knowledge about known problems (that is, DMAIC does not offer a strategy akin to the symptomatic search strategy). Academic research to improve this critical part of the DMAIC method, and incorporate a wider range of diagnostic strategies, is all the more important given the prominence of this model in teaching quality improvement to practitioners.

The standard techniques in quality management for visualization in group meetings, such as cause-and-effect diagrams, are mainly geared to funneling and proximate causes strategies. A welcome extension would consist of techniques for facilitating branch-and-prune and other pruning strategies in group meetings. Diagnostic trees, as used in Shainin's method and Steiner and MacKay (2005, 119ff.), are a promising option, visualizing the successive branching and pruning steps.

A final suggestion for further research is to try to learn about quality problem solving from empirical research, as is done occasionally in the fields of medical diagnosis and troubleshooting (for example, Norman et al. 2000; Schaafstal, Schraagen, and van Berlo 2000). Studying how experienced and successful problem solvers work may enrich the theory about diagnostic problem solving. And also, it would be valuable to study the effect of teaching diagnostic strategies on the prowess of professionals in solving realistic problems.

REFERENCES

Anderson-Cook, C.M., L. Lu, G. Clark, S. P. DeHart, R. Hoerl, B. Jones, R. J. MacKay, D. C. Montgomery, P. A. Parker, J. Simpson, R. Snee, S. Steiner, J. Van Mullekom, G. G. Vining, and A. G. Wilson. 2012. Statistical engineering—Forming the foundations. *Quality Engineering* 24: 110-132.

Balakrishnan, A., R. Kalakota, P. Si Ow, and A. B. Whinston. 1995. Document-centered information systems to support reactive problem-solving in manufacturing. *International Journal of Production Economics* 38: 31-58.

Bhote, K. R. 1991. *World class quality.* New York: Amacom.

Boreham, N. C. 1986. A model of efficiency in diagnostic problem solving: Implications for the education of diagnosticians. *Instructional Science* 15: 191-211.

Breyfogle, F. W. 2003. *Implementing Six Sigma: Smarter solutions using statistical methods,* second edition. New York: Wiley.

Chakravorty, S. S. 2009. Six Sigma programs: An implementation model. *International Journal of Production Economics* 119: 1-16.

Chittaro, L., and R. Ranon. 2004. Hierarchical model-based diagnosis based on structural abstraction. *Artificial Intelligence* 155: 147-182.

Clancey, W. J. 1988. Acquiring, representing and evaluating a competence model of diagnostic strategy. In *The Nature of Expertise,* eds. M. Chi, R. Glaser, and M. Farr, 343-419. London: Lawrence Erlbaum Associates.

Custers, E. J. F. M., G. Regehr, and G. R. Norman. 1996. Mental representations of medical diagnostic knowledge. *Academic Medicine* 71: S55-S61.

Dasgupta, T., N. Adiga, and C. F. J. Wu. 2011. Another look at Dorian Shainin's variable search technique. *Journal of Quality Technology* 43: 273-287.

Davis, R., and W. C. Hamscher. 1988. Model-based reasoning: Troubleshooting. In *Exploring Artificial Intelligence,* edited by H. E. Shrobe, 297-346. San Francisco: Morgan Kaufmann.

De Mast, J. 2011. The tactical use of constraints and structure in diagnostic problem solving. *Omega* 39: 702-709.

De Mast, J., and J. Lokkerbol. 2012. An analysis of the Six Sigma DMAIC method from the perspective of problem solving. *International Journal of Production Economics* 139: 604-614.

De Mast, J., K. C. B. Roes, and R. J. M. M. Does. 2001. The multi-vari chart: A systematic approach. *Quality Engineering* 13: 437-447.

De Mast, J. and A. Trip. 2007. Exploratory data analysis in quality improvement projects. *Journal of Quality Technology* 39: 301-311.

Elstein, A. S., and A. Schwarz. 2002. Clinical problem solving and diagnostic decision making: Selective review of the cognitive literature. *British Medical Journal* 324: 729-732.

Eva, K. W. 2005. What every teacher needs to know about clinical reasoning. *Medical Education* 39: 98-106.

García-Muñoz, S., T. Kourti, J. F. MacGregor, A. G. Mateos, and G. Murphy. 2003. Troubleshooting of an industrial batch process using multivariate methods. *Industrial Engineering and Chemistry Research* 42: 3592-3601.

George, M. L., D. T. Rowlands, M. Price, and J. Maxey. 2004. *The Lean Six Sigma pocket toolbook: A quick reference guide to 100 tools for improving quality and speed.* New York: McGraw-Hill.

Gitlow, H. S., and D. M. Levine. 2004. *Six Sigma for green belts and champions: Foundations, DMAIC, tools, cases, and certification.* New Jersey: FT Press.

Gryna, F. M. 1988. Quality improvement. In Juran's *Quality Handbook,* 4th edition, ed. J. M. Juran, 22.1-74. New York: McGraw-Hill.

Ho, J. K. K., and D. Sculli. 1997. The scientific approach to problem solving and decision support systems. *International Journal of Production Economics* 48: 249-257.

Hopp, W. J., and M. K. Spearman. 2008. *Factory physics,* third edition. New York: McGraw-Hill.

Juran, J. M. 1998. The quality improvement process. In *Juran's Quality Handbook,* fifth edition, eds. J. M. Juran, and A. B. Godfrey, 5.1-74. New York: McGraw-Hill.

Kassirer, J. P., and A. Gorry. 1978. Clinical problem solving: A behavioral analysis. *Annals of Internal Medicine* 89: 245-255.

Kepner, C. H., and B. B. Tregoe. 1997. *The new rational manager.* Princeton: Kepner-Tregoe.

Keravnou, E. T., and L. Johnson. 1989. Towards a generalized model of diagnostic behaviour. *Knowledge-based Systems* 2: 165-177.

Keravnou, E. T., and J. Washbrook. 1989. Deep and shallow models in medical expert systems. *Artificial Intelligence in Medicine* 1: 11-28.

Kolodner, J. L., R. L. Simpson, and E. Sycara. 1985. A process model of case-based reasoning in problem solving. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence,* 284-290. Los Altos, CA: Kaufmann.

Lucas, P. J. F. 1997. Model-based diagnosis in medicine. *Artificial Intelligence in Medicine* 10: 201-208.

MacDuffie, J. P. 1997. The road to 'Root Cause': Shop-floor problem-solving at three auto assembly plants. *Management Science* 43: 479-502.

Milne, R. 1987. Strategies for diagnosis. *IEEE Transactions On Systems, Man, and Cybernetics* 17: 333-339.

Mooren, J., J. de Mast, and R. J. M. M. Does. 2012. Quality quandaries: The case of premature drill wear-out. *Quality Engineering* 24: 354-359.

Morris, N. M., and W. B. Rouse. 1985. Review and evaluation of empirical research in troubleshooting. *Human Factors* 27: 503-530.

Mozetič, I. 1991. Hierarchical model-based diagnosis. *International Journal of Man-Machine Studies* 35: 329-362.

Newell, A., and H. A. Simon. 1972. *Human problem solving.* Englewood Cliffs, NJ: Prentice-Hall.

Niiniluoto, I. 1999. Defending abduction. *Philosophy of Science* 66: S436-S451.

Norman, G. R. 2005. Research in clinical reasoning: Past history and current trends. *Medical Education* 39: 418-427.

Norman, G. R., L. R. Brooks, C. L. Colle, and R. M. Hatala. 2000. The benefit of diagnostic hypotheses in clinical reasoning: Experimental study of an instructional intervention for forward and backward reasoning. *Cognition and Instruction* 17: 433-448.

Pople, H. E. 1982. Heuristic methods for imposing structure on ill-structured problems: The structuring of medical diagnosis. In *Artificial Intelligence in Medicine,* ed. P. Szolovits, 119-190. Boulder, CO: Westview Press.

Portinale, L., P. Torasso, C. Ortalda, and A. Giardino. 1994. Using case-based reasoning to focus model-based diagnostic problem solving. In *Lecture Notes in Computer Science,* vol. 837: *Topics in Case-Based Reasoning,* eds. S. Wess, K. D. Althoff, M. M. Richter, 325-337. Berlin: Springer.

Pyzdek, T. 2003. *The Six Sigma handbook: The complete guide for greenbelts, blackbelts, and managers at all levels,* 2nd edition. New York: McGraw-Hill.

Rasmussen, J. 1981. Models of mental strategies in process plant diagnosis. In *Human detection and diagnosis of system failures,* eds. J. Rasmussen and W. B. Rouse, 241-258. New York: Plenum Press.

Rouse, W. B. 1983. Models of human problem solving: Detection, diagnosis, and compensation for system failures. *Automatica* 19: 613-625.

Schaafstal, A., J. M. Schraagen, and M. van Berlo. 2000. Cognitive task analysis and innovation of training: The case of structured troubleshooting. *Human Factors* 42: 75-86.

Shainin, R. 1993. Strategies for technical problem solving. *Quality Engineering* 5: 433-448.

Smith, G. F. 1988. Towards a heuristic theory of problem structuring. *Management Science* 34: 1489-1506.

Smith, G. F. 1998. Determining the cause of quality problems: Lessons from diagnostic disciplines. *Quality Management Journal* 5: 24-41.

Smith, P. J., W. C. Giffin, T. H. Rockwell, and M. Thomas. 1986. Modeling fault diagnosis as the activation and use of a frame system. *Human Factors* 28: 703-716.

Steiner, S. H., and R. J. MacKay. 2005. Statistical engineering. Milwaukee: ASQ Quality Press.

Steiner, S. H., and R .J. MacKay. 2013. *Statistical engineering* and variation reduction. Quality Engineering, in press.

Steiner, S. H., R. J. MacKay, and J. S. Ramberg. 2008. An overview of the Shainin System for quality improvement. *Quality Engineering* 20: 6-19.

Torasso, P., and G. Torta. 2005. Compact diagnoses representation in diagnostic problem solving. *Computational Intelligence* 21: 27-68.

Venkatasubramanian, V., R. Rengaswamy, S. N. Kavuri, and K. Yin. 2003a. A review of process fault detection and diagnosis, Part I: Quantitative model-based methods. *Computers & Chemical Engineering* 27: 293-311.

Venkatasubramanian, V., R. Rengaswamy, and S. N. Kavuri. 2003b. A review of process fault detection and diagnosis, Part II: Qualitative models and search strategies. *Computers & Chemical Engineering* 27: 313-326.

Wagner, C. 1993. Problem solving and diagnosis. *Omega* 21: 645-656.

BIOGRAPHY

Jeroen de Mast is professor of methods and statistics for operations management at the University of Amsterdam. He also holds a position as principal consultant at the Institute for Business and Industrial Statistics (IBIS UvA). He is recipient of ASQ's Brumbaugh, Nelson, and Feigenbaum awards. He can be reached by email at J.deMast@uva.nl.