

Simulating Families of Studies to Build Confidence in Defect Hypotheses¹

Forrest Shull¹, Daniela Cruzes^{2,3}, Victor Basili², Manoel Mendonça³

¹ Fraunhofer Center - Maryland, 4321 Hartwick Road, Suite 500,
College Park, MD, 20740, USA
fshull@fc-md.umd.edu

² Dept. of Computer Science, University of Maryland,
College Park, MD, 20742, USA
basili@cs.umd.edu

³ Computer Networks Research Group (NUPERC), Salvador University (UNIFACS),
Rua Ponciano de Oliveira, 126 Salvador, BA 41950-275, Brazil
{daniela, mgmn}@unifacs.br

Abstract. While it is clear that there are many sources of variation from one development context to another, it is not clear a priori what specific variables will influence the effectiveness of a process in a given context. For this reason, we argue that knowledge about software process must be built from families of studies, in which related studies are run within similar contexts as well as very different ones. Previous papers have discussed how to design related studies so as to document as precisely as possible the values of likely context variables and be able to compare with those observed in new studies. While such a planned approach is important, we argue that an opportunistic approach is also practical. The approach would combine results from multiple individual studies after the fact, enabling recommendations to be made about process effectiveness in context.

In this paper, we describe two processes with which we have been working to build empirical knowledge about software development processes: One is a manual and informal approach, which relies on identifying common beliefs or “folklore” to identify useful hypotheses and a manual analysis of the information in papers to investigate whether there is support for those hypotheses; the other is a formal approach based around encoding the information in papers into a Structured Hypothesis Base that can then be searched to organize hypotheses and their associated support. We test these processes by applying them to build knowledge in the area of defect folklore (i.e., commonly accepted heuristics about software defects and their behavior). We show that the formal methodology can produce useful and feasible results, especially when it is compared to the results output from the more manual, expert-based approach. The formalized approach, by relying on a reusable Hypothesis Base, is repeatable

¹ SHULL, Forrest ; CRUZES, Daniela Soares ; BASILI, Victor Robert ; MENDONÇA NETO, M. G. . Simulating Families of Studies to Build Confidence in Defect Hypotheses. Information and Software Technology, v. 47, n. 15, p. 1019-1032, 2005.

and also capable of producing a more thorough basis of support for hypotheses, including results from papers or articles that may have been overlooked or not considered by the experts.

1 Introduction

Empirical studies have long been used to provide confidence in assertions about what is true and not true in the software engineering domain. By providing rigorous observation of the effects of a development technique under specific conditions, empirical study allows an analysis of the conditions under which practices yield similar effects on a project's cost, quality, or schedule. Where different results are obtained under different contexts, the contexts can be analyzed to hypothesize about which variations led to the differences in results².

The ability to build up rigorous abstractions of information about practices not only provides confidence in individual assertions about specific techniques, but is an important capability in providing an engineering basis for software development. This capability is an essential part of the Experience Factory approach [Basili94], for example, as well as more recently suggested in the idea of Evidence-Based Software Engineering [Kitchenham04].

Due to the large number of possible variations from one development environment to another, we have argued [Basili99] that this process of knowledge building about practices must therefore be based on families of related studies, designed so that a range of context variations can be explored. Although this approach is logically appropriate, it does pose some practical problems. First, it is not always clear *a priori* what the important context variables are, meaning that important sources of variation may go unmeasured. Second, because there are so many potential context variables, we cannot design experiments that cover all of them and may not be able to identify environments which offer coverage of all the variables. In other words, to build an effective family of studies, multiple experimenters, without having a clear concept of all the contributing factors, must agree *a priori* on a set of variables to collect and identify environments that cover the complete set of variables, so that all studies are fully comparable.

Providing robust decision support for software development – i.e. making a statement about what development practices can help achieve goals related to cost, quality, or schedule for a given environment – requires the collection of data across multiple contexts so that we can begin to elicit these variables.

To make this approach work in practice, and build a suitably large and varied dataset, we need to leverage existing work wherever possible. That is, we have to be able

² The set of potential context variables is quite large, but includes for example such issues as: Team size (larger teams require more communication overhead, making certain practices more or less effective); team experience (certain practices may require a level of skill development that makes them better suited for experts); lifecycle model (if a project is contractually required to follow a waterfall lifecycle then practices that better fit a spiral lifecycle may not be appropriate); etc.

to mine information about the relevant variables and the effect of practices from previous studies that were not *a priori* designed to fit together. In effect, this requires the ability to simulate a family structure over independent studies that were not explicitly designed to build directly on one another.

In this paper, we extend the work in our earlier paper [Basili99] by presenting a process for creating a “post-hoc” family of studies. We present a set of results in the area of software defect phenomenology to demonstrate the feasibility and practicality of this approach.

2 Challenges in combining studies post-hoc

Designing a replicated study that fits together with previous work to build a family of studies is difficult (cf. [Brooks96, Miller05]). A mistake in experimental design or application of the practice under study can easily render the study invalid and waste potentially hundreds of person-hours of effort from multiple subjects.

Combining the studies after the fact also suffers from several difficulties (although at least if mistakes are made, only the final body of knowledge is affected and only the analysis needs to be redone).

Borrowing concepts from work on Ontology Integration [Klein01, Visser97], we can see that mismatches between empirical studies are the key type of problems that hinder the combined use of independently developed studies.

Visser *et al.* enumerate the types of mismatches that may occur at the semantic level, which describe differences in the way a domain is modeled. Semantic-level mismatches can arise when two or more papers that describe (partly) overlapping domains are combined. Based on classifications by Visser and Klein, the different types of mismatches include:

- **Conceptualization Mismatches:** A conceptualization mismatch is a difference in the way a domain is *interpreted* (conceptualized), which results in different concepts or different relations between those concepts. There are two types:
 - **Scope:** In this type of mismatch, two results seem to represent the same concept, but do not have exactly the same meaning (although there may be some overlap). For example: Two studies may refer to the “cost” of a practice, although one may include only the cost of the effort to apply the practice in the calculation, while the other may include the start-up costs as well (e.g. sending personnel for training).
 - **Model coverage and granularity:** This type of mismatch describes problems that can arise in trying to combine results when it is unclear to what part of the domain those results are applicable. For example, a study may make claims about a large class of software development projects while only having evidence concerning one or two specific instances of such projects.
- **Explication Mismatches:** An explication mismatch is a difference in the way the conceptualization is *specified*. This can manifest itself in mismatches in definitions, mismatches in terms and combinations of both. There are three types:

- **Synonymous terms:** Synonyms, in this context, are different terms which refer to the same concept. A trivial example is the use of the term “strength” in one study and the term “cohesion” in another, to refer to the same concept (that is, the amount of interaction within components of a system). Although the technical solution for this type of problems seems relatively simple (the use of thesauri), the integration of studies with synonyms or different languages usually requires significant human effort to resolve the semantic issues.
- **Homonym terms:** This type of mismatch occurs when the meaning of a term is different in different contexts. For example, the term “interface defects” can have different interpretations, depending on the context: It can refer to a defect in the Human-Computer Interface or a defect in the interfaces between two software components. This inconsistency is much harder to handle; human knowledge is required to solve this ambiguity. One must also be careful that a mismatch in the explication is not masking a deeper mismatch at the conceptual (scope) level.
- **Encoding:** Values in the studies may be encoded in different formats. For example, a numbers of lines of code may be represented as “KLOC” or as “LOC” or “SLOC,” etc. The first step is to check that the basic definition of a line of code is the same across each source, i.e. that the underlying entity being counted is really the same in each case. If so, and the differences are in fact that the results have just been reported using different levels of granularity, a transformation step or wrapper can simply be applied to eliminate all those differences.

To avoid mismatch problems within this paper, we will use the following IEEE definitions [IEEE87] for defects and related phenomena. Where necessary we annotate direct quotes from the papers cited for data to enforce a consistent terminology that was not always used by the individual authors.

- Error: a defect in the human thought process made while trying to understand given information, to solve problems, or to use methods and tools;
- Fault: a concrete manifestation of errors within the software (note that one error may cause several faults and various errors may cause identical faults);
- Failure: a departure of the operational software system behavior from users’ expected requirements (a particular failure may be caused by several faults and some faults may never cause a failure).

Where necessary we will also use the term *defect* as a generic term, to refer to an error, fault, or failure.

3 Building hypotheses using empirical evidence

The current state of the practice for abstracting information from across several studies is to perform a literature search, reviewing the relevant literature in a rigorous

way and constructing a textual summary of the evidence related to a given issue. If the sources do not agree then it is the reviewer's responsibility to construct a fair summary of the evidence on both sides of the issue. We have been using the concept of "folklore" as a way to focus such a literature search in a way that is useful for combining data and producing well-formed and empirically-supported hypotheses [Asgari05]. We define folklore as informal, subjective lessons learned based on the experiences of people in the field.

This approach involves considering the folklore as a kind of hypothesis and identifying papers that either support or countermand the folklore. In studying a set of related papers, the researcher is able to build evidence for or against the folklore, to refine the folklore into testable hypotheses, or to recognize the context variables that differentiated when and where the hypotheses were true or false (thus creating more specific sub-hypotheses).

In this section, we present an example of folklore testing as an example of how testable hypotheses can be generated by such a literature search. We choose the area of software defects as a topic which is especially rich in folklore; this is not surprising, as debugging software defects is an activity that consumes a significant amount of time for most software developers. Two pieces of folklore that we identified are:

- 1. There are patterns in the defect classes found in classes of projects.**
- 2. The vast majority of defects are interface defects.**

Both items of folklore are important as together they state the basis for much of what is done in defect analysis. If there are patterns in the defects we find in projects, then analyzing the defects in a project in a particular environment will allow us to better understand what techniques, methods and processes to apply in that environment in the future. If the majority of defects are interface defects, then the kinds of quality assurance techniques applied need to have the goal of preventing or identifying interface defects. To keep the example short we make use of only three papers, with which we happen to have significant experience:

- Endres75: One of the earliest papers on software defects, it describes the release of a version of an operating system in which approximately 500 modules (140K source lines of assembly code) were affected by the modification. (The 140K total program size had to be calculated based on average values given for the individual modules, providing an example of resolving an encoding mismatch as defined in Section 2.) Defects were classified as being problem-specific, implementation-specific, or textual-specific. In this analysis, problem-specific defects are considered as due to a misunderstanding of the requirements. [Endres75]
- Weiss/Basili85: This paper deals with a study of three projects in the Software Engineering Laboratory (SEL) at NASA/GSFC. They all dealt with development of requirements, design and code for ground support software for unmanned spacecraft control, and were 50K to 120K source lines of Fortran code. One of the issues specifically explored in the study was the distribution of defects by error origin (i.e., according to the phase in which the misunderstanding took place). In this case the phases recorded were Requirements, Functional Specification, Design or Coding of Multiple Compo-

nents, Design or coding of a single component, Language, Environment, and Other. [Weiss85]

- Basili/Pericone84: This paper deals with the development and evolution over three years of a general purpose program for Satellite Planning Studies in the SEL. The system was 90K source lines of Fortran code and the requirements kept changing and evolving over time. The same question was asked in this study as Weiss/Basili85 but the phases were defined slightly differently. Phases here were Requirements, Functional Specification, Design, Coding, and Evolution over 3 years. [Basili84]

The Endres75 study only had one project, but the largest source of defects (46%) in that project was associated with a misunderstanding of the problem domain. In the Weiss/Basili85 study, all three projects had very similar profiles but the majority of defects (between 56% and 72%) occurred during the design or coding of a single component. In the Basili/Pericone84 study, even though data was gathered in the same SEL environment, the majority of defects was caused by misunderstanding of the requirements (55%).

The difference in results among these studies allows us to examine the impact of the context variables. In Weiss/Basili85, the organization had earlier developed many similar projects, so that although the requirements varied, the organization had experience in developing that type of system. Thus the fact that the majority of the defects were occurring in the coding phases was due to the fact that the requirements and high level design were relatively well understood by the developers but new hires were used to develop the code. On the other hand, the Endres75 and Basili/Pericone84 studies involved projects with relatively new requirements and less understanding of those requirements by the organization. This allows us to postulate two new hypotheses:

1a. In novel projects, the largest source of defects will be due to misunderstanding of the requirements. [Supported by: Endres75, Basili/Pericone84 data sets; Contradicted by: None.]

1b. In projects where the organization has built up experience in the application, the largest source of defects will be coding. [Supported by: Weiss/Basili85 data set; Contradicted by: None]

This analysis also allows us to recognize that the hypothesis as originally stated is also supported.

With regard to the second hypothesis, the issue is the distribution of interface defects. Endres75 reported that only 15% of the defects found were interface defects. However, he defined interface defects based on the number of components changed, i.e., a defect is an interface defect if more than one component (module) must be changed to fix the problem. We can think of this as an “implementation interface defect.” Using the same definition, Basili/Pericone84 found that only 11% of the defects were “implementation interface defects.” However, Basili/Pericone84 used a second definition of interface defect, associated with the number of components examined; i.e., it is an interface defect if more than one component (module) had to be examined in order to design the change. We can call this a “design interface defect”. (The multiple definitions for “interface defects” are an example of the homonyms, as well as a scope mismatch, as discussed in Section 2.) This is a much more inclusive

definition and most definitions would lie somewhere between these two definitions. Using this latter definition, Basili/Pericone84 found that 39% of the defects were “design interface defects.” Thus, in either case the folklore is not true, meaning that a useful hypothesis describing defect behavior should be formulated regarding the exactly opposite condition:

2'. Interface defects will not comprise the vast majority of defects in a system. [Supported by: Endres75, Basili/Pericone84; Contradicted by: None]

These examples demonstrate the benefit of accumulating results from multiple studies in multiple papers, recognizing how context variables can provide insights into hypothesis refinement as well as open new hypotheses. However, this approach also demonstrates the difficulties alluded to in Section 2:

- The lack of formality makes it easy to miss mismatches among studies. In the analysis above, the integration of studies was done by someone with a large degree of context knowledge concerning all of the studies described, but this cannot always be expected to be the case.
- The analysis is not reusable; literature reviews are typically done to address a particular issue and if a related but different issue is of interest to another researcher, then typically the analysis must be redone from scratch.
- Due to the textual nature of the review, it is hard to create a summary that is both rigorously backed up by evidence and that summarizes the current state of knowledge.

4 A more formal methodology

Based upon the results from the folklore-based method of abstracting empirical results, we recognized the need for a more effective and rigorous process for identifying possible hypotheses based not just on folklore but on various forms of information from the unrelated studies themselves. In this section, we describe the process we created, which combines a structured approach to knowledge gathering plus data mining techniques. The outputs of this method are compared to those in Section 3 as a way to check the efficacy of this new method.

As in the folklore-based approach, this methodology has the goal of building a set of hypotheses that represent the knowledge about software development practices contained in multiple studies, which need not have been designed specifically to produce related data. This knowledge is represented as hypotheses to reflect its provisional nature, that is, each study in a family can only contribute evidence *for* or *against* some statement about software development phenomena. No amount of studies ever proves such a relationship outright.

As input, the methodology requires a focus of study, i.e. a (set of) software engineering phenomenon(a) about which information is needed.

The process consists of five steps (Figure 1), starting with a definition of the problem area, then the selection of papers of interest, the extraction of information from these papers, integration of the information and finally the analysis and interpretation of the results. This process is iterative, in that the results of a given step may convince the researcher to go back to a previous step and redo the associated activities. For

example, if the researcher is not satisfied with the information extracted from a set of papers, he or she may use these results to suggest new areas to search in order to select more papers for analysis.

The output of the process is a set of conclusions and new knowledge that arises from the process. As a secondary output the process creates a Structured Hypothesis Base, a structured and searchable repository of issues about which information has been found. The advantage of the creation of the Structured Hypothesis Base is that it can be reused. Other researchers can evolve and reuse it according to new research goals as they arise.

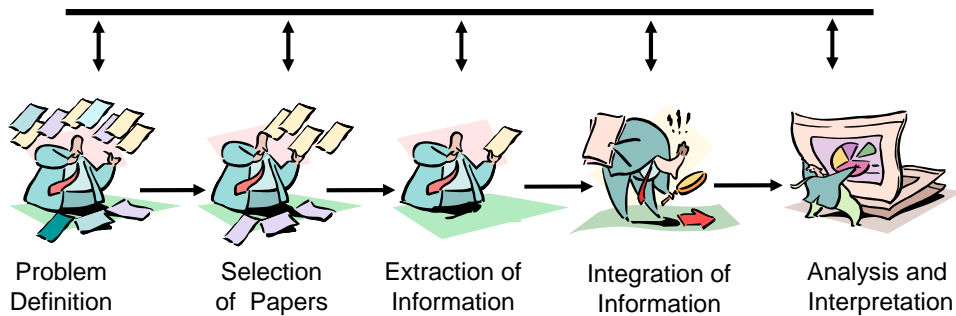


Figure 1 – Building and Effective Set of Structured Information – The Meta Process

The first two steps of the methodology, defining the problem and selecting relevant papers, are performed much the same as they would be in any method, no matter how formal, and are thus not discussed here at length. Defining the problem of course depends largely on the interests of the researcher. The problem definition is also related to the amount of knowledge already accumulated in an area. For example, as more evidence is accumulated we can move from studying how failure-prone software products are to which types of failures are most common; to which types of products display common failure profiles; to which context variables make those failure types more likely to occur. This allows us to evolve our knowledge into more useful models over time.

The second step, selecting papers that can be searched for evidence in the focus area, is also conducted largely in the same way regardless of the individual process being followed. To be suitable, a paper must provide some empirical information or experience-based hypotheses relating to the focus of study.

The remaining steps will be performed in a quite specific way in this methodology, and so are described in more detail in the following subsections.

4.1 Extracting Information from Papers

In this step, the researcher must review each paper, looking for potential hypotheses described in the text. According to Gay, a hypothesis is a tentative explanation for certain behaviors, phenomena, or events that have occurred or will occur. "A good hypothesis states as clearly and concisely as possible the expected relationship (or

difference) between two variables and defines those variables in operational, measurable terms." [Gay96]

Any hypothesis should be stated in such a way that data can be collected that either supports or refutes the hypothesis. For the purpose of our analysis, we classify the hypotheses as tested and untested. A tested hypothesis is a tentative explanation for certain behaviors, phenomena, or events that have occurred in an empirical study. An untested hypothesis (otherwise called a belief or assumption) is a tentative explanation for certain behaviors, phenomena, or events without explicit reference to empirical data.

While reviewing the papers selected in the previous step, the researcher should highlight the important information (so that there can be some traceability to the original source if questions arise later). After highlighting the information, it is important that key details are transferred to data entry forms to create the structured base for analysis. For now we are using forms implemented in Excel, although in future work we intend to create a tool to support the activities of the process. The researcher should focus his or her search specifically on the following categories and types of information: hypotheses, context descriptions, and definitions.

Hypotheses: It is important to note that in this step, all relevant information should be highlighted, without requiring critical appraisal from the researcher regarding the quality of support. For our Structured Hypothesis Base, we want to collect beliefs and assumptions as well as hypotheses with empirical results that support or reject their contention. These various levels will be distinguished from one another by using the "support" field described below. The reason for this practice is that we want to be able to describe the possible relationships among variables as broadly as possible, drawing from the common knowledge of experienced professionals as well as the subset of relationships that have been empirically verified. However, at the same time, we want to carefully record the quality of that support so that the end results can be understood in relation to the associated level of confidence that is warranted.

Our experience is that the section of the paper describing the analysis of the empirical data is where most tested hypotheses can be found. The conclusions are a good place to find hypotheses, although these are many times repetitions from earlier in the text. Some hypotheses will also be found in tables and figures; although not explicitly stated in the text of the paper, relationships that are expressed visually for readers will need to be translated into textual form to be inserted into the Hypothesis Base in a usable way.

The following information should be recorded on the appropriate form for each hypothesis:

- Plain Text: The hypothesis should be stated in such a way that data can be collected that either supports or refutes the hypothesis. A good hypothesis states as clearly and concisely as possible the expected relationship (or difference) between dependent and independent variables and defines those variables in operational, measurable terms. The hypothesis should be written using exactly the wording from the paper so that traceability is assured.
- Original Source of Hypothesis: A reference to the paper in which the information was found.

- Level of Support: The support should be described as one of the following levels:
 - Significantly positive: The results support the hypothesis and the results of a test are included to show that the results are statistically significant, that is, with a high degree of certainty are not the result of pure chance.
 - Positive: The results in the paper support the hypothesis, but no significant statistical results can back this up.
 - Null: The hypothesis has been tested but the results in the paper neither support nor contradict the hypothesis.
 - Negative: The results in the paper contradict the hypothesis, but no significant statistical results can back this up.
 - Significantly negative: The results contradict the hypothesis and the results of a test are included to show that the results are statistically significant, that is, with a high degree of certainty are not the result of pure chance.
 - Belief: The hypothesis is formulated based on assumption or belief but has not been tested.
- Confidence in support: Reflects the level of confidence that can be placed in the results and author's analysis. It is a 4-level scale:
 - High: Results were rigorously measured and in a representative industrial environment.
 - Med: Measurement was not completely rigorous *or* the context was not realistic.
 - Low: Measurement was not completely rigorous *and* the context was not realistic.
 - None: No evidence was presented. (The hypothesis describes a belief and has not been tested.)
- Observations: This is a free-text field where the researcher can keep track of additional information that is important for correctly understanding or interpreting the hypothesis.

Context descriptions: This form holds details concerning the environment from which the measures were drawn. There will be at least one record on this form associated with each source (possibly more, if the paper describes data that was collected from several projects). Our experience is that the context descriptions usually come shortly after the introduction. As different studies report different metrics of interest to them, not every paper will have all of the following information. However, the template should be filled out as completely as possible given the information that has been published. Missing values will need to be accounted for during the analysis, as they limit the strength of the conclusions that can be drawn.

The attributes of the context description form are:

- Source: A reference to the paper in which the information was found.
- Criteria related to the development team:
 - Number of Subjects: Number of subjects from which the author collected the information to do the analysis.
 - Experience of Subjects: Level of skill or expertise, captured using whichever measures the authors described in the paper.

- Criteria related to the software product:
 - Project Size: Size of the project, described using whichever measures the authors described in the paper.
 - Degree of complexity of the application: Complexity, described using whichever measures the authors described in the paper.
 - Notation: If applicable, the language used to encode the model of the system that was the object of study. For example, in a study of code the notation would be the programming language used; in a study of requirements, the notation could be natural language or a formal modeling notation like SCR.
 - Application Description: A brief description of the application. For example: “The development of an on-board flight control program for a new aircraft.”
 - Development Phase: The development phase in which the information was gathered: Requirements, Architecture, Design, Code, Test, Maintenance.
- Criteria related to the data collection process:
 - Data Source: The source of the data, like forms, compilers, tools, etc. The main intent is to provide information during the analysis phase concerning how objectively or subjectively the data was reported. Each source may have multiple sources, for different types of information; the granularity required in the form depends on the source.
 - Collection Period: Number of months for which the collection was done and the year that the information was collected.
- Observations: This is a free-text field where the researcher can keep track of additional information that is important for correctly understanding or interpreting the hypothesis.

The context description can be extended with other fields, specific to the given focus, that help the analysis of hypotheses. For example, using the focus of “Defects” we gathered information including:

- Number of Defects: Number of Analyzed Defects;
- Number of Modules: Number of Analyzed Modules.

Definitions: The extraction of all the terms used in the paper is important for the interpretation and analysis of the Structured Hypothesis Base. To address the challenges regarding mismatches among different sources, described in Section 2, the definitions from each paper must be recorded in a systematic way and then reconciled, to remove the chance that misinterpretations occur during the analysis of results.

There should be one record on this form for each measure collected in each source. Our experience is that the definitions are found in papers usually following the introduction.

The attributes of the definition form are:

- The term itself, for example: Module

- The description, as given in or implied from the text. For example: Module is a named subfunction, subroutine, or the main program of the software system.
- The original source: a reference to the paper.

The hypotheses, context descriptions and definitions are important for comprehending the possible measures and influencing factors for the focus of study. We are also working on ways of recording less directly connected contextual information, like software development challenges and lessons learned about empirical studies, that are also important to gather from the papers because they help the researcher to understand the context of the paper and of the conclusions. However, this work is ongoing and outside the scope of the current paper.

4.2 Integrating Information from Papers

Information integration consists of two steps: 1) Integration of definitions and 2) Formalization of the hypotheses based on the unified definition set. These steps may be iterated as needed.

The integration of definitions includes the creation of one unified definition schema, by finding the places in the definitions where they overlap, relating concepts that are semantically close via equivalence and subsumption relations (aligning) and checking the consistency, coherency and non-redundancy of the result [Klein01, Visser97]. The alignment of concepts is especially difficult, because this requires understanding of the meaning of concepts, but this step has been explored in ontological research, and many techniques and tools exist to facilitate its automation [Noy05].

The unified definition schema has the main goal of solving problems of conceptualization and explication mismatches pointed in section 2. Some of them will remain unsolved, and the researcher has to be aware of them before drawing conclusions from the analysis.

In the **hypothesis formalization** step, the hypotheses recorded from papers are formalized in such a way that structured analysis can be done. For the purposes of such analysis, we define hypotheses as being relationships between exactly two variables (one dependent and one independent). In cases where the hypotheses as stated in the source paper describes multiple relationships, the formalized hypotheses have to be divided into several smaller ones of two variables each.

The formalization of hypotheses is an important step for a better analysis of the results. It is important that the researcher, when executing this step, keeps in mind the focus of the study, because this will direct the choice of the dependent and independent variables. We have found, in our experience, two main types of hypotheses: A statement about a logical relationship between two variables (e.g. “As module size increases, the number of defects injected increases”) and a statement of a mathematical function that relates two variables (e.g. “Interface defects are expected to account for approximately 25% of all defects in a system”). This abstraction of hypothesis types is based on the hypotheses uncovered in our work to date; we do not claim that there may not be others, less often used, which we simply have not encountered yet.

However, these types seem to fit equally well with information coming from case studies (which look at detailed relationships over the life of a single project) as well as controlled experiments (which investigate whether a treatment, such as using a new process, has any impact on some success criteria).

A formalized version of either type has three major components: A dependent variable, an independent variable, and a relation between them. The variables themselves can be any entity discussed in the hypothesis (e.g. “module size,” “number of injected defects,” “number of interface defects”).

Each of the dependent and independent variables also has three additional fields which may be left blank: the **direction** of effect, the **magnitude** of effect, and **constraints** on the effect. These fields will only have values if the hypothesis contains appropriate information. In the case where the hypothesis describes a logical relationship, the direction and magnitude fields for each variable capture information about effects related to changes in value. For example, if the hypothesis is “as module size increases by 10%, the number of defects injected increases by 25%” then the independent variable of “module size” has the direction of “increases” and the magnitude of “10%.” In cases where the hypothesis describes a mathematical function, the direction and magnitude capture the functional calculation. For example, if the hypothesis is “25% of all system defects will be interface defects,” then the dependent variable of “interface defects” has the magnitude of “25%” but no direction.

A constraint on either variable simply captures information about the range of applicability of the variable. For the hypothesis “as the size of C modules increases by 10%...”, for example, the constraint field for the independent variable would contain the information that this is for the C programming language.

The **relation** field always describes the hypothesized relationship between the independent and dependent variables. Values can be: \approx (approximately equal), $<$, $>$, \leq , \geq , $=$, \neq , \Rightarrow (implies), \nRightarrow (does not imply).

As illustration, consider the following examples:

1. Plain text: Larger modules are more complex than smaller modules
 - a. Formalization:
 - i. Increased size in modules implies increased complexity**
 1. Independent Direction : Increased
 2. Independent Magnitude of Difference: N/A
 3. Independent Variable: Size
 4. Independent Context: Modules
 5. Effect : Imply
 6. Dependent Direction : Increased
 7. Dependent Magnitude of Difference: N/A
 8. Dependent Variable: Complexity
 9. Dependent Context: N/A
 - b. Source: [Basili84]
 - c. Level of Support: Significantly positive
 - d. Confidence: High
2. Plain text: 89% of defects can be corrected by changing only one module
 - a. Formalization:

1. Independent Direction: N/A
2. Independent Magnitude of Difference: N/A
3. Independent Variable: Requirements
4. Independent Context: Errors
5. Effect : Imply
6. Dependent Direction : Increased
7. Dependent Magnitude of Difference: Difficulty
8. Dependent Variable: Defects
9. Dependent Context: Correction

ii. Interface errors imply increased difficulty of defect correction

1. Independent Direction: N/A
2. Independent Magnitude of Difference: N/A
3. Independent Variable: Interface
4. Independent Context: Errors
5. Effect : Imply
6. Dependent Direction : Increased
7. Dependent Magnitude of Difference: Difficulty
8. Dependent Variable: Defect
9. Dependent Context: Correction

- b. Source: [Weiss85]
- c. Level of Support: Significantly positive
- d. Confidence: Low

After the formalization of the hypotheses we have a Structured Hypothesis Base that can now be used in a process of analysis.

4.3 Analysis and Interpretation

We argue that different types of analysis (top-down or bottom-up) can be done depending on the researcher's interests and how well-specified the problem can be. If, during the problem definition phase, the researcher can only state a general topic of interest rather than a specific hypothesis, he/she can do a bottom-up analysis, using the data from the papers analyzed to specify the set of important context variables. On the other hand, if the researcher already has detailed hypotheses in mind, with influencing factors already identified, he/she can do a top-down analysis by searching the Structured Hypothesis Base for any evidence describing those influence factors and their effects.

The input for the analysis is the Structured Hypothesis Base, also containing the specified definitions and context descriptions. The output of the process depends on the researcher's goals for the analysis, which may include: new hypotheses with supporting evidence, new beliefs that reflect expert opinion and may not have been previously recognized as such, or the refutation or confirmation of a set of initial hypotheses or folklore.

We assume that in some way the researcher can follow a core of structured steps to do the analysis. We propose here only one possibility of these steps, but recognize that the exact process followed is dependent upon the specific researcher doing the analysis:

1. Filter papers, if desired, selecting a subset that is most interesting for the current analysis. (For example, the researcher may select a subset of papers in which he/she has the most confidence in the results.) As filtering may introduce bias, we recommend this be done only according to objective criteria, such as including only papers in peer-reviewed journals. Any filtering should be reported clearly in the report of the study results, to allow the rationale to be subjected to peer review.
2. Filtering based on level of support and confidence. The researcher can filter the base, for example to remove hypotheses with the support level of “belief” (i.e. removing items with no empirical backing), giving more confidence in the results.
3. Organize the hypotheses that are still included in the analysis according to the dependent variables that they describe.
4. Choose a dependent variable to explore. This choice is based on the focus of analysis described in the definition of the problem.
5. Visualize influencing factors on the dependent variable. In this step the researcher can see which independent variables are related to the chosen dependent variable, and he/she can choose to explore some of these relations in more detail, for example, getting a sense if there is consensus about the direction of the relation, the magnitude of the effect, or the constraints under which those results are observed.
6. Based on context descriptions and definitions, certify that all mismatch problems are solved and that different hypotheses are not being combined. Filter problems not solved or hypotheses that can not be combined to draw conclusions.
7. Draw conclusions based on the patterns observed, create new hypotheses, or refute or confirm initial hypothesis or folklore.

The overall process is iterative, and the researcher can choose to traverse the steps in an order influenced by her/his specific goals, e.g. by doing first a bottom-up analysis to construct a detailed hypothesis, and then a top-down analysis to evaluate the body of evidence available in support of it.

While applying the process, the researcher may need to transform the values of information stored in the Structured Hypothesis Base, as described in Section 2 (e.g. expressing measures in KLOC in LOC instead, or estimating the total size of a program from the sizes of its component modules). This is done in an ad hoc manner, depending on the hypotheses that the researcher is trying to abstract, and so may be largely dependent upon the particular context. Some of the translations may be stored in the Structured Hypothesis Base if they appear likely to be reusable, but many will be too context-specific for this to be useful.

The bottom-up analysis considers only information found in papers published about the focus of study. That is, categories of phenomena and resulting hypotheses are built up based only on the information found in the literature search. The analysis

of the Structured Hypothesis Base is used to identify knowledge that cuts across multiple studies, for example in identifying variables that have an effect on the outcome.

The top-down approach considers that users may have intuition, folklore, or hypotheses; in short, some pre-existing opinions about which types of information and phenomena are interesting enough to collect information on. In this case, the analysis of the Structured Hypothesis Base yields a refutation or confirmation of these preconceptions.

Many visualization techniques can be applied to facilitate the exploration of the Structured Hypothesis Base [Oliveira03, Keim02]. In the next section we present a detailed example of applying the process in the area of software defects, in which we used Treemaps [Shneiderman91, Shneiderman92] to do the data exploration. Treemaps are a space-constrained visualization of hierarchical structures. They assign the available hypotheses into regions on the display, which allows researchers to see the hierarchies in datasets, in this context for example, to see the entire set of hypotheses organized first according to common dependent variables (i.e. specific aspects of the general topic of interest), and secondly in additional subsets around common independent variables (i.e. the factors that influence those dependent variables).

5 A bottom-up analysis case study

5.1 The Problem Definition

To allow a comparison between our more formal methodology and the example given of the folklore-based analysis in Section 3, as a topic of study we select software defects.

We follow a two-pronged approach to compare the results of the methodology to those of other approaches. First, we use a relatively small set of papers in which one of the authors had some experience in the application context, to investigate whether or not the results that could be obtained from applying the methodology are useful compared to an expert's opinion of the contribution of the source articles. These results are described in sections 5.2 and 5.3. Secondly, we reran the analysis on a larger set of papers done by other authors. In section 5.4 we compare the results from our formal approach on this set against an external expert's manual analysis of the knowledge that could be built across studies, to assess the completeness and repeatability of our approach when done without relying on a deep knowledge of the application context of the source articles.

5.2 Results from a Bottom-Up Analysis

To allow comparison with the folklore-based approach described in Section 3, we first apply the formal methodology using only the same three papers that were used in Section 3: Endres75, Basili/Pericone84, and Weiss/Basili85. The methodology was applied by one of the authors who did not have any direct experience with the papers. As was emphasized in our discussion of building up bodies of knowledge from fami-

lies of studies, each study in the family need not be a “strict” replication of one another, i.e., may have a different overall design and data collection strategy [Basili99].

67 plain text hypotheses were extracted from these papers. These were translated into 76 formalized hypotheses in the Hypothesis Base. During the analysis, we used only the hypotheses that were supported by empirical data, and we filtered the analysis to include only the relations that were supported by more than one paper. We used the Treemap tool³ to do the data mining exploration on data.

We first created a hierarchy based on dependent variable, and found three large groups of hypotheses related to defects, changes, and effort. Looking in detail at the variable of “defects” we created a new level of the hierarchy based on independent variable. Analyzing Figure 2, we can see the main independent variables related to defects: interface, size, effort, misunderstandings of specification, changes, and others. (In Figure 2, each of the small rectangles represents an individual hypothesis, with the text shown inside the block. Different shading represents the different source papers from which the hypotheses are drawn. The larger groupings of hypotheses collect together all hypotheses about a common independent variable.)

At this stage in the analysis, then, we have a systematic list of potential relations on the focus of interest (software defects), which may or may not represent new information to the researcher. Each relationship may be explored in more detail to examine the context in which the relationship held and the level of accumulated supporting data. Choosing to analyze “interface” as independent variable yields the results shown in Figure 3. We can see that two of the three papers fail to support the hypothesis that most defects are interface defects.

³ <http://www.cs.umd.edu/hcil/treemap/>

Defect						
The number of errors that occurs per module is, on the average, very small	Modified modules had a high percentage of errors of commission and a small percentage of errors of omission with a higher percentage of data and initialization errors	The most common type of error is one made in the design or implementation of a single component of the system.	Errors Corrections are more frequent than modifications	89 Percent of errors can be corrected by changing only one module	Interfaces appear to be the major source of errors regardless of the module type.	Interfaces are not a major source of errors
Modified and new modules behave similarly except for the type of errors prevalent in each and the amount of effort required to correct an error	More than 85% of the errors could be corrected by changing only one module per error	Modified and New Modules have similar number of Errors.	The greater the number of requirements changes does not imply greater number of requirements errors	Relatively few changes result in errors	Interfaces and requirements understandings errors are more difficult to correct than others	39% of the errors are Interface Errors
New modules have and equal number of errors of omission and commission and a higher percentage of control errors	For each type of error there exist different causes and therefore different measures have to be taken to prevent them. There is no single cure at all.		Despite a significant number of requirements changes imposed on some projects, there is no corresponding increase in frequency of requirements		85% of the errors could be corrected by changing only one module per error. It contradicts the preconception we have of the interdependence of the	
Misunderstandings of Specifications			Effort		Size	
The majority of errors are the result of functional specification (incorrect or misinterpreted)	Modified modules appear to be more susceptible to errors due to the misunderstanding of the specifications	Software development environment, the programming language used, misunderstandings of interface, project size, and misunderstandings of specifications do not	50% of the total effort required for error correction occurred in modified modules	Errors contained in modified modules requires more effort to correct than those in new modules, although the two classes contained approximately the same number of	Larger modules have more errors	There is a higher error rate in smaller sized modules
The majority of errors are the result of functional specification (incorrect or misinterpreted) and within this the majority of the errors involved modified modules	Misunderstanding of a module's specifications or requirements constituted the majority of detected errors.	36% of Errors have the source on Functional Specification Incorrectness or Misinterpreted. 24% of these errors occurred on Modified Modules and 12	Errors occurring in new modules required less effort to correct than those in modified modules	SEL programmers tend to spend their time finding and correcting many small errors made while designing or implementing single routines, rather than struggling with a few	Software development environment, the programming language used, misunderstandings of interface, project size, and misunderstandings of	

Figure 2 - Hypotheses about six independent variables affecting software defects.

Defect	
Interface	
Equal	Not imply
Interfaces appear to be the major source of errors regardless of the module type.	Interfaces are not a major source of errors
39% of the errors are Interface Errors	85% of the errors could be corrected by changing only one module per error. It contradicts the preconception we have of the interdependence of the modules in an operating system and the frequently heard observation that interfaces between modules, in particular are source of errors

Details of selected node

Attri... Value

De... []

Dir... []

Indi... []

Main Legend Filters Hierarchy

Nodes

Hide Filtered Hide Selected

Automatically Hide Filtered Nodes

19 nodes out of 573 nodes (3%) are shown

Original_Source_Hypothesis

- 1- Software Errors and Complexity: An
- 11- Software Requirements: Are they
- 12- Software Engineering in Avionics A
- 13- EXPERIMENTS WITH COMPUTERS
- 14- Identifying Error-Prone Software
- 15- QUANTITATIVE ASPECTS OF SOFT
- 16- Error Density and Size in Ada Soft
- 17- Persistent Software Errors - 1981
- 19- Prediction of Software Failures -
- 2- Evaluation of a Software Requirem
- 20- Collecting and Categorizing Softwa
- 3- Evaluating Software Development k
- 4- Quantitative Assessment of Softw
- 5- Error Localization During Softwa
- 6- Analyzing Error-Prone System Stru
- 7- Using the GQM Paradigm to Investi
- 8- An Analysis of Errors and Their Cau

Figure 3 - Analyzing specific relationships between interfaces and software defects.

To try to solve this discrepancy, we looked in more detail at the definitions of “interface” in each paper. In the two papers that do not support the hypothesis, the definition of interface defects is based on the number of components changed, i.e., a defect is an interface defect if more than one component (module) must be changed to fix the problem. In the third paper, interface defects are associated with structures existing outside the module's local environment. We can say then that, using the first definition, interface defects are not in fact the majority of the defects in this study.

Exploring the remaining independent variables, we could draw some additional conclusions:

- A. Interface defects are not the majority of defects, using the definition of interface defects that is based on the number of components changed, i.e., a defect is an interface defect if more than one component (module) must be changed to fix the problem [Supported by: Basili84, Weiss85 and Endres75; Contradicted by: None]
- B. Misunderstandings of specifications make up the majority of defects when the developers are not experienced with the application domain. [Supported by: Basili84 and Endres75; Contradicted by: Weiss85]
- C. Size, alone, is not a factor to determine defect-proneness. [Supported by: Basili84, Endres75, and Weiss85; Contradicted by: None]

Comparing these hypotheses to the results of the folklore-based approach, we see that hypothesis A above matches exactly to the refined hypothesis 2' in Section 3. Hypothesis B above also maps directly to hypothesis 1a in Section 3. Hypothesis C above is a concrete instance of the general case described in hypothesis 1 (namely, that there are consistent patterns that can be found describing software defect occurrences).

5.3 An expanded hypothesis base

One of the advantages of the formal approach is that the Structured Hypothesis Base is reusable and expandable as needed. To demonstrate, we can easily rerun the above analysis by adding additional papers into the Base.

Specifically, we next added defect data from the following papers, which brought the total number of hypotheses (including belief statements) to 223:

- 1- Rombach/Basili87: This study was conducted in the maintenance environment of a major computer company in commercial systems. [Rombach87]
- 2- Selby/Basili88: A single release of a code library tool. [Selby88]
- 3- Selby/Basili91: A single release of a code library tool. [Selby91]
- 4- Mashiko/Basili97: A set of four projects dealing with communication software. [Mashiko97]

Again, this large number of hypotheses was severely filtered for this example, by removing belief statements which did not have empirical data addressing them and limiting analysis only to those hypotheses which were addressed by more than one paper.

The analysis process helps to get an understanding of the solution space covered by these studies. For example, while Mashiko/Basili97 has some hypotheses about interfaces and defects, upon resolving the semantic mismatches it becomes apparent

that all of them are related to *external* interfaces (i.e. faults in the interface between the product and its external system) or *human* interfaces, thus covering quite separate phenomena from the other three papers which draw conclusions about *internal* interfaces (faults in the interface between modules).

The analysis of the remaining papers added no new hypotheses to the list about misunderstandings of specification, but increased the level of confidence in the hypothesis related to size by providing new hypotheses from additional contexts that matched the initial conclusions.

5.4 Using papers of other authors

To demonstrate that this methodology can use papers from other authors, we decided to replicate the study done by Brian Marick [Marick90], in which several different defect databases were reviewed in order to formulate statements about common issues in defect behaviors. Specifically, we added to our Hypothesis Base defect data from the following papers [Basili81, Bell76, Dniestrowski78, Glass81, Lipow79, Ostrand84, Potier82, Rubey75, Shen85, Withrow90]. (It is important to note that these papers were selected to match the analysis done by Marick and not to produce a comprehensive and up-to-date body of defect knowledge.)

- Basili81: The example application is an evaluation of the first stage of the A-7 flight software redevelopment: the production of the A-7 requirements document.
- Bell76 contains an analysis of two contexts. The first was a project in a graduate software engineering class, in which students wrote a set of software requirements and a preliminary system design for a Student Employment Information System (SEIS). The second case was a large (1 million machine instructions) real-time Ballistic Missile Defense (BMD) system being developed with a top-down approach.
- Dniestrowski78: The data reported in this paper is about an avionics digital flight control for a FALCON 20 variable stability aircraft. The source code consisted of: 800 variables declarations, 1715 constant declarations, 3615 assembly instructions, 3230 assembly instructions, 1050 assembly directives, 9100 documentation.
- Glass81: Operational software systems for military aircraft use was analyzed in this paper. Project A involved 150 programmers at the peak person-load, and contains about a half million instructions in the operational software alone. Project B involved 30 programmers and about 100,000 instructions.
- Lipow79: Data from three large software projects are examined and an analysis of the effect on defects of certain preventive and detective tools and techniques is presented.
- Ostrand84: In this paper they reported the results obtained from collecting defect data for nine months from a special-purpose editor project. The editor's source is about 10,000 lines of high-level language and assembly. The program design and coding were done by three programmers over ten months, after the initial specification had been completed. The implementation represents approximately 18 persons-months of effort.
- Potier82: The software used in the experiments belongs to a family of compilers for the real-time language LTR II developed for a wide range of target machines.

These compilers written in the LTR language consist of a common kernel specific to the LTR language and independent of the target machine, and a code generation part dependent on the target machine. The total number of lines is 64,939, with 1,402 defects found.

- Rubey75: The data presented in this paper were obtained from validation efforts, generally involving relatively small real-time control programs (averaging about 32K machine-language instructions).
- Shen85: This study was based on the detailed analysis of three products developed at IBM's Santa Teresa Laboratory. The products studied were developed and released since 1980. Product A is a software metrics counting tool written primarily in Pascal of total size is 7K TSI (Total Number of Source Instructions in thousands, excluding comments). Product B is a compiler written primarily in PL/S (a derivate of PL/1), of size 94K TSI. Product C is a database system written primarily in Assembly Language, the total size is 326K TSI.
- Withrow90: The object of this study was the Ada software for the command and control of a military communications system. The project is composed of 362 Ada Packages, containing 114,000 lines of code. The average package was 316 lines.

This brought the total number of papers used in the analysis to 18 (representing the work of 26 authors). From these papers, 332 plain text statements of potential hypotheses were found during the extraction step (as described in Section 4.1). During the integration step (Section 4.2), 141 of these statements were found to be at the level of belief/opinion only, and thus were excluded from further analysis; the remaining 191 plain text statements were formalized, resulting in 396 specific hypotheses. These metrics describe the Hypothesis Base as it was found at the beginning of the analysis.

During the analysis step (Section 4.3), visualization and data mining were applied on the formalized hypotheses, which were then traced back to the plain text hypotheses from the various sources. Using this analysis, 80 of the 191 plain text hypotheses were found to be related to the topics highlighted by Marick as the common points found in his analysis. (The remaining plain text hypotheses dealt with a wide range of other topics, such as maintainability, effort, reliability, reparability, memory constraints, reusability, and cohesion, which are outside the scope of our current analysis. For the purpose of comparability, we chose to analyze only the intersection of our hypothesis base with topics identified by Marick, rather than aim for completeness.)

As a result of this analysis, we found almost the same results as in the Marick study, which used a manual approach to abstracting information based on empirical evidence. Table 1 compares the results produced from the two methodologies:

No.	Conclusion	Support listed in Marick analysis	Support found in Hypothesis Base
1	Defects in programming logic (path selection) are common.	[Dniestrowski78], [Potier82], [Shen85],	All of the same, plus: [Basili84], [Ostrand84]

		[Glass81], [Lipow79]	(11 specific hypotheses)
2	Defects of omission are important	[Basili84], [Glass81], [Ostrand84]	All of the same, plus [Bell76]. (Also [Mashiko97], [Selby91] which were not in Marick's input set.) (29 specific hypotheses)
3	Data handling is more defect-prone than computation	[Basili84], [Potier82], [Shen85], [Glass81], [Lipow79], [Ostrand84]	All of the same. (8 specific hypotheses)
4	Modified modules are no more defect-prone than new modules	[Basili84] and additional sources that aren't clearly referenced.	Support from: [Basili84], [Shen85], [Selby91] (3 specific hypotheses)
5	There is evidence that both small and large modules are more defect-prone than medium-sized modules	[Basili84], [Shen85], [Withrow90]	All of the same plus [Endres75], [Potier82]. (Also [Selby91], [Selby88] which were not in Marick's input set.) (10 specific hypotheses)
6	In abstract descriptions (requirements/ functional specifications), beware of incorrectness, omissions and inconsistencies, in that order.	[Basili81], [Bell76]	All of the same plus [Basili84], [Weiss85], [Shen85]. (13 specific hypotheses)
7	Bug fixes cause a small number of new bugs	Unclear	[Basili84], [Weiss85], [Potier82], [Ostrand84] (6 specific hypotheses)

As a result of this comparison, we can see that the formal methodology had the advantage of making the links between the underlying data and the conclusions more explicit (e.g. conclusion number 7). In other cases, the conclusions were made more robust because more sources and more specific hypotheses resulting from them could be explicitly associated (e.g. conclusions number 1 and 2).

Through the comparison to the less formal approach, we did find a flaw in our methodology, namely that we were not analyzing and extracting information from figures. That is, information which happened to be encoded as figures or graphics, and which could be extracted from the paper, had not made it into our Hypothesis Base, resulting in missing support for some conclusions that Marick had found. As a

result, we decided to insert this step into the extraction of information from papers. Rerunning the methodology using this step resulted in the data as show in Table 1.

We could see in this case study that we can replicate results from other manual reports, but with an advantage of that the review was conducted according to explicit and reproducible methodology.

6 Summary

In order to provide useful, not to mention accurate, decision support about software development practices and their effects on projects, we need not only empirical studies, but empirical studies that cover a range of interesting development environments. Due to the wide range of influencing factors, and the fact that we cannot yet confidently specify them all ahead of time, we need a large dataset from which observations about influencing factors can be built bottom up. Building such an adequate dataset would be impossible if we can not make use of existing data, even if it was never designed to contribute to a larger empirical base.

We have demonstrated two processes with which we have been working to construct such an empirical base: One is a manual and informal approach, which relies on identifying common beliefs or “folklore” to identify useful hypotheses and a manual analysis of the information in papers to investigate whether there is support for those hypotheses; the other is a formal approach based around encoding the information in papers into a Structured Hypothesis Base that can then be searched to organize hypotheses and their associated support. There are some advantages of having a structured base, even when it does not produce many new results in comparison to the more informal analysis approach. The advantages are:

- a) The structured base is reusable, in the sense that it can be reused to test additional research questions that may not have been known at the time of its creation.
- b) If another researcher wants to validate the results, the process is repeatable.
- c) The analysis process can be undertaken by researchers who do not have in-depth knowledge about the context and details of the individual studies. The structured information helps the researcher to see the important information to do the analysis.

Having looked at a collection of datasets and abstracted up a set of conclusions on specific topics, we have shown that the formal methodology can produce useful and feasible results, especially when it is compared to the results output from the more manual, expert-based approach. The formalized approach, by relying on a reusable Hypothesis Base, is repeatable and also capable of producing a more thorough basis of support for hypotheses, including results from papers or articles that may have been overlooked or not considered by the experts.

In terms of its contribution to supporting the building of knowledge across studies, we feel that the work described in this paper demonstrates that:

- There is value in multiple studies for both supporting and not supporting hypotheses. There are several instances above where the conclusions from multiple datasets all point in the same direction, thus making the overall conclu-

sion much stronger than if it came from any single study in isolation. And, in several important instances, the results from additional studies identify important caveats by examining processes in new environments.

- Care must be taken to make sure that the objects of the comparison are actually like things that can support the conclusions being drawn.

In short, there are insights to be gained from the collection and analysis of defects according to different classification schemes, independent of the scheme. Our results show that interesting abstractions can be drawn by comparing defect information opportunistically, based on points of similarity where they occur.

Acknowledgements

This work is partially sponsored by NSF grant CCF0438933, "Flexible High Quality Design for Software."

References

- [Asgari05] Sima Asgari, Lorin Hochstein, Victor Basili, Jeff Carver, Jeff Hollingsworth, Forrest Shull, "Generating Testable Hypotheses from Tacit Knowledge for High Productivity Computing." In Proceedings of the Workshop on Software Engineering and High Performance Computing Applications (held at ICSE 2005). St. Louis, MO, USA.
- [Basili81] Basili, Victor R., Weiss, David , "Evaluation of Requirements Document by Analysis of Change Date", Proceedings of the Fifth International Conference on Software Engineering, pp 314-323, March 1981.
- [Basili84] Basili, V.R., Perricone, B. "Software Errors and Complexity: An Empirical Investigation", Communication of the ACM, vol. 27, no. 1, pp 42-52, January 1984.
- [Basili94] Basili, V. R., Caldiera, G., and Rombach, D. H., "The Experience Factory," in *Encyclopedia of Software Engineering - 2 Volume Set*, 1994, pp. 469-476.
- [Basili99] Basili, V. R., Shull, F., Lanubile, F. "Building Knowledge through Families of Experiments", IEEE Transactions on Software Engineering, Vol. 25, No. 4, pp. 456-473, July 1999.
- [Bell76] Bell, T.E., Thayer, T.A.. "Software Requirements: Are They Really a Problem?". Proceedings of the 2nd International Conference on Software Engineering, pp. 61-68, IEEE Press, 1976.
- [Brooks96] Brooks, A., Daly, J., Miller, J., Roper, M., Wood, M. (1996). "Replication of experimental results in software engineering". Technical Report ISERN-96-10, Department of Computer Science, University of Strathclyde, Glasgow.

- [Dniestrowski78] Dniestrowski, A., Guillaume, J.M., Mortier, R. “Software Engineering in Avionics Applications”. Proceedings of the 3rd International Conference on Software Engineering, pp. 124-131, IEEE Press, 1978.
- [Endres75] Endres, A. “An analysis of Errors and their Causes in System Programs”, Proceedings of the International Conference on Reliable Software, pp. 327-336, April 1975, Los Angeles, CA.
- [Gay96] Gay, L. R. “Educational Research: Competencies for Analysis and Application”, Fifth Edition (1996, Prentice-Hall, ISBN 0-02-340814-6). Pp. 61-66, 72-73.
- [Glass81] Glass Robert L., “Persistent Software Errors,” *IEEE Transactions on Software Engineering*, March 1981.
- [IEEE87] IEEE. Software Engineering Standards. IEEE Computer Society Press, 1987.
- [Keim02] Keim, D. “Information visualization and visual data mining”. *IEEE Transactions on Visualization and Computer Graphics*, 8(1):1-8, January 2002.
- [Kitchenham04] Kitchenham, Barbara A., Dybå, Tore , Jørgensen, Magne, “Evidence-Based Software Engineering,” Proceedings of the 26th International Conference on Software Engineering (ICSE 2004), Edinburgh, Scotland, 23-28 May, IEEE Computer Society, 2004, pp. 273-281.
- [Klein 01] Klein, M, “Combining and relating ontologies: an analysis of problems and solutions”. A. Gomez-Perez, M. Gruninger, H. Stuckenschmidt, and M. Uschold, editors, Workshop on Ontologies and Information Sharing, IJCAI'01, Seattle, USA, Aug. 4--5, 2001
- [Lipow79] Lipow, Myron, Thayer, Thomas A., “Prediction of Software Failures”, *Journal of Systems and Software*, Vol 19, N° 7, July, 1979, pp 71-75.
- [Marick90] Marick, Brian, “A Survey of Software Fault Surveys”, Technical Report UIUCDCS-R-90-1651, University of Illinois, 1990.
- [Mashiko97] Mashiko, Y., Basili, V.R. “Using the GQM Paradigm to Investigate Influential Factors for Software Process Improvement”, *The Journal of Systems and Software*, Volume 36, Number 1, pp 17-32, January 1997.
- [Miller05] Miller, James. “Replicating software engineering experiments: A poisoned chalice or the Holy Grail”. *IST*, 47(4), 2005, pp. 233-204.
- [Noy05] Noy, Natasha F., Stuckenschmidt, Heiner. “Ontology Alignment: An annotated Bibliography”. *Semantic Interoperability and Integration 2005*
- [Oliveira03] Oliveira, M. C. F.; Levkowitz, Haim. “From Visualization to Visual Data Mining: A Survey”. *IEEE Transactions on Visualization and Computer Graphics*, United States, v. 9, n. 3, p. 378-394, 2003.
- [Ostrand84] Ostrand Thomas J., Weyuker, Elaine J.. “Collecting and Categorizing Software Error Data in an Industrial Environment”. *Journal of Systems and Software*, Vol. 4, 1984, pp. 289-300.

- [Potier82] Potier, D., Albin, J.L., Ferreol, R., Bilodeau, A.. "Experiments with Computer Software Complexity and Reliability". Proceedings of the 6th International Conference on Software Engineering, pp.94-103, IEEE Press, 1982.
- [Rombach87] Rombach, H. D., Basili, V.R., "A Quantitative Assessment of Software Maintenance: An Industrial Case Study", Conference on Software Maintenance, Austin, Texas, September 1987.
- [Rubey75] Rubey, Raymond J., "Quantitative Aspects of Software Validation". Proceedings of the 1975 International Conference on Reliable Software, in SIGPLAN Notices, vol. 10, No. 6, pp. 246-251, June, 1975.
- [Selby88] Selby, Richard W., Basili, Victor R., "Error localization during software maintenance: Generating hierarchical system descriptions from the source code alone". In Proceedings of the Conference on Software Maintenance, Phoenix, AZ, October 1988.
- [Selby91] Selby, R.W., Basili, V.R. "Analyzing Error Prone System Structure", IEEE Transactions on Software Engineering, pp. 141-152, February 1991.
- [Shen85] Shen, V.Y. & Yu, T. J. & Thebaut, S.M. & Paulsen, L.R. "Identifying Error-Prone Software - An Empirical Study", IEEE Transactions on Software Engineering, Vol. 11, n°4, p. 317-324, April 1985.
- [Shneiderman91] Shneiderman, B., Johnson, B. "Treemaps: A Space-Filling Approach to the Visualization of Hierarchical Information Structures". Proc. of IEEE Information Visualization, pp. 275-282, 1991.
- [Shneiderman92] Shneiderman, B. "*Tree visualization with tree-maps: 2-d space-filling approach*". ACM Transactions on Graphics, vol. 11 , n°. 1, pp. 92-99, 1992.
- [Visser97] Visser, Pepijn R. S., Jones, Dean M. , Bench-Capon, T. J. M., Shave, M. J. R.. "An analysis of ontological mismatches: Heterogeneity versus interoperability". AAAI 1997 Spring Symposium on Ontological Engineering, Stanford, USA, 1997.
- [Weiss85] Weiss, D.M., Basili, V.R. "Evaluating Software Development by Analysis of Changes: The Data from the Software Engineering Laboratory", IEEE Transactions on Software Engineering, pp 157-168. February 1985.
- [Withrow90] Withrow, C., "Error density and size in Ada software, IEEE Software, pp. 26-30, Jan. 1990.