

# Documenting Design-Pattern Instances: a Family of Experiments on Source Code Comprehensibility

GIUSEPPE SCANNIELLO, University of Basilicata  
 CARMINE GRAVINO, University of Salerno  
 MICHELE RISI, University of Salerno  
 GENOVEFFA TORTORA, University of Salerno  
 GABRIELLA DODERO, Free University of Bozen-Bolzano

Design-patterns are recognized as a means to improve software maintenance by furnishing an explicit specification of class and object interactions and their underlying intent [Gamma et al. 1995]. Only a few empirical investigations have been conducted to assess whether the kind of documentation for design-patterns implemented in source code affects its comprehensibility. To investigate this aspect, we conducted a family of four controlled experiments with 88 participants having different experience (i.e., professionals and Bachelor, Master, and PhD students). In each experiment, the participants were divided into three groups and were asked to comprehend a non-trivial chunk of an open-source software system. Depending on the group, each participant was, or was not, provided with graphical or textual representations of the design-patterns implemented within the source code. We graphically documented design-pattern instances with UML class diagrams. Textually documented instances are directly reported in source code as comments. Our results indicate that documenting design-pattern instances yields an improvement in correctness of understanding of source code for those participants with an adequate level for experience.

Categories and Subject Descriptors: D.2.0 [Software Engineering]: General; D.2.7 [Software Engineering]: Maintenance

General Terms: Design; Documentation; Human Factors

Additional Key Words and Phrases: Design-Patterns, Controlled Experiment, Maintenance, Replications, Software Models, Source Code Comprehension

## ACM Reference Format:

ACM Trans. Softw. Eng. Methodol. V, N, Article A (January YYYY), 40 pages.  
 DOI = 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

## 1. INTRODUCTION

In the context of building and town design, Alexander et al. [1977] asserted that a “pattern describes a problem which occurs over and over again in our environment, and then describes the core of the solution to that problem, in such a way that you can use this solution a million times over, without ever doing it the same way twice”. This definition of pattern also holds for the design and development of object-oriented software [Gamma et al. 1995]. In such a context, a design-pattern includes a name, an intent, a problem, its solution, some examples, and more. The core of both kinds of design-pattern definitions concerns the solution to a problem in a given context. In software development, such a solution is named design-pattern instance or also design motif [Guéhéneuc and Antoniol 2008].

---

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

© YYYY ACM 1049-331X/YYYY/01-ARTA \$15.00

DOI 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

ACM Transactions on Software Engineering and Methodology, Vol. V, No. N, Article A, Pub. date: January YYYY.

Copyright © held by the Association for Computing Machinery, Inc. (ACM).

Authors version

The publisher version is available at <https://dl.acm.org/doi/10.1145/2699696>

Gamma *et al.* assert that developers would benefit from the documentation of design-patterns to comprehend source code, making its modification and evolution an easier task [Gamma et al. 1995]. Although there are a number of empirical investigations on design-patterns (e.g., [Prechelt et al. 2001; Bieman et al. 2003; Cepeda Porras and Guéhéneuc 2010; Jeanmart et al. 2009; Khomh and Guéhéneuc 2008; Di Penta et al. 2008; Vokác 2004; Vokác et al. 2004; Krein et al. 2011]), only a few evaluations have been conducted to study this conjecture, namely there is a lack of investigations to assess the benefit of documenting design-pattern instances on source code comprehensibility in the context of software maintenance and evolution [Krein et al. 2011; Gravino et al. 2012; Prechelt et al. 2002; Gravino et al. 2011].

To obtain an initial insight into the usefulness of the documentation of design-pattern instances (or simply pattern instances or instances, from here on), some of the authors (the first four) of this paper carried out two experiments as a pilot study [Gravino et al. 2011]. These experiments aimed to assess the benefit of documenting instances with respect to not documenting them at all. The first experiment studied instances graphically documented by using the class diagrams of the Unified Modeling Language (UML) [OMG 2005], while instances textually documented were considered in the second experiment. These results provided evidence that documenting pattern instances improves source code comprehensibility with respect to not documenting them at all. This effect is even clearer in the case of graphically documented instances. To further investigate the effect of the kind of design-pattern instance documentation on source code comprehensibility and to study any possible effect of the developers' experience [ISO 1991], we have performed a series of empirical investigations as a family<sup>1</sup> of controlled experiments. This family includes four experiments with 88 participants having different levels of experience (i.e., professionals and Bachelor, Master, and PhD students). The original experiment was carried out with professional developers and preliminary findings are shown in [Gravino et al. 2012]. The results of this experiment suggest that professionals provided with graphically and textually documented pattern instances achieved a significantly better comprehension of source code than those provided with source code alone. Therefore, we provide here the following new contributions: (i) three new experiments have been added; (ii) the analysis procedure of the already published experiment has been updated; (iii) the effect of the participants' experience has been analyzed; (iv) a more thorough discussion of the obtained results is reported; and (v) practical implications for the results are presented and discussed. Our family of experiments does not include the pilot study mentioned above [Gravino et al. 2011].

The experiments in our family were grounded on theoretical foundations to avoid casting serious doubts on their validity. To this end, we took advantage of the framework by Aranda et al. [2007], which has been properly adapted here to better handle the problem at hand. In addition, each experiment was carried out by following recommendations provided by Juristo and Moreno [2001] and by Wohlin et al. [2012].

The paper is organized as follows. In Section 2, we discuss how we deal with comprehensibility in our family of experiments. In Section 3, we present the design of the experiments according to the guidelines suggested by Jedlitschka et al. [2008]. In Section 4, we show the achieved results, while their discussion is presented in Section 5. Related work is shown in Section 6, while remarks and future work conclude the paper.

<sup>1</sup>A family is composed of multiple similar experiments that pursue the same goal in order to build the knowledge that is needed to extract and to abstract significant conclusions [Basili et al. 1999]. Families of experiments allow researchers to answer questions that are beyond the scope of individual experiments and to generalize findings across studies. In addition, families of experiments can contribute to the conception of hypotheses that may not be suggested by individual experiments.

## 2. DEALING WITH SOURCE CODE COMPREHENSIBILITY

Any study on comprehensibility should consider many variables [Aranda et al. 2007]. Some variables are affected by comprehensibility (*affected variables* or *dependent variables*), while others affect comprehensibility (*affecting variables* or *moderator variables*). In Table I, some dependent and moderator variables are highlighted with respect to source code comprehensibility. In particular, in this table we summarize our modification to the framework by Aranda et al. [2007].

The larger the number of controlled dependent/moderator variables, the better it is. In controlled experiments some factors may be possible to control, while others may be difficult or impossible to control [Wohlin et al. 2012]. To increase the confidence in the outcomes, we had to select and study a number of dependent and moderator variables. The following are the moderator variables considered in our family of experiments:

- **Type of task.** Tasks performed by people are facilitated or hindered at varying degrees. Comprehensibility for information search or maintenance tasks requires information integration and this could lead to different results. We focus here on aspects related to the general comprehension of source code (e.g., question Q1 in Appendix A) and to the comprehension of source code with the goal of performing modification tasks (e.g., question Q12 in Appendix A).
- **Experience.** Finding competent people is a challenge in most software engineering studies [Bergersen et al. 2014]. In our family of experiments, we conducted four controlled experiments. The participants in each experiment had a similar level of experience with design-pattern development. We classified participants on the basis of their knowledge on high- and low-level design of object-oriented software systems and software development and maintenance.

The dependent variables are:

- **Correctness of understanding.** The degree to which a person correctly answers questions about source code. This variable estimates the comprehensibility achieved by people. To assess correctness of understanding, we used the comprehension questionnaire in Appendix A.

Table I. Moderator (Affecting) and Affected (Dependent) Variables

Kind of Variable	Name	Description
Moderator (Affecting)	Type of task	Tasks performed by people are facilitated or hindered at varying degrees. Comprehensibility requires the integration of information (e.g., source code and documentation) in the people mental model. This could lead to different results.
	Experience	Previous experience with the kind of pattern instance documentation and with programming and maintenance.
	Domain knowledge	Previous knowledge with the solution domain of the system under study and with its application domain.
	Problem size	Different notations might scale up with variable degrees of success.
Dependent (Affected)	Correctness of understanding	The degree to which a person can answer questions about a comprehension task on source code.
	Time	Time required to accomplish a comprehension task on source code.
	Confidence	Subjective confidence that people display regarding their own understanding.
	Perceived difficulty	Subjective judgment about the ease in obtaining information through the provided artifacts (e.g., analysis and design models).

- **Time.** The time required to answer questions about source code. This variable represents an approximation for effort. This is almost customary in literature (e.g., [Ricca et al. 2014]) and it is compliant with the ISO/IEC 9126 standard [ISO 1991], where effort is the productive time associated with a specific project task.
- **Confidence.** The subjective confidence that a person shows about his/her own understanding of the furnished source code. It is expressed for each question in the comprehension questionnaire.
- **Perceived difficulty.** It concerns the subjective judgment about the ease in obtaining information through the documentation of a design-pattern instances (if present) and/or source code.

The theoretical framework by Aranda et al. [2007] was originally conceived to empirically evaluate model comprehensibility. It consists of a sequence of guidelines organized in the following order:

- (1) **Selecting the notation.** There are several decisions to be taken to reduce as much as possible biases in the experimentation.
- (2) **Articulating the underlying theory.** The assumptions under which a notation is useful should be made clear.
- (3) **Formulating the claims of the notation.** The underlying theory is expressed as a set of claims regarding comprehensibility.
- (4) **Choosing a control.** A baseline for the comparison has to be chosen.
- (5) **Turning the claims into hypotheses.** Claims are turned into testable hypotheses that cover most of the dependent variables.
- (6) **Informing the hypotheses.** The defined hypotheses are informed using frameworks from cognitive science theory.
- (7) **Designing and executing the study.** Whatever is the empirical method used for comprehensibility evaluation, some design choices should be taken into account (e.g., comprehension questionnaire).

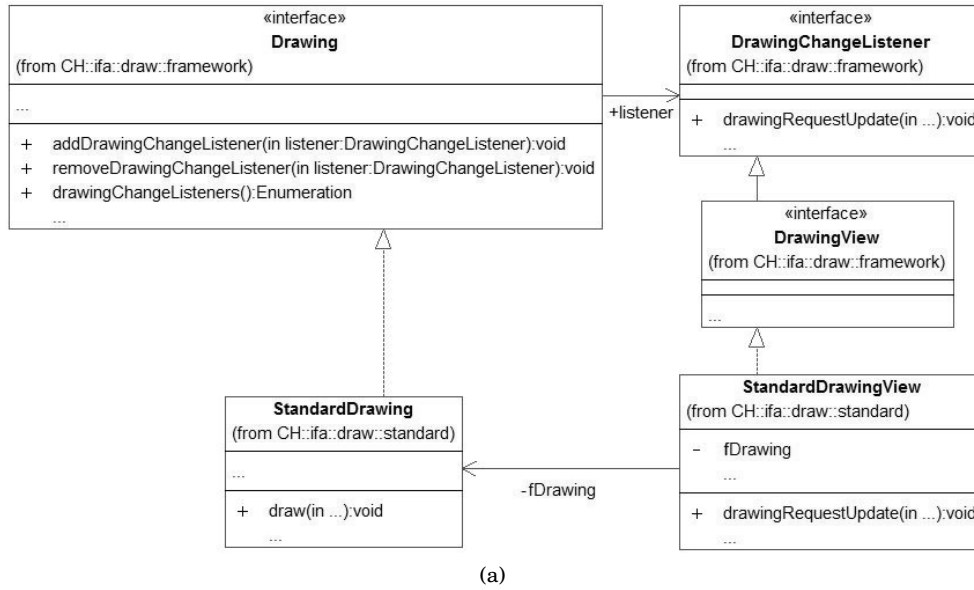
In the following subsections, we show how we have instantiated the guidelines in between *Selecting the notation* and *Informing the hypotheses*, while *Designing and executing the study* is described in Section 3 (i.e., The Family of Experiments).

## 2.1. Selecting the notation

Some notations have been proposed for documenting design-patterns and their instances in source code. Many of these notations are based either on UML or on textual descriptions [Gamma et al. 1995; Heer and Agrawala 2006]. Those based on UML are widely used in software engineering academic and industrial courses to teach design-pattern based development [Gamma et al. 1995; Bruegge and Dutoit 2003]. In our family of experiments, we investigated source code comprehensibility when pattern instances implemented in it are documented graphically, as proposed by Gamma et al. [1995] using UML class diagrams (see for example Figure 1(a)), or textually, as done in JHotDraw<sup>2</sup> (see for example Figure 1(b)). In this latter case, textual descriptions are directly reported in source code as comments (e.g., [Guéhéneuc and Antoniol 2008]), when design-patterns are intentionally used. In many cases, identifier names (e.g., class names) suggest the pattern instances implemented in source code. We selected these two notations because both are widely used and known in academic and in industrial contexts, so improving external validity. The kind of documentation for design-

<sup>2</sup>It is a two-dimensional graphics framework for structured drawing editors. It can be downloaded at [www.jhotdraw.org](http://www.jhotdraw.org).

pattern instances represents the manipulated factor (or independent variable or main factor) in our family of experiments.



```

1: /*
2:  * @(#)DrawingChangeListener.java 5.1
3:  *
4:  * Observer Pattern
5:  * Subject: Drawing
6:  * ConcreteSubject: StandardDrawing
7:  * Observer: DrawingChangeListener
8:  * ConcreteObserver: StandardDrawingView
9:  *
10: */
11:
12: package CH.ifa.draw.framework;
13: ...
14: public interface DrawingChangeListener
15:     extends EventListener {
16: ...
17:     /**
18:      * ...
19:      */
20:     public void drawingRequestUpdate(...);
21: }
22:
23:

```

(b)

Fig. 1. A sample of an instance for the Observer design-pattern: graphically documented (a) and textually documented (b)

## 2.2. Articulating the underlying theory

The underlying theory is ideally obtained from the literature by identifying target users, domains, problem sizes, and required *domain experience* of people who will use

a notation. The effect of design-pattern instance documentation on source code comprehensibility has been scarcely investigated, therefore, we could not rely on the literature. Based on our experience, knowledge, and results from our pilot study [Gravino et al. 2011], we postulated the following claims:

- (1) source code comprehensibility increases when documented pattern instances are used as a complementary information;
- (2) graphically documented pattern instances can be exploited by any user with experience in UML;
- (3) textually documented instances require the target user to have some experience in dealing with source code written by other developers.

Both kinds of documentation are mostly used as a communication mechanism among developers to make a software implementation (and implicitly the rationale behind it) clearer and more understandable.

### 2.3. Formulating the claims

Since source code comprehension increases when documentation is given (e.g., [Arisholm et al. 2006; Budgen et al. 2011; Prechelt et al. 2002; Scanniello et al. 2013]), we claimed that correctness in understanding unknown source code increases when pattern instances are documented and given to developers dealing with comprehension tasks. The use of that documentation should reduce possible misunderstandings among developers even when they have different level of experience.

### 2.4. Choosing a control

Our natural choice for the control/baseline was source code without documented instances, and without any reference to pattern instances from source-code identifiers. We opted for this strategy because it is rather common that professional developers and students unintentionally (or accidentally) use design-patterns when writing source code (e.g., [Guéhéneuc and Antoniol 2008]). Therefore, when a different developer has to comprehend that source code, s/he has to deal with pattern instances that are not documented at all.

### 2.5. Turning the claims into hypotheses

A graphical representation of a pattern instance might show a superset of the information provided by the corresponding textual representation. For example, in both representations the roles each class plays within the pattern instance might be indicated, while in a textually documented instance the relationships among abstract and concrete classes and interfaces might not be reported. Figure 1(a) shows an example of a graphically documented instance of the Observer design-pattern [Gamma et al. 1995] within the source code of JHotDraw v5.1. Figure 1(b) shows how the same instance is explicitly reported within the source code as comment.

The graphical instance shows more information and thus it should improve source code comprehensibility. For example, from the class diagram in Figure 1(a), we understand that, when a drawing is changed, all its views are updated. More experienced developers might find the additional information provided by graphically documented instances to be of little use, because they can directly deduce such an information from the name of the design-pattern and from the role of each class and interface. On the other hand, developers with different levels of experience (e.g., less experienced) would receive some help in using graphically documented instances even if the use of this kind of documentation might require more cognitive resources to switch the perception mode from text to graphic and back.

## 2.6. Informing the hypotheses

We consider external cognition to have particular relevance in our study. This branch of cognitive science treats humans and the artifacts they use to solve a given task as a single cognitive entity. The perspective is that artifacts are part of problem-solving resources, and thus they improve human reasoning. Representations might improve reasoning in several ways [Bauer and Johnson-Laird 1993]. According to Scaife and Rogers [1996], documented instances can reduce cognitive effort by putting knowledge in the representation, rather than in the head of the developer (i.e., offloading). That is, the least information has to be kept in mind and the fewer rules are needed to process them, the better it is. This is especially true for graphically documented instances, because they show relationships among classes and roles that each class plays [Guéhéneuc and Antoniol 2008]. When using textually documented instances, more experienced developers might not be hindered in inferring this information. Indeed, these developers might have concerns going from source code to documentation and vice-versa when using graphically documented instances. Approaching source code comprehension tasks in such a way might affect comprehensibility when a developer is not accustomed to use documentation that is kept externally from source code.

## 3. THE FAMILY OF EXPERIMENTS

The software engineering community has been embracing replications more readily (e.g., [da Silva et al. 2014]). Even with the increased interest in replications, there is no agreement yet on terminology, typology, purposes, operation, and other replication issues [Kitchenham 2008; Shull et al. 2008]. The lack of shared terminology allows authors to use different definitions for the same kind of replication and to use the same definition to refer to different kinds of replication [Baldassarre et al. 2014]. In general, we can define a replication as the repetition of an experiment, either as closely following the original experiment as possible, or with a deliberate change to one or several of the original experiment parameters [Gómez et al. 2010]. In particular, there are two important factors that underlie the definition of replications: *the procedure* (i.e., the steps followed) and *the researchers* (i.e., the experimenters conducting the replication). As for the procedure, the kinds of replications range in between *close* and *conceptual*. A close replication is a replication whose procedure is as close to the original procedures as possible [Shull et al. 2008], while a replication is conceptual if the research question is the same, but the experimental procedure is completely different from that of the original experiment. As for the researchers, we can distinguish between *internal* and *external replications*. Internal replications are conducted by the same group of researchers as the original experiment [Mendonça et al. 2008], while external replications are executed by different experimenters also using the same procedure as the original experiment. In case of external replications, the results might slightly be affected by experimenters' biases [Juristo and Moreno 2001; Shull et al. 2008].

Experiments need to be replicated in different contexts, at different times, and under different conditions before they can produce generalizable knowledge [Shull et al. 2008]. From the scientific and industrial perspectives, we can identify two primary motivations for conducting replications: (i) they are necessary to solve problems and collect evidences because they bring credibility to a given research and (ii) they are valuable because they provide evidence on the benefits of a software engineering practice thus allowing industrial stakeholders to use this information to support adoption decisions [Baldassarre et al. 2014; Colosimo et al. 2009; Pfleeger and Menezes 2000].

Differently from a single replication, a family of controlled experiments can contribute to the conception of important and relevant hypotheses that may not be suggested by an individual experiment/replication (e.g., any possible effect of the develop-

Table II. Participants in the family of experiments grouped by experiment

Experiment	Context	Description	Participants	Type	Year
E1-Prof	Practitioners	Original experiment	25	-	2011
E2-UniSA	University of Salerno	Internal replication	25	2nd year Master students	2012
E3-UniBAS	University of Basilicata	Internal replication	23	3rd year Bachelor students	2012
E4-UniBZ	Free University of Bozen-Bolzano	External replication	15	2nd/3rd year PhD students	2013

ers' experience) [Basili et al. 1999]. This is why we carried out a family of controlled experiments composed of: the original experiment, two internal replications, and one external replication.

In Table II, we summarize information on the four experiments we conducted in our family. The original experiment, named E1-Prof, was carried out with professional developers in 2011. Some results from this experiment were preliminarily presented in [Gravino et al. 2012]. All the participants in E1-Prof had at least a Bachelor degree in Computer Science. Among the participants in E1-Prof, there were neither PhD students nor PhD graduates. By varying participants' experience, E1-Prof was replicated in 2012 at the University of Salerno with 25 second year students from the Master program in Computer Science. This experiment is named E2-UniSA. In the same year, we conducted the second replication named E3-UniBAS. The participants in this replication were 23 third year students from the Bachelor program in Computer Science at the University of Basilicata. The third replication was E4-UniBZ. It was performed at the Free University of Bozen-Bolzano with 15 PhD students in Computer Science. We considered this experiment an external replication because an investigator (the fifth author) performed this experiment on the basis of an experimental package provided by the original investigators (the first four authors). The new investigator participated neither in the design of the experiments nor in the preparation of the experimental material. The experiments E2-UniSA, E3-UniBAS, and E4-UniBZ are presented for the first time in this paper.

Features that made the experiments in our family distinct from the pilot study [Gravino et al. 2011] were: (i) the experiment design and (ii) the kind of the involved participants<sup>3</sup>. We also renewed and improved the experimental material.

For replication purposes, we made available on the Web an experiment package, the raw data, and a technical report [Scanniello et al. 2014b].

### 3.1. Goal

According to the Goal Question Metrics (GQM) template by Basili and Rombach [1988], the goal of our family of experiments can be formalized as follows:

**Analyze** source code comprehensibility **for the purpose** of evaluating graphically and textually documented design-pattern instances **with respect to** correctness of understanding, time, confidence, and perceived difficulty **from the point of view** of the researcher and **from the point of view** of the practitioner **in the context of** object-oriented software.

### 3.2. Context Selection

For each experiment, its participants had the following characteristics:

<sup>3</sup>The participants in the first experiment of the pilot study were Master students in Computer Science from the University of Basilicata, while the participants in E3-UniBAS were students from a Bachelor program in Computer Science from the same university. In the second experiment of the pilot study, the students were enrolled to the 1st year of a Master program in Computer Science at the University of Salerno, while in our family of experiments the Master students came from the same university, but they were enrolled to the 2nd year of the same program.



- **E1-Prof.** The participants were Italian software professionals. They have been working for software companies in the contact network of the research groups of the original investigators. These companies were partners in research projects, or hosted students from the universities of Basilicata and Salerno for internships or employed alumni from either universities.
- **E2-UniSA.** The participants were students from the Software Engineering II course of the Master program at the University of Salerno. They all passed the following courses: Software Engineering I, Object Oriented Programming I and II, and Data Bases. In particular, they knew the basics of high- and low-level design of object-oriented software systems (based on UML), software development, and software maintenance. The majority of the participants took an internship in a company, as the final part of their Bachelor degree. The experience level of these students is considered inferior to that of the participants in E1-Prof.
- **E3-UniBAS.** The participants were students from the Software Engineering I course of a Bachelor program at the University of Basilicata. They all passed the following courses: Procedural Programming, Object Oriented Programming I, and Data Bases. In the Software Engineering I course, they learned object-oriented modeling and UML. They had a limited experience in developing and maintaining non-trivial software systems. These students can be considered to be the participants in our family of experiments with the lowest level of experience.
- **E4-UniBZ.** The participants were students from an international PhD course at the Free University of Bozen-Bolzano. They took their Master degree in Computer Science from universities worldwide. In particular, they came from: Italy (6), Lithuania (2), Russia (2), Romania (1), Morocco (1), Serbia (1), Thailand (1), and Brazil (1). The participants in E4-UniBZ can be considered those with the highest level of experience in our family of experiments or at least their experience is not inferior to that of the participants in E1-Prof.

The prior knowledge and experience of the participants in each experiment can be considered rather homogeneous (especially in those experiments conducted with students [Verelst 2004]). The classification of participants in levels of experience follows an approach similar to that proposed by Ricca et al. [2010] and Abrahão et al. [2013].

The participation in the experiments was on a voluntary basis. Each participant took part in only one experiment. Results did not influence the grades in the Software Engineering courses where the replications E2-UniSA and E3-UniBAS took place. Professionals participated in the experiment as part of their work hours. This choice was made to encourage professionals to participate. No participants were obliged to participate in the experiments.

We selected a chunk (i.e., vertical slice) of JHotDraw v5.1 that included: (i) a non-trivial number of instances and (ii) well-known and widely adopted design-patterns. In the source code selection process, we have taken into account a trade-off between complexity of the implemented functionality in the chunk and the effort to comprehend source code. One of the authors (i.e., the third) detected and documented the pattern instances present in the source code of the experimental object. To this end, both the JHotDraw documentation and the PMARt data set [Guéhéneuc 2007] were used. This allowed us to document both intentional and unintentional instances (see Appendix B, where it is also possible to understand which classes participated in more than one instance). We also took care to not exclude classes in the selected source code that were related to other classes, which played a role in a pattern instance. The selection process was iterative and allowed us to select a total number of 10 instances. The following instances were present in the selected chunk of JHotDraw: State, Adapter, Strategy, Decorator, Composite, Observer, Command, Template Method, and Prototype. For the

State design-pattern, there were two instances, while only one instance was present for all the remaining design-patterns. The chunk consisted of 1326 Lines of Code, 26 Classes, and 823 Lines of Source Code Comment.

The use of incomplete documentation and of a subset of the entire software on which a maintenance operation impacts is quite common in the software industry [Scanniello et al. 2014a]. For example, it happens when only a part of the documentation actually exists (e.g., in lean development processes) or is updated [Bruegge and Dutoit 2003; McDermid 1991], or only a subset of the entire documentation is useful to perform a given maintenance operation [Asuncion et al. 2007; Lindvall and Sandahl 1996].

As for E1-Prof, E2-UniSA, and E3-UniBAS, we translated the experimental material including the comments in the original source code of JHotDraw from English into Italian. This allowed us to reduce biases related to different levels of participants' familiarity with English. For E4-UniBZ, the comments were not translated because the official language of the PhD program at the Free University of Bozen-Bolzano is English. As for the participants provided with source code alone and graphically documented instances, we removed any possible information on pattern instances present in comments and identifiers (e.g., *CompositeFigure* was renamed *ArrayFigure*). We took this decision to analyze only the effect of pattern instance documentation.

### 3.3. Selection of the Variables

To determine Time, we used the overall time (expressed in minutes) to answer the comprehension questionnaire. The higher the value of Time, the higher the effort to accomplish the comprehension tasks.

To quantify Correctness of understanding, we used an approach based on the information retrieval theory [Salton and McGill 1983]. We defined:  $A_{s,i}$  as the set of string items provided as answer to the question  $i$  in the comprehension questionnaire by the participant  $s$ ;  $C_i$  as the correct set of items expected for the question  $i$  (i.e., the oracle). Then, we computed precision and recall for each answer to a given question in the questionnaire as follows:

$$precision_{s,i} = \frac{|A_{s,i} \cap C_i|}{|A_{s,i}|} \quad (1)$$

$$recall_{s,i} = \frac{|A_{s,i} \cap C_i|}{|C_i|} \quad (2)$$

Precision (i.e., the fraction of items in the answer that are correct) and recall (i.e., the fraction of correct items in the answer) measure accuracy and completeness of the answer to a given question, respectively. To get a trade-off between accuracy and completeness, we used the balanced F-measure of precision and recall:

$$F-Measure_{s,i} = \frac{2 \cdot precision_{s,i} \cdot recall_{s,i}}{precision_{s,i} + recall_{s,i}} \quad (3)$$

For each participant, Correctness of understanding is computed as the average of the F-Measure values of all the questions in the comprehension questionnaire. This variable takes values in between 0 and 1. A value close to 1 means that the participant achieved a good source code comprehension because s/he answered rather well to all the questions in the comprehension questionnaire. Conversely, a value close to 0 means that source code comprehension was not good. We chose the measures above because they are widely used in the empirical software engineering field (e.g., [Ricca et al. 2010; Abrahão et al. 2013]).

Figure 3.3 shows a sample question. It expected as correct answer the following set of items: `CreationTool`, `ArrayFigure`, `StandardDrawingView`, `createFigure()`, `draw()`,

<b>Q2.</b> Indicate the class/es and the method/s in charge of creating, drawing, and updating the instances of the class Figure.				
How much do you trust your answer <sup>+</sup> ?				
<input type="checkbox"/> Unsure	<input type="checkbox"/> Not Sure Enough	<input type="checkbox"/> Sure Enough	<input type="checkbox"/> Sure	<input type="checkbox"/> Very Sure
How do you assess the question <sup>+</sup> ?				
<input type="checkbox"/> Very Difficult	<input type="checkbox"/> Difficult	<input type="checkbox"/> On Average	<input type="checkbox"/> Simple	<input type="checkbox"/> Very Simple
What is the source of information used to answer the question <sup>+</sup> ?				
<input type="checkbox"/> Previous Knowledge (PK)		<input type="checkbox"/> Internet (I)	<input type="checkbox"/> Source Code (SoC)	
<sup>+</sup> Mark only one answer				

Fig. 2. A sample question from comprehension questionnaire

Table III. Experiment Design

Experiment	GD	TD	SC	Total
E1-Prof	8	9	8	25
E2-UniSA	8	9	8	25
E3-UniBAS	8	8	7	23
E4-UniBZ	5	5	5	15

and `drawingRequestUpdate()`. These items can be derived by the instance of the Observer pattern shown in Figure 1 and by the instances of the Prototype and Composite patterns. The Observer instance is in charge of painting and/or repainting objects that are instances of Figure. The Prototype instance helps in understanding the object that would manage the creation of a template for Figure, while the Composite instance draws each base element of an object Figure. Details on the latter two instances can be found in the experimental material we made available on the web [Scanniello et al. 2014b]. If, for example, a participant provided `CreationTool`, `ArrayFigure`, and `drawingChangeListeners()` as his/her answer, the corresponding value for Correctness of understanding would be 0.44. This results from the precision and recall values, respectively being 0.66 and 0.33. In fact, the number of correct items provided is 2 (`CreationTool` and `ArrayFigure`), the total number of items provided is 3, and the expected number of correct items is 6.

As mentioned in Section 2.1, *Method* is the main factor (or manipulated factor) in our family of experiments. It is a nominal variable and assumes values in {SC, TD, GD}. SC indicates that the participant was provided with source code without documented design-pattern instances. On the other hand, TD and GD indicate that participants were provided with source code added with textually or graphically documented instances, respectively. The moderator variable Experience assumes values in {E1-Prof, E2-UniSA, E3-UniBAS, E4-UniBZ}.

### 3.4. Experiment Design

We used the one factor with three treatments design [Wohlin et al. 2012] shown in Table III. The use of a different experimental design (such as a within-participant counterbalanced design) with non-trivial experimental objects (as in our family of experiments) may bias results introducing a possible carry-over effect: if a participant is tested first under the experimental condition *A* and then under the experimental condition *B*, s/he could potentially exhibit better or worse performances under the condition *B*.

We used the participants' working experience<sup>4</sup> as the blocking factor for E1-Prof, while the participants' ability<sup>5</sup> was the blocking factor for both E2-UniSA and E3-UniBAS. As for E4-UniBZ, we used the final grade of the Master degree<sup>6</sup>. Due to the used experiment design, we use hereafter group and treatment interchangeably.

### 3.5. Hypotheses

We have defined and investigated two null hypotheses. Each of them is described as a parameterized hypothesis as follows:

- **Hn1 X.** With respect to the dependent variable X (i.e., Correctness of understanding or Time), there is no difference among source code with undocumented design pattern instances, source code with textually documented design pattern instances, and source code with graphically documented design pattern instances.
- **Hn2 X.** With respect to the dependent variable X and the participants' experience, there is no difference among source code with undocumented design pattern instances, source code with textually documented design pattern instances, and source code with graphically documented design pattern instances.

When null hypotheses are rejected, it is possible to accept the alternative ones that can be easily derived (e.g., Ha1 X: With respect to the dependent variable X, there is difference among source code with undocumented design pattern instances, source code with textually documented design pattern instances, and source code with graphically documented design pattern instances).

### 3.6. Experiment Tasks

We asked participants to perform the following tasks:

- **Comprehension task.** Independently from the treatment, the comprehension questionnaire was composed of 14 open questions (see Appendix A). These questions were divided into three groups to let participants take a break if needed when passing from a group of questions to the next one. This choice aimed to reduce fatigue effect biases. The questions were arranged into these groups on the basis of the recommendations by Sillito et al. [2008]. In particular, the questions in the first group focused on expanding points in the source code believed to be relevant and related to software comprehension, often by exploring relationships among entities (e.g., classes and method). The questions in the second group are concerned with understanding concepts in the source code that involved multiple relationships and software entities. Answering these questions required understanding the overall structure of a subgraph<sup>7</sup>. To answer the questions in the third group, the participant had to deal with over related groups of subgraphs. That is, these questions involved

<sup>4</sup>A professional was considered as junior developer when his/her working experience was less than or equal to 3 years, senior otherwise. The junior developers were 10, while 15 were the senior developers.

<sup>5</sup>In Italy, the exam pass grades are expressed as integers and take values in between 18 and 30. The lowest grade is 18, while the highest 30. Students with a Grade Point Average (GPA) less than or equal to 24 were considered low ability participants; otherwise high [Abrahão et al. 2013].

<sup>6</sup>A PhD student was considered to be low when his/her final grade was less than or equal to 8, high otherwise. All the final grades were normalized in between 1 and 10 because of differences in the organization of the universities where the PhD students took their Master grade. For example, final Master grades in Italy are expressed as integers and take values in between 66 and 110, so a student with a final grade less than or equal to 88 was considered to be a low ability participant.

<sup>7</sup>Sillito et al. [2008] see a code base as a graph of entities (classes, methods, and fields, for example) and relationships between those entities (references and calls, for example). To answer any specific question related to the comprehension of a software, the authors believe that it requires considering some subgraph of the code base.

Table IV. Post-experiment survey questionnaire

Id	Question	Answers
S1	I had enough time to perform the task	(1-5)
S2	The tasks I performed were perfectly clear to me	(1-5)
S3	The task objectives were perfectly clear to me	(1-5)
S4	Comments included in the source code were clear	(1-5)
S5	I found the arrangement of the comprehension questions in three parts to be useful	(1-5)
S6	The design-pattern instances were useful to answer the questions	(1-5)
S7	The design-pattern instances were well documented	(1-5)
S8	How much time (in percentage) you believe to have spent in reading source code?	(A-E)
S9	How much time (in percentage) you believe to have spent in analyzing design-pattern instances?	(A-E)
1 = Strongly agree, 2 = Agree, 3 = Neutral, 4 = Disagree, 5 = Strongly disagree		
A. < 20%; B. $\geq 20\%$ and < 40%; C. $\geq 40\%$ and < 60%; D. $\geq 60\%$ and < 80%; E. $\geq 80\%$		

understanding relationships between multiple subgraphs or the understanding of the interaction between a subgraph and the rest of the system. All the questions in the comprehension questionnaire were formulated using a similar form/schema. For each question, the participants in the groups GD and TD had to specify whether the answer was deduced from: (DPI) Design-Pattern Instances, (PK) Previous knowledge, (I) Internet, or (SoC) Source Code. If participants specified DPI, they were also asked to indicate the used instance/s. Participants in the group SC could choose among: PK, I, or SoC (see Figure 3.3). This was the only difference introduced in the comprehension questionnaires used for SC, GD, and TD. As for confidence level, participants had to chose among the following possibilities: Unsure, Not sure enough, Sure enough, Sure, and Very Sure. Furthermore, perceived difficulty level admitted values were {Very Difficult, Difficult, On Average, Simple, and Very Simple}.

- **Post-experiment task.** We asked participants to fill in the post-experiment survey questionnaire shown in Table IV. The goal of this questionnaire was to obtain feedback about participants' perceptions of the experiment execution. Answers to questions S1-S7 were based on a five-point Likert scale [Oppenheim 1992]: from strongly agree (1) to strongly disagree (5). Questions S8 and S9 demanded answers according to a different five-point Likert scale based on the percentage of time used to deal with source code and documented pattern instances (see last row of Table IV). The participants in the GD and TD groups had to answer all the questions in Table IV, while other participants answered just questions in between S1 and S5.

### 3.7. Experiment Procedure

The experiment procedure adopted for the replications was different from that in the original experiment, where at most two professionals were involved in each experimental session. In fact, it is practically impossible to conduct a single experimental session gathering together professionals from different companies. Before the controlled experiment, all the professionals had to fill in a pre-questionnaire. This questionnaire was sent and returned by email. The gathered information was used to classify participants as junior or senior developers. Just before an experiment session, experiment

supervisors (two of the authors, namely the first and the third) highlighted the study goal without providing details on its hypotheses. Experiment sessions were carried out under controlled conditions to avoid biasing results. Experiment supervisors were the same at each session. This allowed us to reduce possible observation biases. No training session on tasks similar to the ones in the experiment was carried out. The main reasons were: (i) they should have an adequate experience in performing maintenance operations on source code written by others and (ii) time and logistic constraints made the execution of a training session impossible (the use of professionals might cause this kind of concern).

As for the replications, we conducted them in research laboratories under controlled conditions. A few days before each experiment, participants attended a training session where supervisors presented detailed instructions on the experiment and recalled basic notions on design-pattern based development. All the participants in a single experiment attended this session together.

We monitored participants to prevent communication with each other. This was also done in E1-Prof, when more professionals participated in one session. At the end of each experiment, we asked participants to return material. No time limit to perform the experiment was imposed. This implies that quality and time are considered separately when analyzing results [Bergersen et al. 2011; Bergersen et al. 2014]. As suggested by Bergersen et al. [2011] this is acceptable when the goal of a study is not to characterize individual differences among participants [Bergersen et al. 2011].

The participants used a general purpose and well known text editor, namely UltraEdit ([www.ultraedit.com](http://www.ultraedit.com)). We opted for this editor because it was used in basic programming language courses at both the University of Salerno and University of Basilicata. Therefore, the prior knowledge of the participants in E2-UniSA and E3-UniBAS should be rather homogeneous on this tool. Finally, we did not allow the participants to choose their preferred IDE (Integrated Development Environment) because familiarity with the tool might skew results in an undesirable way.

We allowed all the participants to use the Internet to accomplish the experiment session since developers usually exploit this medium as support for their daily work activities. We tried to reproduce actual working conditions.

We asked all the participants to use the following experimental procedure: (i) writing down their name and start-time; (ii) answering the questions of the comprehension questionnaire without executing source code; and (iii) writing down the end-time. The experiment procedure was the same for each group of questions within the comprehension questionnaire. We did not suggest any approach to deal with source code. To avoid wasting time, we only discouraged to read all the source code.

We provided all the participants with a printed copy of the following experiment material: (i) the comprehension questionnaire and (ii) a post-experiment survey questionnaire. The participants in the group GD were provided with the source code (without any reference to pattern instances) and a printed copy of the document where each pattern instance was graphically reported (see Figure 1(a)). The participants in the group TD were provided with source code that included references to the instances in the comments (see Figure 1(b)). As for the group SC, the participants were provided with source code without any kind of documentation.

### 3.8. Analysis Procedure

To perform data analysis, we used the R environment ([www.r-project.org](http://www.r-project.org)) for statistical computing and we carried out the following steps:

- (1) We undertook the descriptive statistics of the dependent variables.

- (2) To test null hypotheses, we adopted non-parametric tests due to the sample size and mostly the non-normality of the data. As for Hn1 X, we used the Kruskal-Wallis test. This statistical test is a non-parametric alternative to ANOVA in case of one factor with more than two treatments [Wohlin et al. 2012]. We used the Kruskal-Wallis test because of the design of the experiments (only unpaired analyses were possible) in the family and for its robustness [Wohlin et al. 2012]. As for Hn2 X, we opted for a two-way permutation test [Baker 1995], a non-parametric alternative to the two-way ANOVA test. In all the performed statistical tests, we decided (as it is customary) to accept a probability of 5% of committing Type-I-error [Wohlin et al. 2012].
- (3) The chosen statistical tests allow the presence of a significant difference between independent groups to be verified, but they do not provide any information about the magnitude of such a difference. To measure this difference (if present), we used the Cliff's  $d$  non-parametric effect size measure [Kampenes et al. 2007]. Effect size is considered negligible for  $|d| < 0.147$ , small for  $0.147 \leq |d| < 0.33$ , medium for  $0.33 \leq |d| < 0.474$ , and large for  $|d| \geq 0.474$  [Kampenes et al. 2007]. We also computed the average percentage improvement for the defined dependent variables (i.e., Time and Correctness of understanding). For example, given the variable Time and the methods SC and TD, we computed  $\frac{MeanTime(GD) - MeanTime(SC)}{MeanTime(SC)} \%$ . If the obtained value is negative, participants spent on average less time with GD. Indeed, the lower the value, the lower the effort to comprehend source code with GD is. As for Correctness of understanding, it estimates the average percentage improvement in source code comprehension. The higher the value, the better source code comprehension with GD is.
- (4) We also analyzed the statistical power of each performed test. Statistical power is defined as the probability that a statistical test will correctly reject a null hypothesis [Dybå et al. 2006]. Knowledge of statistical power can influence the planning, the execution and the results of an empirical study (e.g., [Arisholm et al. 2007; Cohen 1988]). Statistical power can be also determined post-hoc [Dybå et al. 2006; Scanniello et al. 2014a]. The highest value is 1, while 0 the lowest. Therefore, the higher the statistical power value, the higher the probability to reject a null hypothesis when it is actually false. The value 0.80 is considered as a standard for adequacy [Ellis 2010]. The statistical power is computed as 1 minus the Type-II-Error (i.e.,  $\beta$ -value). This type of error indicates that a null hypothesis is false, but the statistical test erroneously fails to reject it. In the discussion of the results, we used the  $\beta$ -value when a statistical test was unable to reject a null hypothesis. The higher the  $\beta$ -value, the lower the probability of erroneously not rejecting a null hypothesis is.
- (5) We used mosaic plots to graphically summarize the results for the confidence and perceived difficulties as well as the source of information the participants exploited to answer questions in the comprehension questionnaire.
- (6) To summarize raw data and to support their discussion we also used boxplots.

### 3.9. Differences and Issues

*3.9.1. Differences among Experiments.* The choice and the numbers of factors to be varied among experiments is relevant because a large number of variations could make it less likely to trace observed results to the factor of interest (manipulated factor) and thus the result validity and interpretation could be strictly affected [Shull et al. 2008; Scanniello et al. 2014a]. In the following, we summarize variations among the four experiments in our family:

- **Participants.** Variations in the participants' experience and in environmental factors can contribute some confidence that the effect of the factor under study is not limited to one particular setting or experimental environment.
- **Training session.** Professionals did not carry out a training session on tasks similar to those used in the experiment.
- **Experiment material.** Sources of possible confusion in materials were removed after each experiment. As for E4-UniBZ, the experimental material was in English. Independently from the treatment, the participants in E4-UniBZ were asked to answer two additional questions to understand the *common sense* in the representation and use of design-patterns. In particular, the following two open questions were asked:
  - (i) Based on your experience and knowledge, please explain why (or why not) the documentation of design-pattern instances should improve source code comprehensibility and maintainability;
  - (ii) Please suggest how design-pattern instances should be documented to improve source code comprehensibility and maintainability.
 We added this difference in the replication E4-UniBZ because we believe that more significant comments might be given from the participants with an experience and maturity level higher than those of the participants in E2-UniSA and E3-UniBAS.

*3.9.2. Documentation and Communication Issues.* The success or the failure of replications might be influenced by issues such as documentation [Shull et al. 2004] and communication among experimenters [Vegas et al. 2006]. To deal with these issues, we used laboratory packages, knowledge sharing mechanisms, and communication media. In particular, with regard to documentation, original experimenters translated all the material (not initially written in English) from Italian into English. The investigator of the external replication (the fifth author) was provided with all the experiment material and asked clarifications and implemented modifications if needed. In addition, the experimenters shared: (i) the papers where the pilot study and original experiment were presented [Gravino et al. 2011; Gravino et al. 2012] and (ii) a document to provide a common background.

We used instant messaging tools, mobile phones, teleconference tools, and email to establish a communication channel among experimenters. To reduce consistency issues across the experimenters, the results from the most important interactions were reported in a shared document where experiment decisions were recorded.

## 4. RESULTS

In this section, we present the obtained results following our analysis procedure.

### 4.1. Descriptive Statistics

Table V shows the values of mean and standard deviation for Time and Correctness of understanding grouped by Method. The distribution of the values for these variables grouped by Method and Experiment is graphically shown by the boxplots in Figure 3.

We can observe that in all the experiments (except for E2-UniSA) the participants belonging to the group/treatment SC spent on average more time than the participants in the groups GD and TD. A small difference can be observed between the groups GD and TD in all the experiments with respect to the mean time to accomplish the task. Boxplots in Figure 3(a) show that the same trend is present among the experiments for each group (e.g., participants in E1-Prof spent more time than those in E4-UniBZ).

As for Correctness of understanding, the participants in the groups GD and TD achieved on average better scores. This does not hold in E4-UniBZ for the group TD. On the other hand, a small difference is observable in favor of the group TD for E1-



Table V. Descriptive statistics for GD, TD, and SC

Experiment	Dependent variable	GD		TD		SC	
		Mean	St. Dev.	Mean	St.Dev.	Mean	St.Dev.
E1-Prof	Time	142.2	31.17	136.4	37.00	147.4	42.81
	Correctness of understanding	0.51	0.10	0.54	0.07	0.40	0.10
E2-UniSA	Time	85.5	24.089	93.44	22.07	83.5	33.98
	Correctness of understanding	0.406	0.116	0.43	0.1	0.31	0.12
E3-UniBAS	Time	112	30.836	125.125	21.397	148.857	40.268
	Correctness of understanding	0.379	0.073	0.349	0.131	0.384	0.111
E4-UniBZ	Time	94.6	13.686	77.8	20.777	108.4	28.5
	Correctness of understanding	0.511	0.034	0.602	0.177	0.388	0.124

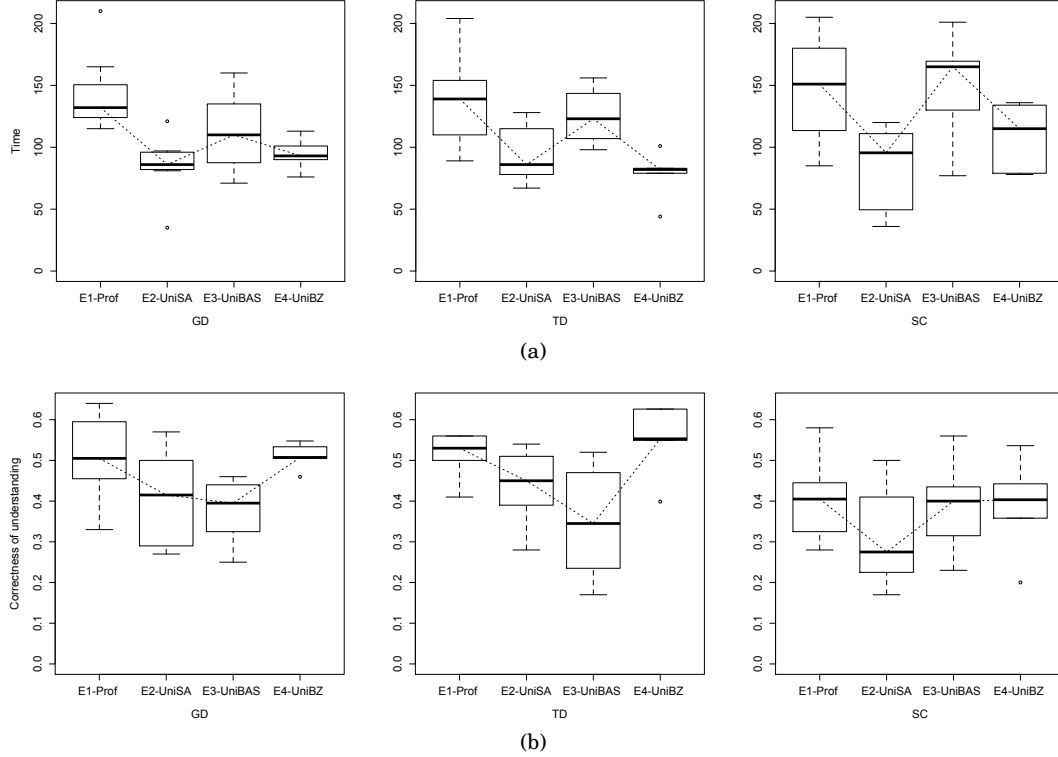


Fig. 3. Distributions for Time (a) and Correctness of understanding (b) grouped by Method and experiment

Prof. Boxplots in Figure 3(b) show that more experienced participants benefit from documented pattern instances more than low experienced participants. As for what concerns the treatment SC, there is a non-noteworthy difference among participants in all experiments.

#### 4.2. Hypotheses Testing

4.2.1.  $H_{n1.X}$ : SC vs. GD vs. TD. The results for the null hypothesis  $H_{n1.X}$  are summarized in Table VI. The Kruskal-Wallis<sup>8</sup> test indicates that this null hypothesis can be

<sup>8</sup>To evaluate if data were not normally distributed, we applied the Shapiro-Wilk W test [Shapiro and Wilk 1965]. If this test returns a p-value smaller than the chosen threshold (0.05 in our case), data are considered to be non normally distributed. The Shapiro-Wilk W test returned 0.02 as p-value for the group GD in E1-Prof. Although p-values larger than 0.05 were obtained in all other cases, we performed non-parametric analyses for uniformity reasons and to be as much conservative as possible.

Table VI. Results for Hn1.X (SC vs. GD vs. TD)

Experiment	Dependent variable	p-value
E1-Prof	Time	No (0.791)
	Correctness of understanding	<b>Yes</b> (0.035)
E2-UniSA	Time	No (0.943)
	Correctness of understanding	No (0.091)
E3-UniBAS	Time	No (0.097)
	Correctness of understanding	No (0.836)
E4-UniBZ	Time	No (0.343)
	Correctness of understanding	<b>Yes</b> (0.046)

rejected on Correctness of understanding for E1-Prof and E4-UniBZ. The p-values are 0.035 and 0.046, respectively.

When three different groups are compared, a significance test only shows whether there is a significant difference in general [Wohlin et al. 2012]. A post-hoc analysis is needed afterwards, to evaluate between which conditions the significant effect occurs. Therefore, we performed a post-hoc analysis on Correctness of understanding for the experiments E1-Prof and E4-UniBZ, namely those experiments where a significant difference was observed. For each of these two experiments, we tested the following null hypotheses: (i)  $H_{nSCvs.D}$  - Source code comprehensibility does not increase when using D (i.e., TD or GD) with respect to SC; and (ii)  $H_{nTDvs.GD}$  - Source code comprehensibility does not increase when using GD with respect to TD. These hypotheses are one-sided on the basis of the results of the descriptive statistics.  $H_{nSCvs.D}$  is a parametrized null hypothesis.

The above hypotheses have been verified by the Mann-Whitney test [Conover 1998] (also known as the Wilcoxon rank-sum test). The results of the performed post-hoc analysis are summarized in Table VII and Table VIII.

The Mann-Whitney test indicated the presence of statistically significant differences in E1-Prof and E4-UniBZ for both  $H_{nSCvs.TD}$  and  $H_{nSCvs.GD}$  (see Table VII). The obtained p-values range in between 0.008 (SC vs. TD in E1-Prof) and 0.047 (SC vs. GD in E4-UniBZ). In both experiments, the effect size is large and the values range in between 0.562 and 0.76 and the average improvement values are very high (from 87% to 116.2%). These results indicate a better source code comprehension for TD and GD than for SC, in practical terms. The statistical power values are not always adequate. This could be due to the number of participants in the groups.

Table VII. Post-hoc analysis on Correctness of understanding ( $H_{nSCvs.TD}$  and  $H_{nSCvs.GD}$  in E1-Prof and E4-UniBZ)

Experiment	Hypothesis	$p - value$	Cliff's $d$	Statistical Power	Average Improvement
E1-Prof	$H_{nSCvs.TD}$	<b>Yes</b> (0.008)	-0.708	0.903	93.3%
	$H_{nSCvs.GD}$	<b>Yes</b> (0.033)	-0.562	0.607	87%
E4-UniBZ	$H_{nSCvs.TD}$	<b>Yes</b> (0.03)	-0.76	0.6	116.2%
	$H_{nSCvs.GD}$	<b>Yes</b> (0.047)	-0.68	0.528	92.8%

Table VIII. Post-hoc analysis on Correctness of understanding ( $H_{nTDvs.GD}$ ) in E1-Prof and E4-UniBZ)

Experiment	$p - value$	Cliff's $d$	$\beta$ -value
E1-Prof	No (0.314)	-0.153	0.865
E4-UniBZ	No (0.072)	-0.6	0.782

Table IX. Results for Hn2.X

Dependent variable	Method	Experience	Method vs. Experience
	<i>p - value</i>		
Time	No (0.33)	<b>Yes</b> (<0.001)	No (0.495)
Correctness of understanding	<b>Yes</b> (<0.001)	<b>Yes</b> (<0.001)	No (0.157)

As for  $Hn_{TDvs.GD}$ , the Mann-Whitney test indicated that this null hypothesis cannot be rejected (see Table VIII). Therefore, the difference between the participants who used graphically documented and textually documented pattern instances is not statistically significant. The  $\beta$ -values are close to the standard for adequacy (0.865 and 0.782). The effect size is either small (in E1-Prof) or large (in E4-UniBZ).

These results suggest that the documentation of pattern instances aids in correctness of understanding for object-oriented source code in E1-Prof and E4-UniBZ. In addition, we observed that there is a non-noteworthy difference in comprehending source code when pattern instances are documented either graphically or textually.

**4.2.2. Hn2.X: Participants' Experience Effect.** Table IX summarizes the results from the two-way permutation test. The effect of participants' experience is statistically significant on Time and no interaction between Method and Experience is present. As shown in Figure 3(a), professionals spent more time to accomplish the comprehension task when pattern instances are documented. There is a non-noteworthy difference in the other experiments for the group GD. As for the group TD, the participants in E3-UniBAS spent more time than the participants in the other two replications. The results from the two-way permutation test also indicate a positive effect of Method and Experience on Correctness of understanding. The p-values are less than 0.001, respectively. More experienced participants benefited more than less experienced ones when pattern instances are documented (see Figure 3(b)). No interaction between Method and Experience is shown on Correctness of understanding.

The data analysis results suggest that the participants experience has a statistically significant effect on source code comprehensibility with respect to both the variables Time and Correctness of understanding, while there is no significant interaction effect.

#### 4.3. Confidence

Figure 4 shows the mosaic plots for the confidence that the participants displayed regarding their own understanding on the source code in the comprehension task. The data are grouped by Experience and Method. The participants generally indicated "sure enough", "sure", and "very sure" as their own confidence level of source code understanding. This holds in all the experiments and for all the methods (i.e., GD, SC, and TD). Indeed, the participants in E3-UniBAS expressed the worst level of confidence in understanding. A possible reason could be related to the participants' knowledge on high- and low-level design of object-oriented software systems and software development and maintenance.

#### 4.4. Perceived Difficulty

The subjective judgment on the ease in obtaining information through the documentation of pattern instances (if present) and/or source code is visually summarized by the mosaic plots in Figure 5. Independently from the method, the most frequent answer given in all the experiments of our family is "On Average" (i.e., 3). Although there is a non-noteworthy difference among the experiments in our family, it seems that low experienced participants found it less easy to obtain information from the documentation provided to accomplish comprehension tasks.

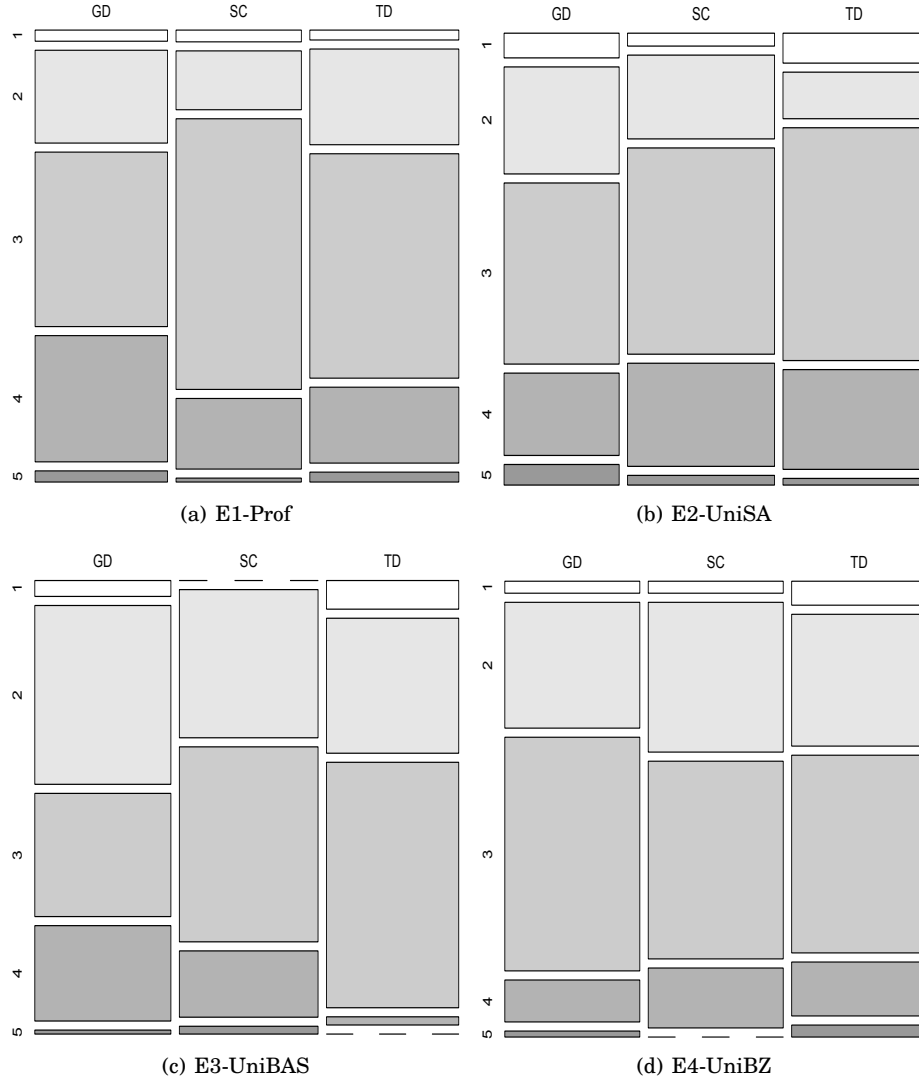


Fig. 4. Mosaic plot for Confidence (1 - Unsure; 2 - Not Sure Enough; 3 - Sure Enough; 4 - Sure; 5 - Very Sure)

#### 4.5. Source of Information

Figure 6 shows the mosaic plots regarding the source of information for each experiment. The participants indicated source code as the main source of information to answer questions in the comprehension questionnaire. This finding held in all experiments except for the group GD in E1-Prof, where the participants indicated DPI as the main source of information. As for E3-UniBAS and E4-UniBZ, the difference between DPI and SoC is less evident for the group TD. The participants in these two experiments indicated that they did not use the Internet to answer the questions of the comprehension questionnaire.

To accomplish a source code comprehension task, the participants perceived documented instances as a relevant, but not predominant source of information. This is




Fig. 5. Mosaic plot for Perceived Difficulty (1 - Very Difficult; 2 - Difficult; 3 - On Average; 4 - Simple; 5 - Very Simple)

a case where controlled experiments provide insight into the difference between perceived usefulness of a given method and the actual advantage of using it.

#### 4.6. Post-experiment Survey Questionnaire

The time needed to carry out experiments was considered appropriate (question S1). The median value was 1 (strongly agree) for all experiments except for E2-UniSA, where the median was 3 (neutral). The objectives (S2) were considered to be clear in all experiments except in E2-UniSA: the median values were 2 (agree) and 3 (neutral), respectively. The participants found comprehension tasks to be clear. The median value was 2 (agree) in all experiments. The answers to S4 suggest that the participants found comments to be clear. In particular, the median value was 1 (strongly agree) in E3-UniBAS, while it was 2 (agree) in all remaining experiments. All the participants

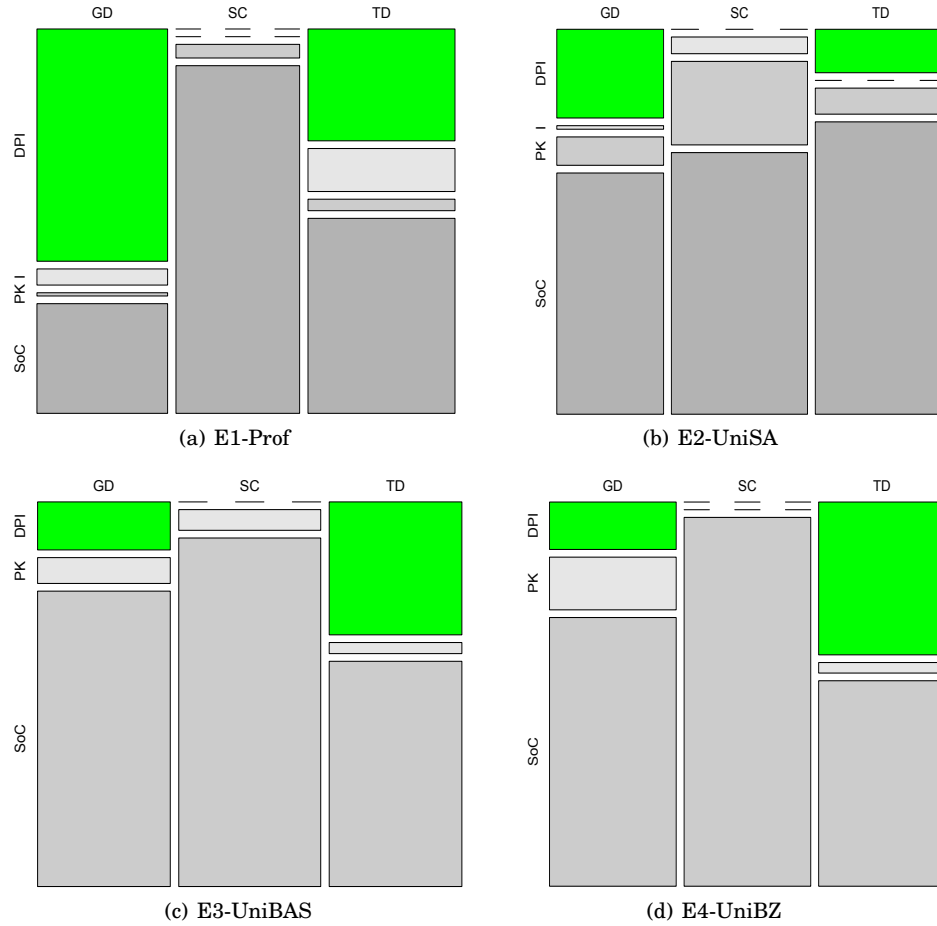


Fig. 6. Mosaic plot for source of information - (DPI) Design-Pattern Instances, (PK) Previous knowledge, (I) Internet, or (SoC) Source Code.

agreed on the usefulness of the arrangement of the comprehension questionnaire in three parts (S5). In all experiments, the median value was 2 (agree).

As for E2-UniSA, the participants that carried out the experiment in the groups GD or TD expressed a neutral judgment (the median value) on the usefulness of pattern instances to accomplish the comprehension task (S6). The participants in other experiments agreed on the usefulness of pattern instances. In these experiments, the median value was 2. The participants found pattern instances to be well documented. The median value for S7 was 2 (agree) in all experiments.

Regarding S8, the median value was equal to D in all experiments. The participants indicated that the time they spent to analyze source code ranges from 60% to 80% of the total time. The participants in E2-UniSa specified that the time to consult pattern instances was less than 20% (S9) of the total time to accomplish the comprehension task. All the other participants specified a range between 20% and 40%.

Table X. Correctness of understanding grouped by CI and NCI

Experiment	Method	CI		NCI		p-value
		Mean	St.Dev.	Mean	St.Dev.	
E1-Prof	GD	0.671	0.309	0.499	0.359	<b>Yes</b> (0.004)
	TD	0.73	0.277	0.337	0.333	<b>Yes</b> (< 0.001)
E2-UniSA	GD	0.721	0.252	0.392	0.339	<b>Yes</b> (< 0.001)
	TD	0.694	0.342	0.34	0.358	<b>Yes</b> (< 0.001)
E3-UniBAS	GD	0.602	0.328	0.268	0.315	<b>Yes</b> (< 0.001)
	TD	0.626	0.265	0.35	0.363	<b>Yes</b> (0.006)
E4-UniBZ	GD	0.644	0.28	0.584	0.376	No (0.39)
	TD	0.725	0.298	0.422	0.324	<b>Yes</b> (< 0.001)

#### 4.7. Further Analyses

In the following, with respect to Correctness of understanding, we compare the results of the participants who correctly identified the instances needed to answer the questions in the comprehension questionnaire and those who did not. We conclude with the feedback that the participants in E4-UniBZ gave on how design pattern instances should be documented to improve source code comprehensibility and maintainability.

*4.7.1. Correctly vs. Incorrectly Identified Instances.* The answers to the questions in the comprehension questionnaire have been classified as: CI (Correctly Identified pattern instances) and NCI (Non Correctly Identified pattern instances). To understand whether the identification of design-pattern instances affects Correctness of understanding, we compute the mean and standard deviation values on these two groups of answers. These statistics are summarized in Table X.

The descriptive statistics suggest that Correctness of understanding improves when pattern instances are correctly recognized. This result held in all experiments with instances documented both graphically and textually. In addition, there is a non-noteworthy difference between the groups GD and TD when the instances are either correctly or incorrectly identified.

The results of a Mann-Whitney test for unpaired analyses indicate that the *Correctness of understanding is significantly greater when design-pattern instances are correctly identified with respect to when these instances are not correctly identified.* This is true in all the experiments with the only exception of E4-UNiBZ on GD, where the statistical test returned a p-value equal to 0.39. It is also worth mentioning that the statistical power is always high, when the Mann-Whitney test returned a p-value less than 0.05. The statistical power assumes values in between 0.87 and 1. We also computed the effect size for GD and TD when pattern instances were correctly or not correctly recognized by participants. In particular, the effect size is small in E1-Prof on GD (Cliff's  $d = 0.277$ ), while it is medium in E3-UniBAS on TD. In all the other experiments (except E4-UNiBZ on GD), the effect size is large (values range between 0.432 and 0.623).

*4.7.2. A Qualitative Look Inside the Use of Design-Pattern Instances.* In this section, we discuss answers that the participants in E4-UniBZ gave to the two additional open questions presented in Section 3.9.1. The following are the main outcomes:

- The participants provided slightly different interpretations of why the documentation of pattern instances should make source code comprehension and maintenance easier. However, there was consensus that instance documentation made software implementation more understandable (e.g., the link among classes can be found quickly) and that it implicitly clarifies the rationale behind that implementation. The participants did not find instance documentation to be useful as a communication mechanism among developers.

- 10 out of 15 participants stated that pattern instances should be graphically represented with UML class diagrams to improve source code comprehensibility and maintainability. This answer was given by 4 participants in the group SC and by 4 participants in the group TD, while it was given by 2 participants in the group GD. A possible justification is that the participants were familiar with the method of representing instances proposed by Gamma et al. [1995] and thus when instances were not documented in accordance with this method they were less comfortable with comprehension tasks. When instances were documented in accordance with the method by Gamma et al. [1995], the participants believed that a different kind of documentation could better fit comprehensibility. We can conjecture that this finding might be due to how participants learned the basics of design-pattern development and not to the true benefits deriving from the use of one kind of documentation rather than another. This point might represent a future direction for our research.

#### 4.8. Threats to Validity

To comprehend strengths and limitations of our family of experiments, threats that could affect results and their generalization are presented and discussed in this subsection. Despite our effort in mitigating as many threats as possible, some of them are unavoidable. We discuss the threats to validity using the guidelines proposed by Wohlin et al. [2012].

*4.8.1. Internal Validity.* In empirical studies internal validity threats are relevant. They are even more relevant in a study like ours that try to establish a causal relationship, because threats to internal validity may invalidate such a relationship.

- *Interaction with selection.* It has been mitigated by the experiment design. Each group of participants worked on one task only, with or without documented design-patterns instances. Furthermore, the participants in each experiment had similar background and experience.
- *Maturation.* The adopted experimental design should mitigate possible learning effect since participants worked on one task only. We have mitigated the fatigue effect allowing participants to take a break when passing from a group of questions in the comprehension questionnaire to the next one.
- *Diffusion or imitation of treatments.* This threat concerns the information exchanged among participants within each experiment. We prevented this in several ways. While accomplishing each experiment, the participants were monitored by the experiment supervisors to prevent that they communicated with one another. The supervisors also ensured that no communication among participants occurred during each break. This point is less serious in E1-Prof, where at most two professionals were present in each laboratory session. The use of the Internet to exchange information might represent another bias. Participants could also use the Internet to get additional information on the experimental object. For example, the participants in the group TD could have fetched from the web the graphical representations of the pattern instances in the experimental object. To deal with threats related to the use of the Internet, we monitored participants, while performing experimental tasks. As an additional measure to prevent the diffusion of material among participants, we asked participants to return material when each experiment was accomplished. This allowed reducing biases related to the communication among participants in different experiments.

*4.8.2. External Validity.* External validity primarily refers to the generalizability of the obtained results. Questions about external validity may arise in controlled experi-



ments. However, this kind of empirical strategy is often conducted in the early steps of a longterm empirical investigation to reduce failure risks [Arisholm et al. 2006].

- *Interaction of selection and treatment.* The use of students as participants may affect external validity [Carver et al. 2003; Ciolkowski et al. 2004; Hannay and Jørgensen 2008]. To deal with this threat, we conducted replications with professionals and PhD students. Regarding professionals, they are Italian junior/senior software developers. The majority of the professionals (15 out of 25) worked for companies located in Southern Italy. The remaining participants worked for companies located in Central (8) and Northern Italy (2). Southern Italy and Central Italy are over-represented with respect to Northern Italy. This point might represent a threat to external validity. Regarding the size of the sample (i.e., number of professionals), it was hard to find a larger number of participants because of their limited availability and restrictions from their employers. PhD students were enrolled in an international PhD course in Computer Science. We were not able to find a larger number of PhD students because of their limited availability. As for the participants in E2-UniSA and E3-UniBAS, they were trained on UML, design-pattern development, and took several programming courses (e.g., procedural programming and object oriented programming). In addition, some of the participants in E2-UniSA (i.e., about 40%) had an internship in industry as a final part of their Bachelor degree. However, working with students also has various advantages. For example, the students prior knowledge is rather homogeneous [Verelst 2004]. To increase our confidence in the results, we are also promoting further independent/external replications making a laboratory package available on the web [Scanniello et al. 2014b]. This will allow researchers to easily conduct replications of our experiments [Lindvall et al. 2005]. External differentiated replications will further increase the confidence in our results. These kinds of replication will reduce the likelihood of the researcher bias possibly present in our experimental procedure [Sjøberg et al. 2005].
- *Interaction of setting and treatment.* In our study, the size and complexity of the experimental object (i.e., documentation and system itself) may affect result validity. We selected a portion of an open-source software system large enough to be considered not excessively easy. Larger experimental objects could overload participants, thus biasing the experimental results. In addition, the source code of the experimental object was not familiar to participants, thus allowing us to simulate a realistic comprehension task. Another possible threat to external validity related to the experimental object is that JHotDraw is well designed and design-patterns are mostly intentionally used. This might make the experimental object not representative. Future work is needed to verify if our results can be generalized in case of poorly designed software systems developed without the intentional use of design-patterns. The use of UltraEdit might represent another threat to the validity of the results as well as the translation of the experimental material. However, it is worth mentioning that these two possible threats might affect the results for all the three treatments (SC, GD, and TD). However, it is still possible that some of the participants were more familiar with UltraEdit than others, especially in E1-Prof and E4-UniBZ. Identifier renaming might also bias the results. To deal with this concern, source code modifications were performed by one of the authors (the first), who was not involved in the definition of the comprehension questionnaire.

**4.8.3. Construct Validity.** Construct validity concerns the link between the concrete experimental elements and the abstract concepts of the experimental goal. Some construct validity threats are also related to the experiment design and to social factors.

- *Interaction of different treatments.* To mitigate this threat we adopted a one factor with three treatments design [Wohlin et al. 2012]. Each participant experimented only one method on the same experimental object.
- *Confounding constructs and level of construct.* The procedure used to divide the participants into the groups in Table III could affect the validity of the obtained results. We are aware that the use of a different criterion to assess participants' working experience and ability could lead to different results. The fact that the participants in our family of experiments received the documentation in their native language (i.e., the participants in E1-Prof, E2-UniSA, and E3-UniBAS), while others did not (i.e., the participants in E4-UniBZ) might also affect the validity of results. Another threat is related to the fact that some of the participants received a training session on design-pattern development, while others did not. Professionals did not receive a training for time constraints.
- *Evaluation apprehension.* We mitigated this threat because the participants were not evaluated on the results achieved in the experiments. All students in E2-UniSA and E3-UniBAS were equally rewarded with one extra point in the exam grade, regardless of their actual performance. All the participants in our family of experiments were unaware of the objectives of our family of experiments and of our experimental hypotheses.
- *Experimenters' expectancies.* Questions in the comprehension questionnaire were formulated in order to favor none among groups GD, TD or SC. The reliability of the comprehension questionnaires was tested by applying the Cronbach's  $\alpha$  test [Cronbach 1951]. In Table XI, we summarize the results obtained for each set of the closed-questions intended to measure the three subjective variables (i.e., Source of Information, Perceived Difficulty, and Confidence). All the obtained values were higher than the acceptable minimum threshold [Maxwell 2002] ( $\alpha=0.70$ ). The post-experiment survey questionnaire was designed to capture the participants' perception of the tasks. We designed this questionnaire using standard methods and scales [Oppenheim 1992].

Table XI. Results of the Cronbach's  $\alpha$  test for the questionnaires reliability

Variable	E1-Prof	E2-UniSA	E3-UniBAS	E4-UniBZ
Source of Information	Acceptable (0.85)	Acceptable (0.95)	Acceptable (0.84)	—
Perceived Difficulty	Acceptable (0.86)	Acceptable (0.86)	Acceptable (0.9)	Acceptable (0.84)
Confidence	Acceptable (0.82)	Acceptable (0.79)	Acceptable (0.89)	Acceptable (0.83)

"—" signifies that the Cronbach's  $\alpha$  test has not been computed because all the participants indicated source code (see also Figure 6(d)).

- *Inadequate preoperational explication of construct.* The constructs are not sufficiently defined, before they are translated into measures or treatments. To deal with this kind of threat to construct validity, we took advantage of the framework we introduced in Section 2.

**4.8.4. Conclusion Validity.** Conclusion validity concerns issues that may affect the ability of drawing a correct conclusion.

- *Reliability of measures.* This threat is related to how the dependent variables were measured. The method to assess correctness of understanding is widely used in experiments with purposes similar to ours (e.g., [Ricca et al. 2010]). In addition, one of the authors (i.e., the second) not involved in the task definition built the comprehension questionnaire and analyzed the participants' answers to this questionnaire. Regarding the second considered dependent variable, we asked the participants to write their start and stop times on their questionnaire. We qualitatively validated the provided information. Even if this method of measuring the task completion

- time can be considered questionable [Hochstein et al. 2005], this practice is very common [Huang and Holcombe 2009; Ceccato et al. 2014; Scanniello et al. 2014a].
- *Random heterogeneity of participants.* We drew fair samples and conducted our experiments with participants belonging to these samples. Regarding practitioners, their experience could be heterogeneous.
  - *Fishing and the error rate.* Our hypotheses have been rejected considering proper p-values.
  - *Statistical tests.* We opted for non-parametric tests (e.g., Mann-Whitney test for unpaired analyses) to statistically reject the defined null hypotheses. We used non-parametric test also when parametric test could be applied. This choice might increase the risk of Type II errors.

## 5. DISCUSSION

We first discuss the results of each experiment individually and then together. The section is concluded with some practical implications from our family of experiments.

### 5.1. Individual Experiments

We observed a clear improvement in Correctness of understanding deriving from the use of documented instances. This finding holds for all the experiments in the family with the only exception of E3-UniBAS, namely the experiment having participants with the lowest level of experience. For E1-Prof and E4-UniBZ, the difference in Correctness of understanding is statistically significant. The participants in these experiments were those with the highest level of experience. On the basis of the observed results, we can conclude that: if the maintainer has an adequate level of experience with design-pattern development and computer programming, source code comprehensibility is statistically greater when instances are textually or graphically documented with respect to not having these instances documented at all. As for E2-UniSA, the average improvement achieved by the participants with graphically documented instances with respect to those having instances not documented at all was 30.1%. The average improvement in comprehension was 38.7% when instances were textually documented. In other words, source code comprehensibility is greater, but not statistically significant, when design-pattern instances are textually or graphically documented with respect to not having them documented at all. This implies that a maintainer benefits from the documentation of design-pattern instances also in case he/she has a Bachelor degree and knows the basics of high- and low-level design of object-oriented software systems, software development, and software maintenance.

The presence of documented instances did not significantly affect the time to accomplish comprehension tasks. Indeed, the participants spent slightly less time to accomplish that task when the implemented instances were documented either textually or graphically. This result might be considered unexpected because more documents/information to read and interpret should require more time to execute a task. This is even more evident for graphically documented instances. A justification for this result is that explicitly reported instances allow finding implementation details quicker and this reduces the time to comprehend source code.

In all the experiments, we observed that the average values for Correctness of understanding were higher when participants were able to correctly identify pattern instances needed to answer the comprehension questionnaire. In many cases (7 out of 8), we observed that the Correctness of understanding is significantly greater when design-pattern instances are correctly identified (see Table X). This finding suggests that it is important to help the maintainer to easily recognize pattern instances in source code because this would improve source code comprehensibility.

As for E1-Prof, we observed that the participants in the group GD indicated pattern instances as the first source of information. In contrast, the participants in the group TD indicated source code as the first source of information. Independently from the group, source code is the first source of information in all the other experiments of our family. This slight difference in the results achieved by the participants in the groups GD and TD could be due to the fact that in the latter case, pattern instances are documented in source code. Then, some participants could have considered documented pattern instances as an integral part of source code.

## 5.2. Participants' Experience

Participants with the greater experience had a more correct understanding of source code. In addition, more experienced participants benefit more than less experienced ones when instances are documented either textually or graphically. As for Time, we observed no noteworthy data pattern.

We observed that the participants who used textually documented instances and with a given experience (i.e., at least a Bachelor degree) obtained on average better results with respect to those participants who used graphically documented instances (see Table V). The effect size is also in favor of TD. It is negligible in E2-UniSA, small in E1-Prof, and large in E4-UniBZ. In case of the highest experienced participants (E1-Prof and E4-UniBZ), we also found that the average time to accomplish comprehension tasks is lower for TD than for SC and GD. These outcomes suggest that the participants that accomplished comprehension tasks with TD achieved an improved comprehension of source code with a lower effort. This is even more interesting if we compare TD with SC: the participants analyzed more information in less time and obtained an improved correctness of understanding of the source code. As for E2-UniSA, the participants that accomplished comprehension tasks with GD achieved an improved comprehension of source code with a higher effort with respect to those that accomplished tasks with TD. The outcomes suggest that a given experience is needed to benefit from the documentation of design-pattern instances in the comprehension of source code and to reduce the average time to accomplish comprehension tasks when this kind of documentation is given in addition to source code.

## 5.3. Implications

To assess the implications of our family of experiments, we adopted a perspective-based approach [Basili et al. 1996]. In particular, we focus here on the practitioner/consultant (simply practitioner in the following) and researcher's perspectives [Kitchenham et al. 2008]:

- The use of textually or graphically documented instances improves the correctness of understanding of source code given that the maintainer has an adequate level of experience. This result is clearly relevant from the researcher's perspective. For example, the researcher could be interested in studying how source code comprehensibility is affected when using together textually and graphically documented pattern instances. The practitioner could also be interested in our results because it is worth documenting pattern instances to get an improved source code comprehension. However, this outcome opens a managerial dilemma: are the additional effort and cost to create and maintain the documentation of pattern instances, adequately paid back by an improved source code comprehensibility? Indeed, from a managerial point of view, the adoption of graphical and textual documentation, as means to represent pattern instances, should take into account the costs it will introduce. Furthermore, what method is less expensive in representing design-pattern

- instances? On the basis of our findings, these points represent future directions for the research.
- The presence of documented instances does not seem to be a cause of distraction for the maintainer (Hn1.Time was not rejected in all the experiments), while performing comprehension tasks. In each experiment, the use of textual and graphical instances reduces the average time to comprehend source code. This result is obviously relevant for the practitioner, but also for the researcher. In particular, the researcher could be interested in investigating how to better support the maintainer in quickly finding implementation details from the pattern instances to reduce the effort to comprehend source code.
  - Benefits deriving from the documentation of pattern instances appear to be dependent on the participants' experience. This result is practically relevant for the professional: maintainers with different experiences shall benefit from the documented instances to a different degree. This aspect is clearly relevant also for the researcher, who could be interested in studying what is the possible experience threshold to benefit from the use of instance documentation in source code comprehension. From the researcher perspective, this outcome also adds an empirical voice to the debate over the use of students as valid proxies in developer-based studies [Carver et al. 2003; Höst et al. 2000; Arisholm and Sjøberg 2004]. That is, can students be considered as appropriate as professionals in controlled experiments? On the basis of our results, it seems that students can if they have a sufficient knowledge on the technology/method that is being evaluated. Therefore, the use of students might be an issue as long as we are interested in evaluating the adoption of a technology/method by expert software engineers and if they do not have a sufficient knowledge on that technology/method. On the other hand, the use of students might be valuable when our goal is to study how easy to use is a technology/method and its possible benefit without or with a limited training or experience (e.g., [Erra and Scanniello 2010; Scanniello and Erra 2014]).
  - The experiment with PhD students replicates the findings of the original experiment. This implication is related to the implication in the prior bullet point. In fact, PhD students can be as appropriate as professionals in controlled experiments. This implication is relevant for the researcher because it suggests that PhD students might substitute for professionals in developer-based studies. Although this point deserves future investigations, we can postulate that our findings pose the basis for future work in this direction.
  - Comprehensibility improves when pattern instances are correctly recognized in source code. This opens the following future research directions: *(i)* investigating issues that led to certain patterns to be better comprehended and recognized than others, *(ii)* defining new methods/notations to make the identification of pattern instances in source code easier, and *(iii)* investigating source code comprehensibility when graphically documented instances are complemented by other notations (e.g., UML sequence diagrams or textual documentation). The results from our family of experiments delineate these future research directions.
  - Since comprehensibility improves when source code is complemented by textually or graphically documented instances, the number of defects should decrease correspondingly as a consequence of a maintenance operation. This is true provided that the maintainer has an adequate level of experience (i.e., at least a Bachelor degree). The use of documented instances should also positively impact the communication among developers. These implications are relevant for the practitioner.
  - The majority of the participants found the source code to be the most relevant source of information to perform comprehension tasks. This finding is relevant for the re-

searcher, who could be interested in assessing if and how this concern affects the benefits stemming from the use of documented pattern instances.

- We considered pattern instances to be graphically documented using UML class diagrams. For each instance, we exploited the same graphical layout as proposed by Gamma *et al.* [Gamma et al. 1995]. The use of different layouts might lead to different results. This aspect is relevant for the practitioner interested in understanding the best way for representing pattern instances, and for the researcher interested in investigating how and why a different layout for a pattern instance affects source code comprehensibility.
- The diffusion of a new technology/method is made easier when empirical evaluations are performed and their results show that such a technology/method solves actual issues [Baldassarre et al. 2014; Pfleeger and Menezes 2000]. Therefore, the results from our family of experiments could increase the practice of documenting pattern instances in the software industry. In addition, even if the costs related to software development, maintenance, and evolution might increase, the introduction or the enforcement of the systematical use of documented instances should not require a complete and radical process change in a given company. This point has particular interest for the practitioner.

## 6. RELATED WORK

Software maintenance is essential in the evolution of software systems and represents one of the most expensive, time consuming, and challenging phases of the whole development process. Maintenance starts after the delivery of the first version of the system and lasts much longer than the initial development process [Bennett and Rajlich 2000; Zelkowitz et al. 1979]. As shown in the survey by Erlikh [2000], the cost needed to perform maintenance operations ranges from 85% to 90% of the total cost of a software project. Whatever is the maintenance operation, the greater part of its cost and effort is due to the comprehension of source code [Mayrhauser 1995]. In particular, Pfleeger and Atlee [2006] estimated that up to 60% of software maintenance is spent on comprehension. There are several reasons that make comprehension even more costly and complex, namely the size of a software system and its available documentation [Selfridge et al. 1993].

Only few studies have been conducted to assess the support that the documentation of pattern instances provides to software maintenance [Prechelt et al. 2002; Gravino et al. 2011]. For example, Prechelt et al. [2002] studied whether pattern instances textually documented in source code (through comments) improve maintainers' performance in the execution of maintenance tasks with respect to well-commented source code without explicit references to pattern instances. The study involved 74 German graduate students and 22 undergraduate students from the USA, who performed tasks on small Java and C++ programs, respectively. For example, the size of these programs ranged from 495 to 762 lines of code (comment lines included), only two kinds of pattern instances were present in each program, and maintenance operations were executed on sheet papers. The results suggested that maintenance tasks supported by explicitly documented pattern instances were completed faster or with fewer errors. This is the study most similar to ours. However, there are many differences that make the study by Prechelt et al. [2002] different from the study we present in this paper. The main difference is that we conducted a family of controlled experiments involving participants having four different levels of experience: from Bachelor students to software professionals. In addition, we analyzed the effect of graphically documented pattern instances, we used a larger experimental object from a real open-source software system (2149 lines of code, comment lines included, with 10 pattern instances implemented in), and the participants conducted their task inside a more realistic en-

vironment, thus reducing threats to external validity. The results achieved in our long term investigation together with those from Prechelt et al. [2002] demonstrate the usefulness of explicitly documenting instances when dealing with source code comprehension. Therefore, the main “take away” lesson from our family of experiments and from the two experiments by Prechelt *et al.* is: documenting instances can improve source code comprehensibility. That is, it is not only important to use design-patterns to develop an object-oriented software system, but it is even more important to properly document the instances implemented in its source code. These results are perhaps not overly surprising, but it is acceptable if not significant, as evidence needs to be verified/reaffirmed through empirical studies [Basili et al. 1999; Kitchenham 2008; Kitchenham et al. 2002; Shull et al. 2008].

The research work presented in this paper is also different from other investigations conducted to examine software design-patterns [Prechelt et al. 2001; Bieman et al. 2003; Cepeda Porras and Guéhéneuc 2010; Jeanmart et al. 2009; Khomh and Guéhéneuc 2008; Di Penta et al. 2008; Vokác 2004; Vokác et al. 2004; Krein et al. 2011], since we pursued a different goal. In particular, the focus here is on the kind of documentation for pattern instances and on source code comprehensibility. For example, Prechelt et al. [2001] investigate whether the use of design-patterns is useful or harmful. To this end, their experiment involved four different, rather small, C++ programs and six software design-patterns. Each of these programs existed in two versions: one based on design-patterns and another implementing a simplified design without pattern instances. This empirical investigation has been next replicated by Vokác et al. [2004] and Krein et al. [2011]. All the investigations reinforce the conclusion that each kind of design-pattern has its own nature and proper place of use. The results achieved in these investigations also suggest that some kinds of design-patterns are much easier to understand and use than others. The implication is that design-patterns are not universally good or bad, they must be used in a way that matches the problem and the maintainer. To this respect, our work complements these investigations by explicitly considering the source code comprehension that four kinds of maintainers achieved when source code is developed exploiting design-patterns and the instances implemented in the source code are textually and graphically documented.

## 7. CONCLUSION

In this paper, we have presented results from a family of four controlled experiments, carried out with students and practitioners. We used controlled experiments because a number of confounding and uncontrollable factors might be present in real project settings. The achieved results suggest that at least a given experience is needed to benefit from the documentation of design-pattern instances as a complementary information to comprehend source code. For professionals and PhD students, the use of graphically and textually documented pattern instances significantly improves correctness of understanding of source code. We did not observe any significant difference between textually and graphically documented pattern instances with respect to both correctness of understanding and task completion time. Our results also suggest that correctness of understanding improves when pattern instances are correctly identified. It seems crucial to ease the identification of instances in source code.

Possible future directions for our research are: (i) performing different kinds of empirical investigations (e.g., case studies); (ii) investigating the effect of documentation for specific design-patterns on source code comprehensibility; (iii) studying new representations for documenting pattern instances; and (iv) investigating if providing textual documentation alongside with the graphical one improves source code comprehensibility.

## ACKNOWLEDGMENTS

The authors would like to thank all the participants in the family of experiments.

## REFERENCES

- ABRAHÃO, S. M., GRAVINO, C., PELOZO, E. I., SCANNIELLO, G., AND TORTORA, G. 2013. Assessing the effectiveness of sequence diagrams in the comprehension of functional requirements: Results from a family of five experiments. *IEEE Transactions on Software Engineering* 39, 3.
- ALEXANDER, C., ISHIKAWA, S., SILVERSTEIN, M., JACOBSON, M., FIKSDAHL-KING, I., AND ANGEL, S. 1977. *A Pattern Language - Towns, Buildings, Construction*. Oxford University Press.
- ARANDA, J., ERNST, N., HORKOFF, J., AND EASTERBROOK, S. 2007. A framework for empirical evaluation of model comprehensibility. In *Proceedings of Modeling in Software Engineering*. IEEE Computer Society, 7–13.
- ARISHOLM, E., BRIAND, L. C., HOVE, S. E., AND LABICHE, Y. 2006. The impact of UML documentation on software maintenance: An experimental evaluation. *IEEE Transactions on Software Engineering* 32, 6, 365–381.
- ARISHOLM, E., GALLIS, H., DYBÅ, T., AND SJØBERG, D. I. K. 2007. Evaluating pair programming with respect to system complexity and programmer expertise. *IEEE Transactions on Software Engineering* 33, 2, 65–86.
- ARISHOLM, E. AND SJØBERG, D. I. K. 2004. Evaluating the effect of a delegated versus centralized control style on the maintainability of object-oriented software. *IEEE Transactions on Software Engineering* 30, 8, 521–534.
- ASUNCION, H. U., FRANÇOIS, F., AND TAYLOR, R. N. 2007. An end-to-end industrial software traceability tool. In *Proceedings of the Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on The Foundations of Software Engineering*. ACM, 115–124.
- BAKER, R. 1995. Modern permutation test software. In *E. Edgington Randomization Tests*, New York, Marcel Dekker.
- BALDASSARRE, M. T., CARVER, J., DIESTE, O., AND JURISTO, N. 2014. Replication types: Towards a shared taxonomy. In *Proceedings of International Conference on Evaluation and Assessment in Software Engineering*. ACM, 18:1–18:4.
- BASIL, V., SHULL, F., AND LANUBILE, F. 1999. Building knowledge through families of experiments. *IEEE Transactions on Software Engineering* 25, 4, 456–473.
- BASIL, V. R., GREEN, S., LAITENBERGER, O., LANUBILE, F., SHULL, F., SØRUMGÅRD, L. S., AND ZELKOWITZ, M. V. 1996. The empirical investigation of perspective-based reading. *Empirical Software Engineering* 1, 2, 133–164.
- BASIL, V. R. AND ROMBACH, H. D. 1988. The TAME project: Towards improvement-oriented software environments. *IEEE Transactions on Software Engineering* 14, 6, 758–773.
- BAUER, M. I. AND JOHNSON-LAIRD, P. N. 1993. How diagrams can improve reasoning. *Psychological Science* 4, 372–378.
- BENNETT, K. H. AND RAJLICH, V. T. 2000. Software maintenance and evolution: a roadmap. In *Proceedings of the Conference on the Future of Software Engineering*. ICSE '00. ACM, New York, NY, USA, 73–87.
- BERGERSEN, G., SJØBERG, D. I. K., AND DYBÅ, T. 2014. Construction and validation of an instrument for measuring programming skill. *IEEE Transactions on Software Engineering*.
- BERGERSEN, G. R., HANNAY, J. E., SJØBERG, D. I. K., DYBÅ, T., AND KARAHASANOVIC, A. 2011. Inferring skill from tests of programming performance: Combining time and quality. In *Procs. of the International Symposium on Empirical Software Engineering*. IEEE, 305–314.
- BIEMAN, J., STRAW, G., WANG, H., MUNGER, P., AND ALEXANDER, R. 2003. Design patterns and change proneness: an examination of five evolving systems. In *Procs. of Software Metrics Symposium*. IEEE Computer Society, 40–49.
- BRUEGGE, B. AND DUTOIT, A. H. 2003. *Object-Oriented Software Engineering: Using UML, Patterns and Java, 2nd edition*. Prentice-Hall.
- BUDGEN, D., BURN, A. J., BRERETON, O. P., KITCHENHAM, B. A., AND PRETORIUS, R. 2011. Empirical evidence about the UML: a systematic literature review. *Software: Practice and Experience* 41, 4, 363–392.
- CARVER, J., JACCHERI, L., MORASCA, S., AND SHULL, F. 2003. Issues in using students in empirical studies in software engineering education. In *Proceedings of International Symposium on Software Metrics*. IEEE Computer Society, Washington, DC, USA, 239–250.



- CECCATO, M., DI PENTA, M., FALCARIN, P., RICCA, F., TORCHIANO, M., AND TONELLA, P. 2014. A family of experiments to assess the effectiveness and efficiency of source code obfuscation techniques. *Empirical Software Engineering* 19, 4, 1040–1074.
- CEPEDA PORRAS, G. AND GUÉHÉNEUC, Y.-G. 2010. An empirical study on the efficiency of different design pattern representations in UML class diagrams. *Empirical Software Engineering* 15, 5, 493–522.
- CIOLKOWSKI, M., MUTHIG, D., AND RECH, J. 2004. Using academic courses for empirical validation of software development processes. In *Proceedings of EUROMICRO Conference*. IEEE Computer Society, Washington, DC, USA, 354–361.
- COHEN, J. 1988. *Statistical Power Analysis for the Behavioral Sciences*. L. Erlbaum Associates.
- COLOSIMO, M., DE LUCIA, A., SCANNIELLO, G., AND TORTORA, G. 2009. Evaluating legacy system migration technologies through empirical studies. *Information & Software Technology* 51, 2, 433–447.
- CONOVER, W. J. 1998. *Practical Nonparametric Statistics*. Wiley, 3rd Edition.
- CRONBACH, L. 1951. Coefficient alpha and the internal structure of tests. *Psychometrika* 16, 3, 297–334.
- DA SILVA, F. Q. B., SUASSUNA, M., FRANÇA, A. C. C., GRUBB, A. M., GOUVEIA, T. B., MONTEIRO, C. V. F., AND DOS SANTOS, I. E. 2014. Replication of empirical studies in software engineering research: a systematic mapping study. *Empirical Software Engineering* 19, 3, 501–557.
- DI PENTA, M., CERULO, L., GUÉHÉNEUC, Y.-G., AND ANTONIOL, G. 2008. An empirical study of the relationships between design pattern roles and class change proneness. In *Procs. of the International Conference on Software Maintenance*. IEEE Computer Society, 217–226.
- DYBÅ, T., KAMPENES, V. B., AND SJØBERG, D. I. K. 2006. A systematic review of statistical power in software engineering experiments. *Information & Software Technology* 48, 8, 745–755.
- ELLIS, P. 2010. *The Essential Guide to Effect Sizes: Statistical Power, Meta-Analysis, and the Interpretation of Research Results*. Cambridge University Press.
- ERLIKH, L. 2000. Leveraging legacy system dollars for e-business. *IT Professional* 2, 17–23.
- ERRA, U. AND SCANNIELLO, G. 2010. Assessing communication media richness in requirements negotiation. *IET Software* 4, 2, 134–148.
- GAMMA, E., HELM, R., R. JOHNSON, AND VLISSIDES, J. 1995. *Design Patterns: Elements of Reusable Object Oriented Software*. Addison-Wesley.
- GÓMEZ, O. S., JURISTO, N., AND VEGAS, S. 2010. Replications types in experimental disciplines. In *Proceedings of International Symposium on Empirical Software Engineering and Measurement*. ACM, 3:1–3:10.
- GRAVINO, C., RISI, M., SCANNIELLO, G., AND TORTORA, G. 2011. Does the documentation of design pattern instances impact on source code comprehension? Results from two controlled experiments. In *Proceedings of the Working Conference on Reverse Engineering*. IEEE Computer Society, 67–76.
- GRAVINO, C., RISI, M., SCANNIELLO, G., AND TORTORA, G. 2012. Do professional developers benefit from design pattern documentation? a replication in the context of source code comprehension. In *Procs. of International Conference on Model Driven Engineering Languages and Systems*. Lecture Notes in Computer Science. Springer, 185–201.
- GUÉHÉNEUC, Y.-G. 2007. P-mart: Pattern-like micro architecture repository. In *Procs. of EuroPLoP Focus Group on Pattern Repositories*.
- GUÉHÉNEUC, Y.-G. AND ANTONIOL, G. 2008. Demima: A multilayered approach for design pattern identification. *IEEE Transactions on Software Engineering* 34, 5, 667–684.
- HANNAY, J. AND JØRGENSEN, M. 2008. The role of deliberate artificial design elements in software engineering experiments. *IEEE Transactions on Software Engineering* 34, 2, 242–259.
- HEER, J. AND AGRAWALA, M. 2006. Software design patterns for information visualization. *IEEE Trans. Visualization & Comp. Graphics* 12, 853–860.
- HOCHSTEIN, L., BASILI, V. R., ZELKOWITZ, M. V., HOLLINGSWORTH, J. K., AND CARVER, J. 2005. Combining self-reported and automatic data to improve programming effort measurement. *ACM SIGSOFT Software Engineering Notes* 30, 5, 356–365.
- HÖST, M., REGNELL, B., AND WOHLIN, C. 2000. Using students as subjects—a comparative study of students and professionals in lead-time impact assessment. *Empirical Software Engineering* 5, 3, 201–214.
- HUANG, L. AND HOLCOMBE, M. 2009. Empirical investigation towards the effectiveness of test first programming. *Information & Software Technology* 51, 1, 182–194.
- ISO. 1991. *Information Technology—Software Product Evaluation: Quality Characteristics and Guidelines for their Use, ISO/IEC IS 9126*. ISO, Geneva.
- JEANMART, S., GUÉHÉNEUC, Y.-G., SAHRAOUI, H., AND HABRA, N. 2009. Impact of the Visitor Pattern on program comprehension and maintenance. In *Procs. of the Symposium on Empirical Software Engineering and Measurement*. IEEE Computer Society, 69–78.

- JEDLITSCHKA, A., CIOLKOWSKI, M., AND PFAHL, D. 2008. Reporting experiments in software engineering. In *Guide to Advanced Empirical Software Engineering* (Eds) F. Shull and J. Singer and D. Sjöberg. Springer London, 201–228.
- JURISTO, N. AND MORENO, A. 2001. *Basics of Software Engineering Experimentation*. Kluwer Academic Publishers.
- KAMPENES, V., DYBA, T., HANNAY, J., AND SJOBERG, I. 2007. A systematic review of effect size in software engineering experiments. *Information & Software Technology* 49, 11–12, 1073–1086.
- KHOMH, F. AND GUÉHÉNEUC, Y.-G. 2008. Do design patterns impact software quality positively? In *Proceedings of Conference on Software Engineering and Maintenance*. 274–278.
- KITCHENHAM, B. 2008. The role of replications in empirical software engineering - a word of warning. *Empirical Software Engineering* 13, 2, 219–221.
- KITCHENHAM, B., AL-KHILIDAR, H., BABAR, M., BERRY, M., COX, K., KEUNG, J., KURNIAWATI, F., STAPLES, M., ZHANG, H., AND ZHU, L. 2008. Evaluating guidelines for reporting empirical software engineering studies. *Empirical Software Engineering* 13, 97–121.
- KITCHENHAM, B., PFLEEGER, S., PICKARD, L., JONES, P., HOAGLIN, D., EL EMAM, K., AND ROSENBERG, J. 2002. Preliminary guidelines for empirical research in software engineering. *IEEE Transactions on Software Engineering* 28, 8, 721–734.
- KREIN, J. L., PRATT, L. J., SWENSON, A. B., MACLEAN, A. C., KNUTSON, C. D., AND EGGETT, D. L. 2011. Design patterns in software maintenance: An experiment replication at brigham young university. *Proceedings of International Workshop on Replication in Empirical Software Engineering Research*, 25–34.
- LINDVALL, M., RUS, I., SHULL, F., ZELKOWITZ, M. V., DONZELLI, P., MEMON, A. M., BASILI, V. R., COSTA, P., TVEDT, R. T., HOCHSTEIN, L., ASGARI, S., ACKERMANN, C., AND PECH, D. 2005. An evolutionary testbed for software technology evaluation. *Innovations in Systems and Software Engineering* 1, 1, 3–11.
- LINDVALL, M. AND SANDAHL, K. 1996. Practical implications of traceability. *Software, Practice and Experience* 26, 10, 1161–1180.
- MAXWELL, K. 2002. *Applied Statistics for Software Managers*. Software Quality Institute Series. Prentice Hall PTR.
- MAYRHAUSER, A. V. 1995. Program comprehension during software maintenance and evolution. *IEEE Computer* 28, 44–55.
- MCDERMID, J. 1991. *Software Engineer's Reference Book*. Butterworth-Heinemann, Linacre House, Jordan Hill, Oxford, UK.
- MENDONÇA, M. G., MALDONADO, J. C., OLIVEIRA, M. C. F. D., CARVER, J., FABBRI, S. C. P. F., SHULL, F., TRAVASSOS, G. H., HÖHN, E. N., AND BASILI, V. R. 2008. A framework for software engineering experimental replications. In *Proceedings of International Conference on Engineering of Complex Computer Systems*. IEEE Computer Society, 203–212.
- OMG. 2005. Unified modeling language (UML) specification, version 2.0. Tech. rep., Object Management Group. July.
- OPPENHEIM, A. N. 1992. *Questionnaire Design, Interviewing and Attitude Measurement*. Pinter, London.
- PFLEEGER, S. AND ATLEE, J. 2006. *Software engineering - theory and practice* (3. ed.). Ellis Horwood.
- PFLEEGER, S. L. AND MENEZES, W. 2000. Marketing technology to software practitioners. *IEEE Software* 17, 1, 27–33.
- PRECHELT, L., UNGER, B., TICHY, W. F., BRÖSSLER, P., AND VOTTA, L. G. 2001. A controlled experiment in maintenance comparing design patterns to simpler solutions. *IEEE Transactions on Software Engineering* 27, 12, 1134–1144.
- PRECHELT, L., UNGER-LAMPRECHT, B., PHILIPPSEN, M., AND TICHY, W. 2002. Two controlled experiments assessing the usefulness of design pattern documentation in program maintenance. *IEEE Transactions on Software Engineering* 28, 6, 595–606.
- RICCA, F., DI PENTA, M., TORCHIANO, M., TONELLA, P., AND CECCATO, M. 2010. How developers' experience and ability influence web application comprehension tasks supported by uml stereotypes: A series of four experiments. *IEEE Transactions on Software Engineering* 36, 1, 96–118.
- RICCA, F., SCANNIELLO, G., TORCHIANO, M., REGGIO, G., AND ASTESIANO, E. 2014. Assessing the effect of screen mockups on the comprehension of functional requirements. *ACM Transactions on Software Engineering and Methodology* 24, 1, 1:1–1:38.
- SALTON, G. AND MCGILL, M. J. 1983. *Introduction to Modern Information Retrieval*. McGraw Hill, New York.

- SCAIFE, M. AND ROGERS, Y. 1996. External cognition: how do graphical representations work? *International Journal of Human-Computer Studies* 45, 2, 185–213.
- SCANNIELLO, G. AND ERRA, U. 2014. Distributed modeling of use case diagrams with a method based on think-pair-square: Results from two controlled experiments. *Journal of Visual Languages and Computing* 25, 4, 494–517.
- SCANNIELLO, G., GRAVINO, C., GENERO, M., CRUZ-LEMUS, J. A., AND TORTORA, G. 2014a. On the impact of UML analysis models on source code comprehensibility and modifiability. *ACM Transactions on Software Engineering and Methodology* 23, 2.
- SCANNIELLO, G., GRAVINO, C., RISI, M., TORTORA, G., AND DODERO, G. 2014b. Technical report: please copy and paste the following url; if it does not work, please check it carefully because there could be an extraspace to be removed. <http://www2.unibas.it/gscanniello/DP/>.
- SCANNIELLO, G., GRAVINO, C., AND TORTORA, G. 2013. An early investigation on the contribution of class and sequence diagrams in source code comprehension. In *European Conference on Software Maintenance and Reengineering*. IEEE Computer Society, 367–370.
- SELFRIDGE, P., WATERS, R., AND CHIKOFSKY, E. 1993. Challenges to the field of reverse engineering. In *Proceedings of the Working Conference on Reverse Engineering*. IEEE Computer Society.
- SHAPIRO, S. AND WILK, M. 1965. An analysis of variance test for normality. *Biometrika* 52, 3-4, 591–611.
- SHULL, F., CARVER, J. C., VEGAS, S., AND JUZGADO, N. J. 2008. The role of replications in empirical software engineering. *Empirical Software Engineering* 13, 2, 211–218.
- SHULL, F., MENDONÇA, M. G., BASILI, V., CARVER, J., MALDONADO, J. C., FABBRI, S., TRAVASSOS, G. H., AND FERREIRA, M. C. 2004. Knowledge-sharing issues in experimental software engineering. *Empirical Software Engineering* 9, 1-2, 111–137.
- SILLITO, J., MURPHY, G. C., AND DE VOLDER, K. 2008. Asking and answering questions during a programming change task. *IEEE Transactions on Software Engineering* 34, 4, 434–451.
- SJØBERG, D. I. K., HANNAY, J. E., HANSEN, O., KAMPENES, V. B., KARAHASANOVIC, A., LIBORG, N., AND REKDAL, A. C. 2005. A survey of controlled experiments in software engineering. *IEEE Transactions on Software Engineering* 31, 9, 733–753.
- VEGAS, S., JURISTO, N., MORENO, A., SOLARI, M., AND LETELIER, P. 2006. Analysis of the influence of communication between researchers on experiment replication. In *Proceedings of the International Symposium on Empirical Software Engineering*. ACM, New York, NY, USA, 28–37.
- VERELST, J. 2004. The influence of the level of abstraction on the evolvability of conceptual models of information systems. In *Proceedings of the International Symposium on Empirical Software Engineering*. IEEE Computer Society, Washington, DC, USA, 17–26.
- VOKÁČ, M. 2004. Defect frequency and design patterns: An empirical study of industrial code. *IEEE Transactions on Software Engineering* 30, 12, 904–917.
- VOKÁČ, M., TICHY, W. F., SJØBERG, D. I. K., ARISHOLM, E., AND ALDRIN, M. 2004. A controlled experiment comparing the maintainability of programs designed with and without design patterns—a replication in a real programming environment. *Empirical Software Engineering* 9, 3, 149–195.
- WOHLIN, C., RUNESON, P., HÖST, M., OHLSSON, M., REGNELL, B., AND WESSLÉN, A. 2012. *Experimentation in Software Engineering*. Springer.
- ZELKOWITZ, M. V., SHAW, A. C., AND GANNON, J. D. 1979. *Principles of software engineering and design*. Prentice-Hall.

## A. COMPREHENSION QUESTIONNAIRE

In the following, we report the complete list of questions in the used comprehension questionnaire for the treatment SC. These questions are divided into three groups (see Section 3.6). For each question, we also show the possible answers for confidence, perceived difficulty, and source of information. As for the treatments GD and TD, we also asked participants to specify the pattern instances they used to answer the question when design-pattern instances were indicated as the source of information.

**First group**

<b>Q1.</b> What are the concrete classes, the abstract classes, and the interfaces that describe the geometric shape rectangle and its attributes?				
How much do you trust your answer <sup>+</sup> ?				
<input type="checkbox"/> Unsure	<input type="checkbox"/> Not sure enough	<input type="checkbox"/> Sure Enough	<input type="checkbox"/> Sure	<input type="checkbox"/> Very Sure
How do you assess the question <sup>+</sup> ?				
<input type="checkbox"/> Very difficult	<input type="checkbox"/> Difficult	<input type="checkbox"/> On average	<input type="checkbox"/> Simple	<input type="checkbox"/> Very Simple
What is the source of information used to answer the question <sup>+</sup> ?				
<input type="checkbox"/> Previous Knowledge (PK)		<input type="checkbox"/> Internet (I)	<input type="checkbox"/> Source Code (SoC)	
+ Mark only one answer				
<b>Q2.</b> What are the main methods involved in creation and visualization of a Figure object? Indicate the classes to which they belong as well.				
How much do you trust your answer <sup>+</sup> ?				
<input type="checkbox"/> Unsure	<input type="checkbox"/> Not sure enough	<input type="checkbox"/> Sure Enough	<input type="checkbox"/> Sure	<input type="checkbox"/> Very Sure
How do you assess the question <sup>+</sup> ?				
<input type="checkbox"/> Very difficult	<input type="checkbox"/> Difficult	<input type="checkbox"/> On average	<input type="checkbox"/> Simple	<input type="checkbox"/> Very Simple
What is the source of information used to answer the question <sup>+</sup> ?				
<input type="checkbox"/> Previous Knowledge (PK)		<input type="checkbox"/> Internet (I)	<input type="checkbox"/> Source Code (SoC)	
+ Mark only one answer				
<b>Q3.</b> What are the concrete classes, abstract classes, and interfaces that manage the composition of a whole drawing and its changes?				
How much do you trust your answer <sup>+</sup> ?				
<input type="checkbox"/> Unsure	<input type="checkbox"/> Not sure enough	<input type="checkbox"/> Sure Enough	<input type="checkbox"/> Sure	<input type="checkbox"/> Very Sure
How do you assess the question <sup>+</sup> ?				
<input type="checkbox"/> Very difficult	<input type="checkbox"/> Difficult	<input type="checkbox"/> On average	<input type="checkbox"/> Simple	<input type="checkbox"/> Very Simple
What is the source of information used to answer the question <sup>+</sup> ?				
<input type="checkbox"/> Previous Knowledge (PK)		<input type="checkbox"/> Internet (I)	<input type="checkbox"/> Source Code (SoC)	
+ Mark only one answer				
<b>Q4.</b> What are the instance variables (their classes have to be indicated as well) that are modified (i.e., those that change state) when a figure is decorated?				
How much do you trust your answer <sup>+</sup> ?				
<input type="checkbox"/> Unsure	<input type="checkbox"/> Not sure enough	<input type="checkbox"/> Sure Enough	<input type="checkbox"/> Sure	<input type="checkbox"/> Very Sure
How do you assess the question <sup>+</sup> ?				
<input type="checkbox"/> Very difficult	<input type="checkbox"/> Difficult	<input type="checkbox"/> On average	<input type="checkbox"/> Simple	<input type="checkbox"/> Very Simple
What is the source of information used to answer the question <sup>+</sup> ?				
<input type="checkbox"/> Previous Knowledge (PK)		<input type="checkbox"/> Internet (I)	<input type="checkbox"/> Source Code (SoC)	
+ Mark only one answer				
<b>Q5.</b> What are the instance variables (their classes have to be indicated as well) that are modified (i.e., those that change state) when a new geometric shape is added?				
How much do you trust your answer <sup>+</sup> ?				
<input type="checkbox"/> Unsure	<input type="checkbox"/> Not sure enough	<input type="checkbox"/> Sure Enough	<input type="checkbox"/> Sure	<input type="checkbox"/> Very Sure
How do you assess the question <sup>+</sup> ?				
<input type="checkbox"/> Very difficult	<input type="checkbox"/> Difficult	<input type="checkbox"/> On average	<input type="checkbox"/> Simple	<input type="checkbox"/> Very Simple
What is the source of information used to answer the question <sup>+</sup> ?				
<input type="checkbox"/> Previous Knowledge (PK)		<input type="checkbox"/> Internet (I)	<input type="checkbox"/> Source Code (SoC)	
+ Mark only one answer				

**Second group**

<b>Q6.</b> What are the abstract classes and the interfaces that participate in the feature <i>"Adding a geometric figure to another geometric figure"</i> ?				
_____				
How much do you trust your answer <sup>+</sup> ?				
<input type="checkbox"/> Unsure	<input type="checkbox"/> Not sure enough	<input type="checkbox"/> Sure Enough	<input type="checkbox"/> Sure	<input type="checkbox"/> Very Sure
How do you assess the question <sup>+</sup> ?				
<input type="checkbox"/> Very difficult	<input type="checkbox"/> Difficult	<input type="checkbox"/> On average	<input type="checkbox"/> Simple	<input type="checkbox"/> Very Simple
What is the source of information used to answer the question <sup>+</sup> ?				
<input type="checkbox"/> Previous Knowledge (PK)		<input type="checkbox"/> Internet (I)	<input type="checkbox"/> Source Code (SoC)	
+ Mark only one answer				
<b>Q7.</b> What are the concrete classes, abstract classes, and the interfaces that manage localization and visualization of connection points of a geometric figure?				
_____				
How much do you trust your answer <sup>+</sup> ?				
<input type="checkbox"/> Unsure	<input type="checkbox"/> Not sure enough	<input type="checkbox"/> Sure Enough	<input type="checkbox"/> Sure	<input type="checkbox"/> Very Sure
How do you assess the question <sup>+</sup> ?				
<input type="checkbox"/> Very difficult	<input type="checkbox"/> Difficult	<input type="checkbox"/> On average	<input type="checkbox"/> Simple	<input type="checkbox"/> Very Simple
What is the source of information used to answer the question <sup>+</sup> ?				
<input type="checkbox"/> Previous Knowledge (PK)		<input type="checkbox"/> Internet (I)	<input type="checkbox"/> Source Code (SoC)	
+ Mark only one answer				
<b>Q8.</b> What methods allow to modify a given geometric figure through direct manipulation? Indicate the classes to which they belong as well.				
_____				
How much do you trust your answer <sup>+</sup> ?				
<input type="checkbox"/> Unsure	<input type="checkbox"/> Not sure enough	<input type="checkbox"/> Sure Enough	<input type="checkbox"/> Sure	<input type="checkbox"/> Very Sure
How do you assess the question <sup>+</sup> ?				
<input type="checkbox"/> Very difficult	<input type="checkbox"/> Difficult	<input type="checkbox"/> On average	<input type="checkbox"/> Simple	<input type="checkbox"/> Very Simple
What is the source of information used to answer the question <sup>+</sup> ?				
<input type="checkbox"/> Previous Knowledge (PK)		<input type="checkbox"/> Internet (I)	<input type="checkbox"/> Source Code (SoC)	
+ Mark only one answer				
<b>Q9.</b> What classes allow to insert a bitmap image in the drawing?				
_____				
How much do you trust your answer <sup>+</sup> ?				
<input type="checkbox"/> Unsure	<input type="checkbox"/> Not sure enough	<input type="checkbox"/> Sure Enough	<input type="checkbox"/> Sure	<input type="checkbox"/> Very Sure
How do you assess the question <sup>+</sup> ?				
<input type="checkbox"/> Very difficult	<input type="checkbox"/> Difficult	<input type="checkbox"/> On average	<input type="checkbox"/> Simple	<input type="checkbox"/> Very Simple
What is the source of information used to answer the question <sup>+</sup> ?				
<input type="checkbox"/> Previous Knowledge (PK)		<input type="checkbox"/> Internet (I)	<input type="checkbox"/> Source Code (SoC)	
+ Mark only one answer				
<b>Q10.</b> What methods allow to first select a figure and then to move it? Indicate the classes to which they belong as well.				
_____				
How much do you trust your answer <sup>+</sup> ?				
<input type="checkbox"/> Unsure	<input type="checkbox"/> Not sure enough	<input type="checkbox"/> Sure Enough	<input type="checkbox"/> Sure	<input type="checkbox"/> Very Sure
How do you assess the question <sup>+</sup> ?				
<input type="checkbox"/> Very difficult	<input type="checkbox"/> Difficult	<input type="checkbox"/> On average	<input type="checkbox"/> Simple	<input type="checkbox"/> Very Simple
What is the source of information used to answer the question <sup>+</sup> ?				
<input type="checkbox"/> Previous Knowledge (PK)		<input type="checkbox"/> Internet (I)	<input type="checkbox"/> Source Code (SoC)	
+ Mark only one answer				

### Third group

**Q11.** You have to perform a corrective maintenance operation (i.e., errors in the source code must be removed) and/or perfective maintenance operation (i.e., improving the code) on the following source code fragment. It defines and manages the ellipse geometric shape and it allows its usage. Correction should be performed directly on the paper.

**(Note: the following source code fragment is not present in the source code we gave you to perform the comprehension task.)**

```
public class FiguraEllisse extends RectangleFigure {
    private Ellisse fDisplayBox;

    public FiguraEllisse() {
        this(new Point(0,0), new Point(0,0));
    }

    public FiguraEllisse(Point origin, Point corner) {
        for (int i=0; i < fDisplayBox.x; i++);
        basicDisplayBox(origin, corner);
    }

    public void basicDisplayBox(Point origin, Point corner) {
        fDisplayBox = new Rectangle(origin);
        fDisplayBox.add(corner);
        fDisplayBox.add(origin);
    }

    public Vector handles() {
        Vector handles = new Vector();
        BoxHandleKit.addHandles(this, handles);
        handles = new vector();
        return handles;
    }

    public Rectangle displayBox() {
        return new Rectangle( fDisplayBox.x, 0, fDisplayBox.width, fDisplayBox.height );
    }

    protected void basicMoveBy(int x, int y) {
        fDisplayBox.translate(x, y);
    }

    public void drawBackground(Graphics g) {
        Rectangle r = displayBox();
        g.fillOval(r.x, r.y, r.width, r.height);
    }

    public void drawFrame(Graphics g) {
        Rectangle r = displayBox();
        g.drawOval(r.x, r.y, r.width-1, r.height-1);
    }

    public Insets connectionInsets() {
        Rectangle r = fDisplayBox;
        int cx = r.width/2;
        int cy = r.height/2;
        return new Insets(cy, cx, cy, cx);
    }

    public Connector connectorAt(int x, int y) {
        return new ChopEllipseConnector(this);
    }

    public void write(StorableOutput dw) {
        super.write(dw);
        dw.writeInt(fDisplayBox.x);
        dw.writeInt(fDisplayBox.y);
        dw.writeInt(fDisplayBox.width);
        dw.writeInt(fDisplayBox.height);
    }

    public void read(StorableInput dr) throws IOException {
        super.read(dr);
        fDisplayBox = new Rectangle(
            dr.readInt(), dr.readInt(), dr.readInt(), dr.readInt());
    }
}
```

How much do you trust your answer <sup>+</sup> ?				
<input type="checkbox"/> Unsure	<input type="checkbox"/> Not sure enough	<input type="checkbox"/> Sure Enough	<input type="checkbox"/> Sure	<input type="checkbox"/> Very Sure
How do you assess the question <sup>+</sup> ?				
<input type="checkbox"/> Very difficult	<input type="checkbox"/> Difficult	<input type="checkbox"/> On average	<input type="checkbox"/> Simple	<input type="checkbox"/> Very Simple
What is the source of information used to answer the question <sup>+</sup> ?				
<input type="checkbox"/> Previous Knowledge (PK)		<input type="checkbox"/> Internet (I)	<input type="checkbox"/> Source Code (SoC)	
<sup>+</sup> Mark only one answer				
<b>Q12.</b> If the system displays a geometric figure in a position different from where it was originally drawn, what classes should be analyzed to identify the fault?				
_____				
How much do you trust your answer <sup>+</sup> ?				
<input type="checkbox"/> Unsure	<input type="checkbox"/> Not sure enough	<input type="checkbox"/> Sure Enough	<input type="checkbox"/> Sure	<input type="checkbox"/> Very Sure
How do you assess the question <sup>+</sup> ?				
<input type="checkbox"/> Very difficult	<input type="checkbox"/> Difficult	<input type="checkbox"/> On average	<input type="checkbox"/> Simple	<input type="checkbox"/> Very Simple
What is the source of information used to answer the question <sup>+</sup> ?				
<input type="checkbox"/> Previous Knowledge (PK)		<input type="checkbox"/> Internet (I)	<input type="checkbox"/> Source Code (SoC)	
<sup>+</sup> Mark only one answer				
<b>Q13.</b> If the system selects a different tool from the one chosen by the user (e.g., the pencil instead of the selected rubber), what concrete classes need to be analyzed?				
_____				
How much do you trust your answer <sup>+</sup> ?				
<input type="checkbox"/> Unsure	<input type="checkbox"/> Not sure enough	<input type="checkbox"/> Sure Enough	<input type="checkbox"/> Sure	<input type="checkbox"/> Very Sure
How do you assess the question <sup>+</sup> ?				
<input type="checkbox"/> Very difficult	<input type="checkbox"/> Difficult	<input type="checkbox"/> On average	<input type="checkbox"/> Simple	<input type="checkbox"/> Very Simple
What is the source of information used to answer the question <sup>+</sup> ?				
<input type="checkbox"/> Previous Knowledge (PK)		<input type="checkbox"/> Internet (I)	<input type="checkbox"/> Source Code (SoC)	
<sup>+</sup> Mark only one answer				
<b>Q14.</b> If the system does not react to changes on a drawing made by the user, what concrete and abstract classes should be analyzed?				
_____				
How much do you trust your answer <sup>+</sup> ?				
<input type="checkbox"/> Unsure	<input type="checkbox"/> Not sure enough	<input type="checkbox"/> Sure Enough	<input type="checkbox"/> Sure	<input type="checkbox"/> Very Sure
How do you assess the question <sup>+</sup> ?				
<input type="checkbox"/> Very difficult	<input type="checkbox"/> Difficult	<input type="checkbox"/> On average	<input type="checkbox"/> Simple	<input type="checkbox"/> Very Simple
What is the source of information used to answer the question <sup>+</sup> ?				
<input type="checkbox"/> Previous Knowledge (PK)		<input type="checkbox"/> Internet (I)	<input type="checkbox"/> Source Code (SoC)	
<sup>+</sup> Mark only one answer				

## B. DESIGN-PATTERN INSTANCES

In Table XII, we report some information about the original chunk of JHotDraw v5.1 we used in our family of experiments. In particular, we show whether each design-pattern instance was textually documented in the source code. This documentation is present for each pattern instance, but not all classes of a given pattern instance contain that documentation. We also report whether class names were chosen to make the role these classes played in each instance to be clear. In addition, for each class its role in the instance is also indicated.

Table XII. Details on the design-pattern instances in the used source code of JHotDraw

Design-Pattern	Documented	Naming Convention
<b>Adapter Pattern</b> * Target: Handle Adapter: LocatorHandle Adaptee: Locator	Yes	No
<b>Command Pattern</b> Invoker: DrawApplication * Command: Command ConcreteCommand: InsertImageCommand Receiver: DrawingView	Yes	Yes
<b>Composite Pattern</b> Component: Figure * Composite: CompositeFigure Leaf: RectangleFigure	Yes	Yes
<b>Decorator Pattern</b> Component: Figure * Decorator: DecoratorFigure ConcreteComponent: RectangleFigure	Yes	Yes
<b>Observer Pattern</b> Subject: Drawing ConcreteSubject: StandardDrawing * Observer: DrawingChangeListener (DrawingView) ConcreteObserver: StandardDrawingView	Yes	Yes
<b>Prototype Pattern</b> * Client: CreationTool Prototype: Figure ConcretePrototype: AbstractFigure	Yes	No
<b>State<sub>1</sub> Pattern</b> Context: DrawApplication * State: Tool ConcreteState: AbstractTool	Yes	No
<b>State<sub>2</sub> Pattern</b> * Context: SelectionTool * State: Tool ConcreteState: AbstractTool	Yes	No
<b>Strategy Pattern</b> Context: LocatorHandle * Strategy: Locator ConcreteStrategy: AbstractLocator	Yes	No
<b>Template Method Pattern</b> * AbstractClass: AbstractFigure ConcreteClass: RectangleFigure	Yes	No

\* Class/es in which the design-pattern instances was/were textually documented in the source code