

The Impact of Domain Knowledge on the Effectiveness of Requirements Idea Generation during Requirements Elicitation

Ali Niknafs and Daniel M. Berry

David R. Cheriton School of Computer Science
University of Waterloo
Waterloo, Ontario, Canada

27 September 2012

Outline

- 1 Introduction
 - Study
- 2 Methodology
 - Pilot Studies
- 3 Controlled Experiment
 - Design
 - Results

Introduction, Definition of RE

The process of arriving at a specifications of a set of features that need to be developed is referred to as *requirements engineering (RE)*.

The Role of People in RE

- Of the three Ps, *process, product, and people*, in software engineering, people have been least scrutinized.
- Boehm observed that the quality of the development personnel is the most powerful factor in determining an organization's software productivity.
- While there is empirical evidence of the importance of the quality of the personnel in software development, there is not much in RE.

The Role of People in RE

The qualifications of the personnel involved in an RE process highly affects the effectiveness of the process, but most decisions about staffing RE teams arise from anecdotes and folklore, not from scientific studies.

The RE Gap

- One issue in RE is the gap between what the customer wants and what the analyst thinks the customer wants.
- To bridge this gap, many believe that an analyst needs to know the customer's problem domain well to do RE well for a system in the domain.
- However, deep knowledge of the problem domain can lead to falling into the *tacit assumption tarpit*.

Benefits of Domain Ignorance

The benefits of domain ignorance include:

- the ability to think out of the domain's box, leading to ideas that are independent of the domain assumptions,
- the ability to ask questions that expose the domain's tacit assumptions, leading to a common explicit understanding.

First Observations of Benefits of Ignorance

In 1994, Berry observed the benefits of domain ignorance when he performed better than expected when he helped specify requirements for software in domains he was quite ignorant of.

First Observations of Benefits of Ignorance

Probably, the earliest observation of the benefits of ignorance was Burkinshaw's statement during the 1969 Second NATO Conference on Software Engineering:

Get some intelligent ignoramus to read through your documentation and try the system; he will find many "holes" where essential information has been omitted. Unfortunately intelligent people don't stay ignorant too long, so ignorance becomes a rather precious resource. Suitable late entrants to the project are sometimes useful here.

Outline

- 1 Introduction
 - Study
- 2 Methodology
 - Pilot Studies
- 3 Controlled Experiment
 - Design
 - Results

Context of the Study

In each experiment, subjects perform an RE task that generates things, such as requirement ideas for some *computer-based system (CBS)* for some *client*.

- The RE task that is done in an experiment is called a *generative task (GT)*. Example GTs are requirements elicitation and requirements document inspection.
- The unit generated by a GT is called a *desired generated unit (DGU)*. For the two example GTs, the DGUs are requirements ideas and defects in a requirements document.

Context of the Study

- The CBS is situated in some *domain*, and at least one member of the client's organization is at least *aware of* and is often expert in this domain.
- Each member of the software development organization doing the RE activities has a different amount of *knowledge about the domain*. Each is either:
 - *Ignorant of the domain*, i.e., is a *domain ignorant (DI)*.
 - *Aware of the domain*, i.e., is a *domain aware (DA)*.
- Each of domain ignorance and domain awareness is a kind of *domain familiarity*.

Research Questions

Main Question

How does one form the most effective team, consisting of some mix of DIs and DAs, for a RE activity involving knowledge about the domain of the CBS whose requirements are being determined by the team?

Elaborated Questions

- Does a mix of DIs and DAs perform a RE activity more effectively than only DAs?
- Do other factors impact the effectiveness of an individual in performing an RE activity?

Research Questions

Main Question

How does one form the most effective team, consisting of some mix of DIs and DAs, for a RE activity involving knowledge about the domain of the CBS whose requirements are being determined by the team?

Elaborated Questions

- Does a mix of DIs and DAs perform a RE activity more effectively than only DAs?
- Do other factors impact the effectiveness of an individual in performing an RE activity?

Hypothesis

Main Hypothesis

A team consisting of a mix of DIs and DAs is more effective in an RE activity than is a team consisting of only DAs.

Null Hypothesis

The mix of DIs and DAs in a team has no effect on the team's effectiveness in an RE activity.

Hypothesis

Main Hypothesis

A team consisting of a mix of DIs and DAs is more effective in an RE activity than is a team consisting of only DAs.

Null Hypothesis

The mix of DIs and DAs in a team has no effect on the team's effectiveness in an RE activity.

Outline

- 1 Introduction
 - Study
- 2 Methodology
 - Pilot Studies
- 3 Controlled Experiment
 - Design
 - Results

Lessons Learned from Pilot Studies

- 1 Find a suitable problem domain.
- 2 Consider other factors (e.g. industrial experience) in analyzing the results.
- 3 Assess also the quality of the DGUs.
- 4 For many domains, so-called DIs turn out not to be real DIs, and so-called DAs turn out not to be real DAs.

Lessons Learned from Pilot Studies

Lessons 1 and 4 taught us that we need a problem domain that *partitions* the set of subjects with precision into

- DAs
- DIs

with no one in between.

We thought very hard to find such a domain, **bidirectional word processing**:

- CSers from the Middle East are DAs.
- CSers from elsewhere are DIs.

Outline

- 1 Introduction
 - Study
- 2 Methodology
 - Pilot Studies
- 3 **Controlled Experiment**
 - **Design**
 - Results

Experiment Context

- *GT*: The first, idea-generation step in a brainstorming activity to generate requirement ideas for a CBS.
- *DGUs*: Requirement ideas
- *Domain*: Bidirectional word processing
- *Subjects*: Volunteer subjects were recruited from a “Software Requirements and Specification” course and from outside the course, but nevertheless in CS or a related discipline.
- *Teams*:
 - **3I**: a team consisting of **3 DIs** and 0 DAs,
 - **2I**: a team consisting of **2 DIs** and 1 DAs,
 - **1I**: a team consisting of **1 DI** and 2 DAs,
 - **0I**: a team consisting of **0 DIs** and 3 DAs.

Variables

- Independent Variables about a team
 - *Mix of Domain Familiarities*
 - *Creativity Level*
 - *RE Experience*
 - *Industrial Experience*
- Dependent Variable
 - *Effectiveness*

Hypotheses

H_{11} : The effectiveness of a team in requirements idea generation is affected by the team's mix of domain familiarities.

H_{10} : The effectiveness of a team in requirements idea generation is not affected by the team's mix of domain familiarities.

H_{21} : The effectiveness of a team in requirements idea generation is affected by the team's creativity level.

H_{20} : The effectiveness of a team in requirements idea generation is not affected by the team's creativity level.

Hypotheses

H_{11} : The effectiveness of a team in requirements idea generation is affected by the team's mix of domain familiarities.

H_{10} : The effectiveness of a team in requirements idea generation is not affected by the team's mix of domain familiarities.

H_{21} : The effectiveness of a team in requirements idea generation is affected by the team's creativity level.

H_{20} : The effectiveness of a team in requirements idea generation is not affected by the team's creativity level.

Hypotheses

H_{31} : The effectiveness of a team in requirements idea generation is affected by **the team's RE experience**.

H_{30} : The effectiveness of a team in requirements idea generation is not affected by the team's RE experience.

H_{41} : The effectiveness of a team in requirements idea generation is affected by the team's industrial experience.

H_{40} : The effectiveness of a team in requirements idea generation is not affected by the team's industrial experience.

Hypotheses

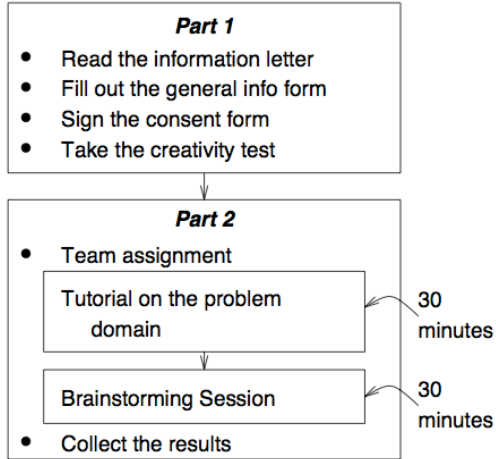
H_{31} : The effectiveness of a team in requirements idea generation is affected by **the team's RE experience.**

H_{30} : The effectiveness of a team in requirements idea generation is not affected by the team's RE experience.

H_{41} : The effectiveness of a team in requirements idea generation is affected by **the team's industrial experience.**

H_{40} : The effectiveness of a team in requirements idea generation is not affected by the team's industrial experience.

Procedure



Evaluation of Generated Ideas

- The **quantitative data** is the **number of raw ideas** generated by each team, which is a good measure for the GT = brainstorming (because quantity is the *goal* of the first stage of brainstorming).
- To better compare the performance of the teams, Niknafs considered **also** the **quality** of their generated **ideas**.

Quality of Generated Ideas

Based on the characteristics of a good requirement in the **IEEE 830 Standard**, each **idea is classified** according to three characteristics:

- 1 **Relevancy**: an idea is considered relevant if it has something to do with the domain.
- 2 **Feasibility**: an idea is considered feasible if it is relevant and it is correct, well presented, and implementable.
- 3 **Innovation**: an idea is considered innovative if it is feasible and it is not already implemented in an existing application for the domain known to the evaluator.

Evaluation of Quality of Generated Ideas

- Berry and Niknafs evaluated the quality of the ideas since we were both experts in bidirectional word processing.
- To eliminate any bias in classifying an idea that might arise from the evaluator's knowing the domain familiarity mix of the team from which the idea came, Niknafs produced a list of all ideas generated by all teams, sorted using the first letters of each idea.
- Each domain-expert evaluator classified the ideas in the full list.
- After both evaluations were done, the each evaluator's classifications of each idea were transferred to the idea's occurrences in the individual team lists.

Outline


- 1 Introduction
 - Study
- 2 Methodology
 - Pilot Studies
- 3 **Controlled Experiment**
 - Design
 - **Results**

Results: Data About the Teams

<i>Type of Teams</i>	<i>Number of Teams</i>	<i>Creativity</i>	<i>RE Experi- ence</i>	<i>Industrial Experience</i>
		Mean	Mean	Mean
<i>3I</i>	9	69.11	0.89	3.06
<i>2I</i>	4	71.75	0.75	3.33
<i>1I</i>	3	70.67	1.00	1.33
<i>0I</i>	3	71.33	1.00	2.00

Outliers

- Boxplots were used to graphically expose any outliers.



ANOVA Prerequisites

- The differences between the teams were determined by means of an analysis of variance (ANOVA).
- In order to be allowed to apply an ANOVA, the data must meet the three prerequisites for an ANOVA:
 - 1 *All dependent variables are normally distributed.*
 - 2 *All variances are homogeneous.*
 - 3 *All observations are independent.*

ANOVA Prerequisites

- An ANOVA was applied to the dependent variables whose values met the prerequisites for an ANOVA; i.e. the numbers of generated raw, relevant, and feasible ideas.
- For innovative ideas, another, non-parametric test was used.

ANOVA Results

Effect	Raw Ideas				Relevant Ideas				Feasible Ideas			
	<i>F</i>	<i>p</i>	<i>f</i> ²	<i>P</i>	<i>F</i>	<i>p</i>	<i>f</i> ²	<i>P</i>	<i>F</i>	<i>p</i>	<i>f</i> ²	<i>P</i>
Mix of Domain Familiarities	.165	.915	.011	.068	8.675	.032	.319	.816	13.486	.015	.449	.941
Creativity	.921	.469	.048	.146	3.918	.114	.159	.459	.984	.449	.051	.153
Industrial Experience	.563	.609	.031	.107	10.089	.027	.331	.833	4.381	.098	.173	.499
RE Experience	.145	.722	.008	.063	.173	.699	.009	.65	.035	.861	.002	.53

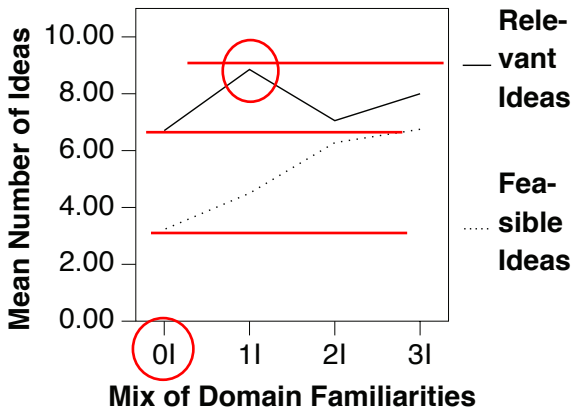
F is *F*-test; *p* is *p*-value of *F*-test; *f*² is Cohen effect size; *P* is post-hoc power.

Focused ANOVA Results

Effect	Relevant Ideas		Feasible Ideas	
	p	P	p	P
Mix of Domain Familiarities	.032	.816	.015	.941
Industrial Experience	.027	.833	.098	.499

p is p -value of F -test; P is post-hoc power.

ANOVA Results: Impact of Domain Knowledge



ANOVA Results: Impact of Industrial Experience



ANOVA Results: Non-Parametric Test on Innovative Ideas

<i>Effect</i>	<i>Kruskal-Wallis</i>
	Significance
Mix of Domain Familiarities	.966
Creativity	.996
Industrial Experience	.240
RE Experience	.749

Threats to **Validity**

- **Conclusion Validity:** Low Statistical Power: 20 teams would be enough to achieve statistical power of 0.80, but, the unequal number of teams in the mixes reduces statistical power.
- **Internal Validity:** Voluntary Subjects: All subjects were voluntary but were randomized to the extent possible while still getting the necessary mixes of domain familiarities among the teams.

Threats to Validity

- **Construct Validity:** Confounding Constructs: Sometimes the value of an independent variable affects the results more than the presence or absence of the variable would.
- **External Validity:** Population Validity: The experiment used student subjects instead of professional analysts, although the students are mostly co-op and work one term per year.

Conclusion About Hypotheses

- Hypothesis H_{11} is strongly accepted:
The effectiveness of a team in requirements idea generation is affected by the team's mix of domain familiarities.
- Hypothesis H_{20} is weakly accepted:
The effectiveness of a team in requirements idea generation is not affected by the team's creativity level.

Conclusion About Hypotheses

- Hypothesis H_{30} is accepted:
The effectiveness of a team in requirements idea generation **is not** affected by the team's RE experience.
- Hypothesis H_{41} is accepted:
The effectiveness of a team in requirements idea generation **is** affected by the team's industrial experience.

Main Result

From these results, considering the threats, the main hypothesis, that

A team consisting of mix of DIs and DAs is more effective in requirements idea generation than a team consisting of only DAs,

appears to be weakly supported.

Expected Application of the Results

Help RE managers in forming teams that are performing knowledge-intensive RE activities, by

- providing a list of RE activities for which domain ignorance is at least helpful and
- providing advice on the best mix of DIs and DAs for any RE activity.

Now!

If we have piqued your interest, then go read the paper for the full details that we did not have time to present here!

But please wait until the end of the session, because the other speakers deserve your attention too!

Enjoy!