

Lecture Notes: Requirements

17-313: Foundations of Software Engineering
Claire Le Goues and Michael Hilston

Requirements play a major role in sinking software projects. Building the wrong thing, failing to build the right thing, or building a thing that would be the right thing under different circumstances, are all common contributors to real-world project failures. *Requirements engineering* is the process of capturing in symbolic representation a worldly computational problem so that the statements obtained by following rules of symbolic manipulation are useful statements once translated back into the language of the world [Scherlis 1989].

In plainer English: requirements engineering seeks to (1) figure out what system should be built, and then (2) express those ideas such that the correct thing can be verifiably built. The first is a question of knowledge *acquisition*, addressing how to capture all relevant detail about a proposed system and then determine whether that knowledge is complete and consistent. The second is a question of knowledge *representation*: Once captured, how do we express it most effectively, such that it is received consistently by all the different people who will interact with it?

(We also need to know how to check requirements for correctness, consistency, and implementability; how to check that they've been represented faithfully; and how to maintain requirements as they change and throughout the implementation process. We focus in this course primarily on elicitation and representation.)

A challenge is that the word "requirements" is used inconsistently throughout the software industry. It can mean everything from the vague idea your cofounder has about an awesome new app to a detailed and lengthy Software Requirements Specification (SRS) used to guide bidding and contracting for a large IT acquisition. In these notes, we drill down into formal requirements engineering to some degree of detail, because in so doing we will better understand the challenges of requirements and the techniques and terminology that can be used to overcome some of those challenges. We won't write a full SRS document, however!

Requirement types and their expression

We divide requirements into three broad categories: (1) functional requirements, (2) quality requirements,¹ and (3) domain assumptions.²

Functional requirements and implementation bias

The functional requirements describe what the machine should do. They cover the inputs and expected outputs of the system, its interface with the world, and how it will respond to events.

¹The literature often uses the term "non-functional requirements", the phrase is misleading; using either term on exams or homework is fine.

²Some writers also contrast between *user* (human-ese) and *system* (computer-ese) requirements, but this distinction is less important for our purposes.




Figure 1: The problem world, the machine solution, and the interface between them (adapted from [van Lamswerde 2009])

A common one-sentence summary of requirements is: “Requirements describe what the software should do, and not how it should do it.” A key goal in requirements engineering is to avoid *implementation bias*: describing the desired system in terms of its implementation. This constrains the software engineers’ choices and may prevent them from exploring alternative (but possibly better) solutions.

For these reasons, functional requirements should be expressed in terms of the relationship between the system to be built (the *machine*) and the *environment* in which it will operate. The goal of software engineering can thus be broadly reformulated as seeking to improve the world by putting some machine into it that solves the problem at hand. Requirements then describe what is observable at the environment-machine interface, and should be expressed as statements about the effect the machine has on the environment (Figure 1).

Thus, requirements are largely expressed using two types of grammatical statements:

- Statements written in the *indicative mood* describes the environment (as-is, independent of the system). Example: “Every product has a unique product code.” We will discuss these statements in the context of the domain knowledge and assumptions, below.
- Statements written in the *optative mood* describes the environment with the system to be built. Example: “The system shall email clients about their shipping status.”

This avoids implementation bias because the requirements don’t actually make statements about the system being built. Instead, they describe the effect the system has on the world.³

Quality requirements

The quality requirements specify the quality with which a system should deliver its functionality (examples include but are not limited to performance, throughput, reliability, availability, usability, etc). Beyond desired functionality, quality requirements serve as constraints on the implemented system and provide design criteria to select between alternative implementations.

Good requirements, of any kind, are verifiable. They serve as contracts, and thus should be testable or falsifiable. We therefore contrast *informal goals*, or general intentions (such as “ease of use”, or even “the system should be easy to use by experienced controllers, and should be organized such that user errors are minimized.”) from *verifiable non-functional requirements*, or statements that use measures that can be objectively tested (such as “Experienced controllers shall be able to use all the system functions after a total of two hours training. After this training, the average number of errors made by experienced users shall not exceed two per day, on average.”)

³Of course, even following these principles, requirements can still be flawed, e.g., if they make invalid statements about the world, if they are vague, etc, but they cannot overconstrain the implementation.

Informal goals may still be informative for developers, as they convey the intentions of the system users, but verifiable quality requirements are key to effective requirements expression.

Domain knowledge

Assumptions or domain knowledge document existing properties of the environment. Such assumptions serve two general purposes: (1) constrain the *scope* of the system by describing the conditions under which it is intended to perform correctly, and (2) define the relationship/interface between the machine model of the world and the actual world.

The first goal is straightforward enough: by documenting assumptions about the environment as it is, we specify (and limit!) the conditions under which the system is expected to perform correctly. Additionally, specifying a clear boundary between machine actions and environment actions clarifies which actions are actually the responsibility of the desired system. For example, consider a system like Netflix's that suggests movies or TV shows to users. The scope of that system would vary tremendously based on whether it is able to query an existing system (within its environment) that stores and manages user watch logs and information, or if, alternatively, it is responsible for storing and managing that information itself.

Regarding the second goal, domain assumptions ground the definitions used in requirements in reality (providing *designations*). This helps ensure functional requirements are correctly defined and implementable. Consider the following example, borrowed from [Zave and Jackson, 1997]:

Ash: Two important basic types are *student* and *course*. There is also a binary relation *enrolled*. It can be expressed as: $\forall s \forall c (\text{enrolled}(s, c) \implies \text{student}(s) \wedge \text{course}(c))$

Brook: Do only students enroll in courses? I don't think that's true.

Ash: But that's what I mean by student!

If a student is simply a person enrolled in a course, Ash's assertion is vacuous, but easy to implement. If a student is a person who has matriculated at the university and not yet graduated or withdrawn, Ash's assertion strongly constrains who can take courses. The difference is important to correct system implementation.

The problem here is that a student is just a concept that the machine will express and manipulate symbolically. A machine doesn't *know* what a student is. We need a designation for the word "student" to check that a specification derived from these requirements faithfully reflects stakeholder desires.

More formally, the meaning of "being a student," is *unshared* knowledge: The machine cannot observe it. This is true of most real-world phenomena. Other actions are observable by both the world and the machine: This is *shared* knowledge. Figure 1 shows an example of the distinction: in a payment system, the "payment record created in database" is an action observable by the machine; both the machine and the world can see if the payment notification was sent to the seller; but the machine cannot possibly know whether the item is actually physically delivered, because this is a property of the physical world. The machine cannot be put in charge of unshared actions or phenomena, and requiring them makes for requirements that are impossible to implement. Domain knowledge and assumptions therefore provide definitions of unshared information, and relate them to shared information that the machine can actually observe and affect.

Eliciting requirements

Stakeholders. We elicit requirements primarily by engaging *stakeholders*, people or groups who will be affected by the system, directly or indirectly. The first step in elicitation is typically *stake-*

holder analysis, to identify relevant stakeholders. That is: who is the system for? Who will decide whether it is acceptable? Missing or hidden stakeholders pose a key risk in requirements elicitation because the system must, by definition, satisfy stakeholder goals. If you miss someone, you risk not missing an important goal of the system.

We distinguish between stakeholders and the *actors* (or *agents*). An actor is an entity that interacts with the system for the purpose of completing an event [Jacobson, 1992] (e.g., a user, organization, or even another system). The actors are not as broad as the stakeholders.

Interviews. The most common way to engage stakeholders to draw out system requirements is via *interviews*. The goals of these interviews are to understand functional requirements, identify and learn domain-specific concepts, and prioritize quality attributes. There are many potential challenges to effective interviews: Stakeholders may disagree; Unhandled or hidden conflicts are another key risk. Stakeholders do not always know what they really want nor how to articulate it, nor how hard it is. They often have domain knowledge, use jargon, or leave out “obvious” requirements that aren’t obvious to a non-expert. Effectively engaging stakeholders can be difficult regardless of their technical sophistication.

Effective interviewing typically begins concretely, with specific questions, proposals, or by working through a mockup or prototype. Humans are typically much better at recognizing correct or incorrect solutions than they are at solving a problem from nothing. Stories, storyboards, scenarios, informal or hypothetical use cases, or examples can be helpful, or contrasting the proposed system with a current system (if applicable). Eliciting requirements at the boundaries of the system—corner, negative, or failure cases—can also help with scope and exceptional situations. We will see examples of these types of tools in class and in the slides.

Cited references

1. William L. Scherlis, responding to E W Dijkstra “On the Cruelty of Really Teaching Computing Science”. *Communications of the ACM* 32:12, p407.
2. Pamela Zave and Michael Jackson. “Four dark corners of requirements engineering.” *ACM Trans. Softw. Eng. Methodol.* 6, 1 (January 1997), 1-30.
3. Axel van Lamsweerde: *Requirements Engineering - From System Goals to UML Models to Software Specifications*. Wiley 2009.
4. Ivar Jacobson, Magnus Christerson, Patrik Jonsson, Gunnar Övergaard: *Object-oriented software engineering - a use case driven approach*. Addison-Wesley 1992, ISBN 978-0-201-54435-0.