



Cognition in Software Engineering: A Taxonomy and Survey of a Half-Century of Research

Downloaded from: <https://research.chalmers.se>, 2026-04-04 12:57 UTC

Citation for the original published paper (version of record):

Fagerholm, F., Felderer, M., Fucci, D. et al (2022). Cognition in Software Engineering: A Taxonomy and Survey of a Half-Century of Research. *ACM Computing Surveys*, 54(11).
<http://dx.doi.org/10.1145/3508359>

N.B. When citing this work, cite the original published paper.

Cognition in Software Engineering: A Taxonomy and Survey of a Half-Century of Research

FABIAN FAGERHOLM, Aalto University, Finland and Blekinge Institute of Technology, Sweden

MICHAEL FELDERER, University of Innsbruck, Austria and Blekinge Institute of Technology, Sweden

DAVIDE FUCCI and MICHAEL UNTERKALMSTEINER, Blekinge Institute of Technology, Sweden

BOGDAN MARCULESCU, Kristiania University College, Norway

MARKUS MARTINI, University of Innsbruck, Austria

LARS GÖRAN WALLGREN TENGBERG, University of Gothenburg, Sweden

ROBERT FELDT, Chalmers University of Technology, Sweden and Blekinge Institute of Technology, Sweden

BETTINA LEHTELÄ, BALÁZS NAGYVÁRADI, and JEHAN KHATTAK, Aalto University, Finland

Cognition plays a fundamental role in most software engineering activities. This article provides a taxonomy of cognitive concepts and a survey of the literature since the beginning of the Software Engineering discipline. The taxonomy comprises the top-level concepts of perception, attention, memory, cognitive load, reasoning, cognitive biases, knowledge, social cognition, cognitive control, and errors, and procedures to assess them both qualitatively and quantitatively. The taxonomy provides a useful tool to filter existing studies, classify new studies, and support researchers in getting familiar with a (sub) area. In the literature survey, we systematically collected and analysed 311 scientific papers spanning five decades and classified them using the cognitive concepts from the taxonomy. Our analysis shows that the most developed areas of research correspond to the four life-cycle stages, software requirements, design, construction, and maintenance. Most research is quantitative and focuses on knowledge, cognitive load, memory, and reasoning. Overall, the state of the art appears fragmented when viewed from the perspective of cognition. There is a lack of use of cognitive concepts that would represent a coherent picture of the cognitive processes active in specific tasks. Accordingly, we discuss the research gap in each cognitive concept and provide recommendations for future research.

Authors' addresses: F. Fagerholm, Aalto University, P.O. Box 15400, Espoo, FI-00076, Finland, Blekinge Institute of Technology, Valhallavägen 1, Karlskrona, 37141, Sweden; email: fabian.fagerholm@aalto.fi; M. Felderer, University of Innsbruck, Technikerstraße 21a, Innsbruck, 6020, Austria, Blekinge Institute of Technology, Valhallavägen 1, Karlskrona, 37141, Sweden; email: michael.felderer@uibk.ac.at; D. Fucci, Blekinge Institute of Technology, Valhallavägen 1, Karlskrona, 37141, Sweden; email: davide.fucci@bth.se; M. Unterkalmsteiner, Blekinge Institute of Technology, Valhallavägen 1, Karlskrona, 37141, Sweden; email: michael.unterkalmsteiner@bth.se; B. Marculescu, Kristiania University College, PB 1190 Sentrum, Oslo, 0107, Norway; email: bogdan.marculescu@kristiania.no; M. Martini, University of Innsbruck, Innrain 52, Innsbruck, 6020, Austria; email: markus.martini@uibk.ac.at; L. Göran Wallgren Tengberg, University of Gothenburg, Box 100, Gothenburg, 40530, Sweden; email: larsgoran.wallgren@psy.gu.se; R. Feldt, Chalmers University of Technology, Chalmersplatsen 4, Gothenburg, 41296, Sweden, Blekinge Institute of Technology, Valhallavägen 1, Karlskrona, 37141, Sweden; email: robert.feldt@chalmers.se; B. Lehtelä, Aalto University, P.O. Box 15400, Espoo, FI-00076, Finland; email: bettina.lehtela@aalto.fi; B. Nagyváradi, Aalto University, P.O. Box 15400, Espoo, FI-00076, Finland; email: balazs.nagyvaradi@aalto.fi; J. Khattak, Aalto University, P.O. Box 15400, Espoo, FI-00076, Finland; email: jehan.khattak@aalto.fi.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2022 Copyright held by the owner/author(s).

0360-0300/2022/09-ART226

<https://doi.org/10.1145/3508359>

CCS Concepts: • **General and reference** → **Surveys and overviews**; • **Software and its engineering** → **Software creation and management**;

Additional Key Words and Phrases: Cognition, cognitive concepts, psychology of programming, human factors, measurement, taxonomy

ACM Reference format:

Fabian Fagerholm, Michael Felderer, Davide Fucci, Michael Unterkalmsteiner, Bogdan Marculescu, Markus Martini, Lars Göran Wallgren Tengberg, Robert Feldt, Bettina Lehtelä, Balázs Nagyvárad, and Jehan Khattak. 2022. Cognition in Software Engineering: A Taxonomy and Survey of a Half-Century of Research. *ACM Comput. Surv.* 54, 11s, Article 226 (September 2022), 36 pages.
<https://doi.org/10.1145/3508359>

1 INTRODUCTION

Cognition is the “*collection of mental processes and activities used in perceiving, remembering, thinking, and understanding, and the act of using those processes*” [4]. These processes include the formation of knowledge, creation and use of memory, problem solving, decision making, attention, and comprehension. Studying the quality of mental processes and how they can be affected is of general interest in psychology, but is also relevant for knowledge-intensive fields such as **Software Engineering (SE)**. From an SE perspective, it is interesting to describe and measure (quantitatively and qualitatively) cognitive concepts, since this enables the systematic evaluation of interventions aimed at improving mental processes related to an engineering task, and subsequently the performance in the task.

Researchers have studied behavioural aspects of SE since the dawn of the discipline [16]. The question of how the human mind processes information related to software development activities has been of interest of SE research [15, 33, 39, 56]. Understanding cognitive processes in a quantifiable manner is a powerful lever when engineering effective processes and tools that support software development activities. Therefore, to advance SE research and understand which cognitive concepts have seen attention in combination with particular software development activities, a systematic categorisation of SE research using a taxonomy of cognitive concepts is valuable. While there are taxonomies that cover aspects of cognition in SE (e.g., human errors [3], cognitive biases [35], and cognitive load [27]), there is not, to the best of our knowledge, taxonomy covering the wide spectrum of concepts that have been studied in the general field of psychology.

This article provides (i) an overview of the vast and fragmented research on cognition in SE, (ii) definitions of central cognitive concepts, processes, and activities, and (iii) approaches to assess them. We propose a taxonomy of cognitive concepts and assessment procedures relevant for SE. We then apply this taxonomy to characterise SE research focusing on cognitive aspects—with articles published between 1973¹ and 2020.

The contribution of this study is threefold, providing:

- (1) A taxonomy of cognitive concepts and a framework to measure and assess these concepts, developed in collaboration with experts from cognitive and organisational psychology.
- (2) An overview of the research in SE, characterised by which cognitive aspects were studied and how they were measured and assessed.
- (3) A discussion on which knowledge gaps exist with regard to cognitive aspects of SE.

This paper focuses on cognitive concepts in professional software development. We do not take into account aspects dedicated to (i) pedagogical, educational, and learning activities [20], as well

¹This date does not indicate a particular event but the publication date of the earliest paper we classified.

as (ii) knowledge management and software design decision [10, 42], which also have cognitive aspects.

The remainder of this paper is structured as follows. In Section 2, we provide definitions of cognition, present prior reviews of the literature related to cognition in SE, argue for the need for the present review, discuss taxonomies in SE, and set out our objectives related to the taxonomy. In Section 3, we describe the research approach and methods used in this article. In particular, we describe the taxonomy development and literature search, selection, and analysis, and discuss the threats to validity of our method choices. In Section 4, we present the taxonomy of cognitive concepts and assessment procedures. We describe the application of the taxonomy and the results of the literature study in Section 5. In Section 6, we discuss the study results and conclude the paper in Section 7.

2 PRELIMINARIES AND RELATED WORK

From technical details to the broadest perspectives on design and requirements, making software involves people thinking. Before software runs on our computing devices, it is imagined in the minds of software professionals—those involved in product design, project management, requirements, architecture, user experience, testing, programming, and other development and design activities.

Researchers have investigated thinking among software professionals for decades (e.g., [18, 19, 34, 51, 55, 59]), providing guidance, methods, and tools for those involved in the development endeavour. Some early studies concerned programming—reading and understanding source code—with implications for variable naming, programming language design, finding, and preventing programming errors. In the past decades, various topics related to human cognition in SE have been investigated.

In this section, we first discuss cognition as a psychological concept to demarcate the focus of this article. We then provide an overview of past work that, similarly to our aim, reviewed studies that pay particular attention to cognitive aspects of SE. We first look at historic reviews that suggest that cognitive aspects were studied from the inception of the SE field. Then, we summarise the findings of contemporary reviews in SE in general and in specialised sub-fields. Accordingly, we provide a motivation as to why our review is necessary and how it contributes to the existing body of knowledge. Finally, we discuss the concept of a taxonomy, how taxonomies are used in SE, and outline the goals of the taxonomy presented in this article.

2.1 What is Cognition?

Cognition is an umbrella term for a host of different conscious and unconscious mental processes. It refers to the acquisition and processing of information [11], and spans processes of perception, attention, memory, thinking, and problem-solving [4]. For example, during programming, we generate, maintain and manipulate a representation of source code in our working memory based on our (expert) knowledge of different programming languages which is stored in our long-term memory. We can direct our focus of attention to specific parts of the code or think about and mentally simulate a solution to a problem. Programming integrates multiple, dynamic, and interacting mental processes which can partly run in parallel. The same is true for other software design and development tasks, such as working with requirements, architectures, and testing, dealing with project management, providing effort estimates, and executing specific development processes and techniques.

Cognition can be conceptually separated from emotion and motivation, although they arise from interacting neural circuits and processes [40, 45]. The human cognitive system and cognitive behaviour, such as decision-making, is also affected by whether an activity occurs for a single individual or in a group [9]. Furthermore, cognition can also be conceptualised from a sociological

perspective that considers interactions between humans and with culturally significant objects to constitute cognitive processes. This is the case in distributed cognition theory, according to which cognition is not only individual, but is embedded and situated in external artefacts, groups, and cultural systems for interpreting reality [29].

In this article, we focus on individual information processing and how it has been studied in SE. We have chosen to put less emphasis on emotion and motivation, and to include only a small subset of cognition in social situations in which the focus is mainly social-psychological and concentrates on the individual. Prior reviews have given overviews of the research on motivation [6, 24] and emotion [46] in SE.

Cognition can be studied using various means, ranging from psycho-physiological measurements to open observations of behaviour. Different assessment procedures serve different purposes—from detailed information on, for example, perceptual processes, to uses as part of process improvement or as input to software development tools. We have chosen to include all kinds of assessment procedures that we have found in the literature, to illustrate the wide variety of available approaches.

2.2 Historic Reviews

In the early 1980s, Sheil [49] reviewed work on the psychology of programming published since the late 1960s, when the field of SE emerged as an independent research area. The review focuses on programming notation (conditionals, control flow, data types), practices (flow-charting, indenting, variable naming, commenting) and tasks (learning to program, coding, debugging), and their impact on developer performance. Sheil criticises the methodological flaws in how these phenomena have been studied, in particular the design of experimental treatments, practice effects (learning), failure to compensate for individual variability, and not considering effect size in statistical tests.

Curtis [16] follows in 1984 by reviewing studies on individual differences of programmers in terms of expertise, problem solving, fault detection, and learning to program. He points to the lack of research, in the reviewed time-period from 1968 to 1984, on problem-solving processes in requirements engineering and software design. With the popularisation of object-oriented programming in the 1980s and 1990s, more empirical studies on software design and its consequences on cognition were performed, as reviewed by Détienne [19].

These early reviews culminate in 1986 in a call for collaboration between psychologists and computer scientists [17]. In this review, Curtis et al. summarise studies on problem solving in unstructured domains, use of programming languages, knowledge representation for programmers, variable naming, and research on the parallels between text and source code understanding. Furthermore, the review discusses methodological approaches (e.g., experiments) and issues in designing and conducting them, and points towards the need for research in conceptual and mental modes, psychological factors affecting programming, transfer of skills, communication media between programmers, and problem solving representations.

2.3 Contemporary Reviews

Here, we discuss reviews pertinent to cognitive aspects that have been published between 2014 and 2019. We begin by looking at four reviews that studied the SE discipline as a whole and two reviews that focused on software project management.

Lenberg et al. [33] conducted a **systematic literature review (SLR)** with the purpose of identifying concepts relevant for the proposed research area of **Behavioural SE (BSE)**. They identified 55 concepts, cognition being one of them, within the area of BSE. Between 1997 and 2013, the most studied BSE concepts were communication, personality, and job satisfaction.

Personality in SE research was the focus of a systematic mapping study by Cruz et al. [13]. They found that, in the reviewed time-frame from 1970 until 2010, personality was studied mostly

in the context of pair-programming, but also in education, team effectiveness, software-process allocation, software engineer personality characteristics, individual performance, team process, behaviour and preferences, and leadership performance. They point out that even after 40 years of research, the field does not yet seem mature. Conflicting evidence indicates that the results presented in the studies cannot be directly applied, as they may not produce the desired results.

Sharafi et al. [48] conducted an SLR of studies utilising eye-tracking to record visual attention as a proxy for cognitive processes. They found 36 studies investigating the engineers performing tasks pertinent to model comprehension, code comprehension, debugging, collaborative interactions (e.g., studying eye-movement as cues in collaborative tasks), and traceability recovery.

Blackwell et al. [7] reviewed the literature pertaining to the psychology of programming, published in the *Journal of Man-Machine Studies* and in the *Journal of Human-Computer Studies* between 1969 and 2019. They describe the evolution of the psychology of programming field during the five decades and observe a change in research focus: in the early days, from individual developers and small scale programs (1970s), through large code bases and programs where professional skills and processes become more important (1980s and 1990s), to contemporary issues where software development is a social activity that is conducted in different contexts (2000s and 2010s).

Advances and novel research directions in psychology often find their way into SE research and suggest new ways of looking at different phenomena in the field. For example, the research program on cognitive biases received mainstream attention in 2002 with a Prize in Economic Sciences in Memory of Alfred Nobel (c.f., Kahneman [30]), likely influencing research in SE as well. In 2014, Shepperd [50] discusses how cost estimations are affected by the planning fallacy, a preference for case-specific over distributional evidence, the peak-end rule, and the anchoring effect. In 2016, Cunha et al. [15] briefly review how cognitive biases have been studied in the context of decision-making in software project management.

2.4 Need for a Review on Cognitive Concepts in SE

Looking at the reviews discussed above, both historic and contemporary ones, we observe that cognition has either (1) not been in the focus of the investigation (e.g., behavioural SE by Lenberg et al. [33] or personality by Cruz [13]), (2) was studied in a narrow area of SE, such as software construction, design and testing (e.g., Blackwell et al. [7], which is also limited to two journals, or Sheil [49] and Curtis et al. [16, 17], which were also done more than 30 years ago) or software project management [15, 50], or (3) focused on a particular method to study cognitive processes (e.g., eye-tracking by Sharafi et al. [48]).

Reviews of cognition in SE have the double challenge of capturing both the breadth of topics in SE, and the breadth of cognition as a psychological concept. Curtis [16] observed already 25 years ago that the research on behavioural aspects of SE has not been organised into a sub-field of SE or psychology. He attributes this to the difficulty of integrating the mixture of paradigms borrowed from psychology. While the **SE Body of Knowledge (SWEBOK)** [8] recognises group dynamics and psychology as an aspect of SE professional practice, topics such as individual cognition and problem solving are merely touched upon. Existing taxonomies in SE do not comprehensively cover people issues in general [57], and thus neither cognition in particular.

To the best of our knowledge, no comprehensive taxonomy of cognitive concepts relevant to SE has yet been proposed and used to characterise the field. With the taxonomy and review presented in this paper, we fill this gap and complement previous reviews.

2.5 Taxonomies in SE

Classifications in general, and taxonomies in particular, provide a common terminology for knowledge sharing, a better understanding of the relationships between concepts in particular fields, and

support in decision making and identifying gaps in knowledge [57]. A key benefit of taxonomies, and of classifications in general, is to “increase cognitive efficiency and facilitate inferences,” as well as facilitating education [41]. An example of the application of a taxonomy to SE research is the SE Body of Knowledge (SWEBOK) [8] that provides an overview of the knowledge areas that are relevant to the field.

Taxonomies provide researchers with the ability to reason about more generally-applicable classes of a domain, rather than restricting them to individual instances [41]. By grouping concepts and their relationships based on some assessment of similarity, researchers can establish useful inferences, including counter-intuitive causal or temporal relationships based on data. Ralph [41] discusses dimensions for analysing and assessing taxonomies: (1) the degree to which the class structure reflects similarities and dissimilarities observed between actual instances, (2) the ability to draw useful conclusions about instances based on their membership in relevant classes, and (3) the degree to which the taxonomy meets its goal.

Ralph also discusses the philosophical grounding for taxonomies and their use [41]. The author notes that teleological, dialectic, and life-cycle process theories include taxonomies and are most suited to being described by such classifications. In addition, the paper provides examples of the cognitive efficiency gained by using taxonomies, and how it can be further used to draw more general conclusions from the available data [41].

We define a taxonomy of psychological concepts used in SE in accordance with the guidelines provided in Usman et al. [57]. In particular, we focus on the definition of a common terminology to improve research communication and education, and on identifying knowledge gaps (c.f. [41]). We first define the taxonomy based on domain knowledge from SE and grounded in psychology literature. Then, we use it as a framework to investigate the extent to which cognitive concepts have been researched in SE, and to identify gaps in the current knowledge.

3 RESEARCH METHODOLOGY

Based on the literature, we observed that a rather large body of work exists at the overlap between SE and psychology, in particular the study of human cognition and SE activities. Accordingly, we set out to systematically create a taxonomy of cognitive concepts and assessment procedures, and to classify the existing scientific literature to understand development, gaps, and future directions of the field.

Accordingly, this study is driven by the following research questions:

RQ1: How can cognitive concepts relevant for SE be classified and assessed?

RQ2: What is the state of the art of cognitive concepts in SE?

RQ2.1: In which context have cognitive concepts been studied in SE?

RQ2.2: How has the study of cognitive concepts in SE developed over time?

RQ2.3: Which SE knowledge areas have been studied under the lens of cognitive concepts?

RQ2.4: How have cognitive concepts in SE been assessed?

RQ2.5: Which theories have been used to study cognitive concepts in SE?

RQ3: What gaps and potential future directions exist in the field of cognition in SE?

We assembled an interdisciplinary team of researchers to work on answering the research questions. Table 1 lists the researchers, their main discipline and the tasks they were involved in in this research. The choice of tasks was first and foremost driven by expertise, personal preference and, lastly, time availability. We refer to the researchers’ acronyms when describing the tasks later in this section. We conducted the research in two phases. During the 1st phase, we executed the initial search that included primary studies until 2017. We conducted the 2nd phase in 2020 to include primary studies published after 2017.

Table 1. Researchers Involved in This Work, Their Main Disciplines, and Tasks

1st Phase			
Name	Acronym	Main discipline	Tasks
Fabian Fagerholm	FFA	SE	taxonomy construction, paper selection, data extraction, validation
Michael Felderer	MFE	SE	taxonomy construction, paper selection, data extraction, validation
Robert Feldt	RFE	SE	paper selection, data extraction, validation
Davide Fucci	DFU	SE	validation
Michael Unterkalmsteiner	MUN	SE	paper selection, data extraction, validation, classification
Bogdan Marculescu	BMA	SE	data extraction, validation
Markus Martini	MMA	Psychology	taxonomy construction
Lars-Göran Wallgren Tengberg	LTE	Psychology	taxonomy construction
2nd Phase			
Name	Acronym	Main discipline	Tasks
Fabian Fagerholm	FFA	SE	paper selection, data extraction, validation, classification
Bettina Lehtelä	BLE	SE	paper selection, data extraction, validation, classification
Balázs Nagyvárad	BNA	SE	paper selection, data extraction, validation, classification
Jehan Khattak	JKH	Human-Computer Interaction	paper selection, data extraction, validation, classification

Table 2. Included and Excluded Primary Studies During the Activities of the 1st and 2nd Phase of the Literature Survey

Activity	1st Phase		2nd Phase	
	Included	Excluded	Included	Excluded
Keyword search	738	–	1030	–
Selection	208	530	96	934
Data extraction	129	79	96	0
Backward snowball sampling	113	0	–	–
Classification	217	25	96	0
Total included primary studies				311

To answer RQ1, we invited experts in cognitive and behavioural psychology to support the taxonomy development, (see Section 3.1). To answer RQ2, we conducted a literature survey driven by keyword search and backward snowball sampling [61] as described in Sections 3.2–3.7. The quantitative results of included and excluded primary studies are summarised in Table 2. Finally, to answer RQ3, we classified the main research contributions of the identified relevant studies using the developed taxonomy on cognitive concepts and assessment procedures, and the knowledge areas presented in SWEBOK, as described in Section 3.6. Threats to the validity of our research methodology are discussed in Section 3.8.

3.1 Taxonomy Development: Cognitive Concepts and Assessment Procedures

To classify existing and future work on cognitive research in SE, FFA and MFE developed, together with two experts (MMA, LTE), a taxonomy of cognitive concepts and assessment procedures. Based on searches of the literature, and according to the best of our knowledge as well as the expertise of the involved psychologists, no suitable taxonomy of cognitive concepts existed. Taxonomies do exist in specific areas of psychology, such as personality and developmental psychology, as well as certain specific topical areas within organisational and cognitive psychology. However, these are too specific to cover the range of cognitive concepts investigated in SE. Since we do not aim to create a taxonomy for use in psychology, we decided to develop a new taxonomy for SE, based on concepts from psychology. The steps of the taxonomy development process are summarised in Table 3 and explained next.

In step 1, the two experts (MMA, LTE) independently created two taxonomies each, one for cognitive concepts and one for cognitive assessment procedures. They used their knowledge of

Table 3. The Taxonomy Development, Refinement and Validation Process

Step	Name	Description	Output
1	Initial development	MMA and LTE independently created candidate taxonomies.	Two candidate taxonomies
2	Initial testing	MMA and LTE applied the candidate taxonomies on sample studies (from the ongoing literature survey) and refined structure and content.	Refined two candidate taxonomies
3	Integration	FFA and MFE integrated the candidate taxonomies.	One integrated taxonomy
4	Review	MMA and LTE reviewed the integrated taxonomy.	Reviewed integrated taxonomy
5	Validation	MUN, FFA, BLE, BNA, and JKH applied the taxonomy on 311 papers originating from the literature review.	Validated taxonomy

psychological research and practice, textbooks, and seminal works to create potential topic areas for the taxonomy. To make the task more specific and concrete, we established constraints during the development of the taxonomy and developed concrete exemplary concepts and assessment procedures.

In the *constraints*, we stated that the purpose of the taxonomies is to classify papers that concern cognitive aspects in SE. Further, we established that the taxonomies should:

- (1) Have distinct categories with minimal overlap so that papers can be unambiguously categorised.
- (2) Have explanations for each category, permitting categorisation as objectively as possible.
- (3) Capture research areas and assessment procedures commonly used in the field of cognitive psychology.
- (4) Have practical utility in supporting an understanding of which areas are covered and which are not, given a set of papers.

In step 2, we provided the experts with *sample papers* and the hint that these papers should be classifiable based on the created taxonomy. The sample papers originated from the ongoing literature survey described in Section 3.2–3.7. While the full validation was conducted in step 6, the purpose of this step was to detect any errors of omission in the taxonomy at an early stage.

We provided *examples* like the following:

Given a paper that investigates the cognitive effort of reading program source code, we have extracted the following information.

Cognitive Concepts: Cognitive Effort

Cognitive Measures: Eye gaze

Assessment: Automatic measurement of eye movement with tool iTrace

This paper could be classified as:

Concept: Working Memory -> Cognitive Load

Assessment procedure: Physiological Measures -> Eye Gaze

However, this is just an example and other classifications could be possible as well. The task is to work out two schemes of classification (one for concepts and another for assessment procedures) that allow this kind of classification to be performed.

In step 3, FFA and MFE integrated the taxonomies received from each psychologist into one taxonomy of cognitive concepts and assessment procedures. The integration was done with the goal of achieving the constraints listed above. Concepts proposed by both experts were retained as such but their position in the taxonomy was chosen considering all other concepts, so as to make the taxonomy logical and parsimonious. Concepts proposed by only one expert were first included but could later be merged with another concept if too similar, made into a sub-concept, or removed. We initially included emotion- and motivation-related concepts but later pruned them from the

Table 4. Terms Used in the Keyword Search

	Terms
Cognitive concepts	“cognitive load” OR “cognitive effort” OR “cognitive resource” OR “cognitive cost” OR “cognitive burden” OR “cognitive overload” OR “cognitive style” OR “cognitive capacity” OR “cognitive work” OR “cognitive psychology” OR “cognitive neuroscience” OR “mental capacity” OR “working memory” OR (cognitive AND attention) OR “task difficulty” OR “cognitive fit” OR “cognitive bias”
AND	
SE knowledge areas	“software engineering” OR “software development” OR “requirements engineering” OR “software testing” OR “software design” OR “software construction” OR “software process” OR “software maintenance” OR “software configuration” OR “software model” OR “software method” OR “software quality” OR “software economics” OR “software project management” OR “software life cycle”

taxonomy since, as discussed in Section 2.1 our emphasis is on individual information processing, and the SE field has investigated emotion and motivation as separate constructs. Proposed concepts related but not the same were analysed using the definitions in the reference literature to determine their relationships, and a corresponding hierarchy was created. There were no concepts where the experts would have proposed differing terminology for essentially the same concept. There were also no directly conflicting concepts. During the merging process, we asked the psychologists for further details and supporting references in case of ambiguity or lack of clarity.

The resulting taxonomy consists of two dimensions (cognitive concepts and assessment procedures), top-level categories and sub-categories. For each category and sub-category, FFA and MFE identified definitions originating from the scientific literature in psychology. The dimension of cognitive concepts consists of nine top-level categories and 15 sub-categories while the taxonomy of cognitive assessment procedures consists of two top-level categories and eight sub-categories.

In step 4, the experts reviewed and accepted the integrated taxonomy (shown in Figure 1).

Finally, in step 5 we classified the set of papers originating from the literature survey on cognitive concepts in SE, which is discussed next.

3.2 Initial Keyword Search

To identify SE studies that investigate cognitive concepts, we followed two strategies. First, MUN, MFE and RFE performed a keyword search in six research databases: ACM Digital Library, Emerald Insight, IEEE Xplore, ScienceDirect, Scopus, and Wiley. Then, after paper selection, we performed backward snowball sampling on the identified primary studies, as explained in Section 3.5. The search string is composed of Boolean expressions including (1) a set of terms related to cognitive concepts and (2) a set of terms related to knowledge areas from SWEBOK (see Table 4).

We iteratively developed the search terms, starting from ‘‘cognitive load’’ AND ‘‘software engineering’’ by performing test searches and adding new terms to the set of cognitive concepts. We also sampled terms from the key chapters in SWEBOK and added them to the set of SE knowledge areas.

We restricted the search, performed in January 2017, to title, abstract and keywords. After removing duplicate results, proceedings and workshop summaries by inspecting the resulting titles, we collected 738 publications.

3.3 Selection of Relevant Papers

We defined explicit criteria to determine the relevance of a study. MUN, MFE, and RFE piloted the selection and refined the inclusion and exclusion criteria during this process. The inclusion criterion (IC1) was: *The paper describes or aims at improving an SE activity, using a cognitive concept as a lens.* We complemented this broad definition with very specific exclusion criteria. A paper should be excluded if:

- EC1. The paper is focused on educational or knowledge management aspects and not on SE activities per se.
- EC2. The paper is about cognitive aspects of the use of the software/system rather than of its development. The dividing line is whether a study looks at the cognitive aspects of performing the design of a user interface, which is included. A study that looks at the effect of a user interface on the user, considering cognitive aspects, is excluded.
- EC3. The paper is about emotional aspects exclusively—i.e., cognitive aspects are not considered.
- EC4. The paper is a method paper and the focus is on a method inspired by, or based on, cognitive psychology rather than on application of such a method.
- EC5. The paper is a literature review (systematic or ad-hoc) and does not add any new empirical results.
- EC6. The cognitive concept is used as a passive motivation, but not actively evaluated or observed in the study.
- EC7. The paper is not written in English.
- EC8. The paper is not peer-reviewed.

We applied these criteria on the title and abstract of a publication during the selection process, but also referred to them during the extraction (Section 3.4) when we read the full-text and needed to exclude a paper based on its (lack of relevant) content or diverging focus that was not clear from reading the abstract alone.

We split the 738 publications resulting from the keyword search into three sets with a 10% overlap which MUN, MFE and RFE reviewed. Within the 75 publications that were reviewed by all researchers, there was an agreement of 79% (59) on the inclusion/exclusion decision. In terms of Fleiss' Kappa [23], the agreement was moderate (0.48). After the entire selection process, we included 208 publications.

3.4 Data Extraction

The data extraction was performed by four researchers (MUN, MFE, RFE, and FFA), using a shared spreadsheet with the data fields shown in Table 5. The meta-data fields were extracted automatically from the search results, while the study setting and cognitive aspects fields required reading the full-text of the paper, together with an interpretation by the respective extractor. The extractor could also decide to exclude a paper, using the exclusion criteria listed in Section 3.3. This decision was reviewed by at least one other extractor and, in case of disagreement, discussed among all extractors. In this stage, we excluded 79 publications and extracted data from 129.

3.5 Backward Snowball Sampling

During the data extraction, we noted that several referenced studies were relevant according to our inclusion criteria, but not found by the keyword search. At this point, we had two options: (1) add new keywords to the original search, or (2) perform snowball sampling to reach relevant papers from the already identified papers. We opted against the first strategy as it would be difficult to decide which keywords should be added and when to stop the addition of new keywords. We chose therefore to do backward snowball sampling from the 129 already extracted papers by reading first title, and if necessary, the abstract of referenced papers. MUN, MFE, FFA, and BMA carried out the snowball sampling, resulting in additional 113 papers from which we extracted data.

3.6 Classification

We used the taxonomies on cognitive concepts and assessment procedures, developed as described in Section 3.1, to further classify the data we extracted from the primary studies, as described

Table 5. Data Extraction Form

Data field	Data type	Value
<i>Meta-data</i>		
Publication venue	free text	Name of the outlet in which the paper was published.
Year	free text	Year in which the paper was first published.
<i>Study setting</i>		
Context	restricted	Based on the study participants, the context was determined as academia, industry, or both. If the study did not involve participants, we set the context to academia.
Type	restricted	We adapted Wieringa et al. [60] to the following study types: solution proposal, empirical study, experience paper, philosophical paper, opinion paper. Since one publication can report on multiple studies, several study types can be selected for one paper.
<i>Cognitive aspects</i>		
Cognitive concepts	free text	Concepts discussed in relation to SE and cognition. E.g., procedural or structural problem solving.
Cognitive measures	free text	The item in relation to cognition that is measured. E.g., comprehension based on structural or procedural solution.
Assessment	free text	The method by which the cognitive concept is measured. E.g., questions relating to the programming task at hand.
Added value	free text	A summary stating to what extent the cognitive concept is central or tangential to the conducted study reported in the paper.
Theory	free text	Name(s) of theories, models, or frameworks related to cognitive concepts that are mentioned in the paper.
Outcome/contribution	free text	A summary or copy of the papers' main contribution. Can be empty.
SWEBOK area	restricted	The sixteen SWEBOK knowledge areas (KAs), with possible multiple selection. A KA was selected if the cognitive concept was studied <i>in</i> or had an impact <i>on</i> the KA: software requirements, software design, software construction, software testing, software maintenance, software configuration management, SE management, SE process, SE models and methods, software quality, SE professional practice, SE economy, computing foundations, mathematical foundations, engineering foundations.
Excluded	restricted	Binary choice; decision based on the full-text reading of the paper, applying the inclusion/exclusion criteria.

in Section 3.4. MUN used the extracted values from the data fields *cognitive concepts*, *cognitive measures*, *assessment*, and *added value* (see Table 5) to associate each primary study with one or more dimensions of the two taxonomies, using the definitions shown in Figure 1. In this step, we further excluded 25 publications with the following motivations:

- The discussion of the cognitive concept was not central to the study but rather used as a motivation/justification.
- The discussion of the cognitive concept was superficial and lacked details to a degree that made classification impossible.

The decision to exclude a publication was discussed among MUN, BMA, MFE, and FFA, resulting in a total of 217 publications that were classified in this step.

3.7 Updated Review

To bring the most recent publications into the result set, we performed a second round of keyword searching, selection, extraction, and classification. We opted for a new search rather than a forward snowballing from the set obtained in the first round because we wanted to bring in more recent publications while remaining consistent with respect to the kinds of papers found, and strengthening the reliability of the first round results by verifying the searches. BLE, BNA, JKH, and FFA first ran the original keyword searches in June 2020 without any publication year restrictions to verify the search protocol. When comparing the results, we noticed that some databases returned additional results from the period covered by the first search. We speculate that some databases may have been updated to include additional publications since the first search, or that their search algorithms may have changed. We analysed the results of each search to determine the shortest date range that would cover all papers not yet found. The ranges varied from 1979–2020 to 2017–2020, with most being closer to the latter. After removing duplicates, we collected 1,030 new publications, extending our review to cover papers published and indexed up to May 2020.

FFA, BLE, BNA, and JKH then selected papers to include based on the inclusion and exclusion criteria. One extractor proposed a selection decision which was then verified by another extractor. In case of disagreements, the decision was resolved through discussion among all extractors. This procedure meant that all paper extractions were verified in this phase. As before, we also excluded papers in the extraction phase using the same verification procedure. After all exclusions, 96 new publications remained, from which FFA, BLE, BNA, and JKH extracted information and classified the publications according to the two taxonomies. The extraction and classification of approximately 17% of these papers was checked by having a different person among these four researchers independently verify them. We resolved conflicts through discussion sessions between the original extractor and the verifier, while FFA resolved any remaining conflicts. Slightly less than 5% of the extracted papers had very minor discrepancies (e.g., formulation of main results of the study). Therefore, we did not perform further validation on the remaining 83% of the papers.

We decided not to perform snowball sampling on the 96 new publications, because the first round of searches and snowball sampling would already cover most of the range where citations from the new publications could be found. Also, this second phase was aimed primarily at extending the range of included papers to cover the years 2017–2020, rather than increasing the set of papers overall.

In total, the final set extracted and classified consists of 311 papers (217 in the first phase and 96 in the second phase) published between 1973 and May 2020.²

3.8 Threats to Validity

In this section we discuss threats to validity of our secondary literature study and applied mitigation actions. We follow the respective guidelines for managing threats to validity of secondary studies in SE by Ampatzoglou et al. [1, 2]. These guidelines distinguish validity threats to study selection, data, and research.

Study selection validity threats deal with searching and including primary studies in the examined set. In the first round, inclusions in the selection stage could be reversed in the extraction stage. However, not all exclusions were verified by a second reviewer in the selection stage. This could have resulted in some papers being erroneously excluded. Still, the agreement on the papers that were checked by another reviewer (see Section 3.3) indicates that this was only a moderate

²The final result set with the extracted data is available as online supplementary material [21].

concern. Furthermore, because of the search verification in the second round, papers erroneously excluded could be brought back into the set. All inclusion and exclusion decisions were verified by a second extractor in the second round.

We opted for a more general search string, but it is possible that some relevant papers could use specific terms and could not be found in the initial search. Through snowballing, we would expect this threat to be reduced. We note that in the second keyword search, we found papers from the time period targeted in the first search. This is a potential threat to validity because it questions the repeatability of the search. We believe this may be due to updates in the databases used, either because of new publications being included, or because the search algorithms were updated. However, this limitation pertains to all literature reviews relying on database searches and is unavoidable in practice, as there is no way to verify what the databases contain nor how they execute the searches. Furthermore, we did not perform snowball sampling on the publications found in the second keyword search. This may have resulted in some publications not being found. While we acknowledge this risk, we consider it acceptable given that the keyword search and the first round of snowballing cover many relevant publications which are likely to be cited in the primary studies found after the second round of search. We already found almost as many papers with snowballing as with keyword search in the first round. This is not common for systematic literature reviews. However, the area of cognitive concepts in SE is quite fragmented. We experimented with different search strings, but could not settle on a definitive one. Therefore, snowballing played a key role in our setting. A further limitation comes from not performing forward snowballing, which could have resulted in some papers being omitted. However, comparing the effort to the uncertain gains spoke for improving the quality of the result set in terms of its coverage of similar papers rather than an expansion of the set towards areas further away from the core topics. This way, the results of the present study should represent a more focused starting point for potential future SLR updates.

Data validity threats deal with the extracted data set and its analysis. To reduce threats to data extraction, this activity was split among several researchers. Each extracted paper in the first round, and at least a sixth of the papers in the second round ($\geq 17\%$) was reviewed by at least one other extractor. Because not all papers in the second round were verified for extraction, there could be some errors in the extracted data. However, only minor discrepancies were found in a very small number of the papers reviewed in the second round ($\leq 5\%$), indicating that such errors are few. Furthermore, each extraction criterion was clearly defined—i.e., using the taxonomy of cognitive concepts and assessment procedures, the SE Body of Knowledge, and simple demographic data like the year of publication. Overall, we had a reasonable size of primary studies (i.e., 311), and applied descriptive statistical analysis, which keeps potential threats to data analysis low.

Research validity threats concern the overall research design. To guarantee repeatability, we documented the entire research method and make the raw data and the analysis procedures available. The research scope is cognition in SE. We do not aim to generalise our results beyond SE. The core contribution of this article is a taxonomy on cognitive concepts and assessment procedures to classify and discuss existing and future work. As this research goal requires expertise from both SE and cognitive psychology, the research was performed by a team of six SE researchers (FFA, MFE, RFE, DFU, MUN, BMA) with BLE, BNA, and JKH joining in the second phase of the literature search, and two researchers in psychology (MMA, LTE).

4 A TAXONOMY OF COGNITIVE CONCEPTS AND ASSESSMENT PROCEDURES IN SE

To address RQ1, we present a taxonomy of cognitive concepts and assessment procedures, shown in Figure 1. The taxonomy includes two dimensions, *Cognitive Concepts* and *Assessment Procedures*.

Cognitive Concepts define *what* can be assessed, while Assessment Procedures define *how* those concepts can be studied. The concepts are arranged hierarchically. The different levels of the hierarchy can be used to classify papers that differ in conceptual preciseness. Papers that study a cognitive concept in more general terms can be classified using a higher-level concept, while papers that are explicitly studying a more narrow concept can be placed in a leaf category. The hierarchically-organised concepts proceed from more basic processes on the top to higher-order ones on the bottom. Some concepts are cross-cutting, spanning the range of all other concepts. Each combination of a cross-cutting concern with another concept constitutes a particular kind of variation on a cognitive process.

The taxonomy serves as an analytical tool to characterise research on cognition in SE and shows the extent to which articles focus on specific topics. We argue that it can inform practitioners by giving an overview of the field and helping to judge how aspects of cognition can affect them and their organisations. Furthermore, it helps to identify knowledge gaps and point researchers to new investigations.

4.1 Cognitive Concepts

The cognitive concepts in the taxonomy capture a collection of important processes and activities of individual cognition at different levels, providing an answer to the classification part of RQ1. We briefly expand on each concept below.

4.1.1 Perception. The processes that organise and interpret sensory information in order to understand it or the environment it presents (e.g., visual information regarding a three-dimensional world) are called *perception*. Through perception, sensory stimuli are interpreted as having been produced by properties of objects or events in the environment. Perceptual processes are not only passive reception of sensory information, but are shaped by other cognitive processes, such as memory and attention. Information processed by perception includes sensory information—vision, sound, touch, taste, and smell—but also social information, such as understanding of several aspects of speech and faces, as well as perception of time, agency, and familiarity.

4.1.2 Attention. Senses (e.g., vision, hearing, and touch) are systems involved in sensation, the physical process where stimuli (e.g., light, sound, and pressure) are collected and provided for perception. *Attention* is a state of focused awareness on a subset of the available perceptual information. In other words, attention refers to the selection of what perceptual information an individual primarily focuses on.

This concept is divided into three sub-concepts: *selective*, *divided*, and *sustained* attention. *Selective attention* is the process of directing the limited resources of attention towards specific stimuli while ignoring other stimuli. This enables humans to focus on what is deemed important in the situation or task. *Divided attention* refers to the ability to integrate multiple stimuli in parallel. It enables paying attention to multiple things at the same time and is required for many everyday tasks as well as complex tasks in the workplace. For example, a software developer may have to juggle between information coming from several different programs to track down a program fault. *Sustained attention* refers to the ability to focus on a specific task or thought for an extended period of time, while avoiding being distracted by other stimuli. This is especially important in complex, knowledge-intensive tasks that require concentration over long periods of time. It is also involved in learning-related processes, such as program comprehension.

4.1.3 Memory. The mental processes of acquiring and retaining information for later retrieval, as well as the mental storage system that enables these processes, are referred to as *memory*. Memory can be divided into sub-concepts based on its role and how long information is retained. *Working memory* is a brain function providing temporary storage and manipulation of information



Fig. 1. Taxonomy of Cognitive Research in SE. The taxonomy is divided into two main dimensions: Cognitive Concepts and Assessment Procedures.

needed for complex tasks such as language comprehension, learning, and reasoning. *Short-term memory* holds current and recently attended information briefly, in the range of seconds or minutes. *Long-term memory* holds information for longer periods, being virtually permanent.

4.1.4 *Cognitive Load*. The amount of information-processing required, especially in terms of attentional and working memory, differs between tasks. The information-processing needs during a specific time period is referred to as *cognitive load*, and can be divided into the load that relates

to task complexity (*intrinsic cognitive load*) and the load that relates to cognitive processes that do not contribute to performing the task itself (*extraneous cognitive load*).

4.1.5 Reasoning. Thinking that is directed toward a given goal or objective is referred to as *Reasoning*. This includes *problem-solving*, which is thinking directed toward solving specific problems, moving from an initial state to a goal state by means of mental operations. Another kind of reasoning is *decision-making*, which refers to the process of choosing between alternatives, selecting or rejecting available options, and potentially creating new options.

4.1.6 Cognitive Biases. *Cognitive biases* refer to patterns of deviation from norm or rationality which originate from reduction of complex judgement tasks to simpler judgement operations. In general, such simplifications are useful and serve to produce quick results that often are accurate enough but that sometimes lead to severe systematic errors. Cognitive biases can be seen as particular kinds of errors (see below) and originate from the limitations of other cognitive processes.

4.1.7 Knowledge. Experience, values, contextual information, and expert insight form a framework referred to as *knowledge*. This framework is used to evaluate new experiences and information, but the framework itself is also altered in the process. Specific kinds of knowledge are identified. *Explicit knowledge* can be expressed as words and numbers, and is readily shared and transmitted between individuals. In contrast, *tacit knowledge* is personal, hard to formalise, and difficult to communicate or share with others. Such knowledge can be divided into *technical tacit knowledge*, or “know-how,” and *cognitive tacit knowledge*—i.e., those beliefs, ideals, values, schemata, and mental models which are deeply ingrained in individuals and often taken for granted.

4.1.8 Social Cognition. Perceptual processes provide information about the social context of an individual, such as facial expressions, actions, and other overt behaviour of other individuals. The individual also displays overt behaviours to others in the social context. The cognitive activity that accompanies and mediates social behaviour, including acquisition of information about the social environment, organisation and transformation of this information in memory, and its effects on the individual’s behaviour are referred to as *social cognition*. It is a cross-cutting concept: it includes many levels of cognition and interacts with different concepts, from lower levels of the taxonomy to higher-order functions.

4.1.9 Cognitive Control. *Cognitive control*, also known as *executive function* is a set of cognitive processes that work to govern behaviour. This includes both basic cognitive processes, such as attentional control, cognitive inhibition, working memory, and cognitive flexibility, and higher-order functions such as planning, reasoning, and problem-solving. Cognitive control is not a single process but multiple processes that help monitor and select behaviours that facilitate attainment of goals. It is thus a cross-cutting concept in the taxonomy.

4.1.10 Errors. Cognitive processes, even when functioning correctly, can produce irrational results. Such *errors* originate from, for example, perceptual distortions, inaccurate judgement, or other irrationalities.

4.2 Assessment Procedures

Assessment procedures are different means of obtaining systematic information about the quality of each cognitive concept in a particular context. This ranges from measures obtained from measurement devices to observations made by the researcher. We divide the research methods into *quantitative* and *qualitative*, and base their characterisation on the *Bibliography* reported in Figure 1. This provides an answer to the assessment part of RQ1.

4.2.1 Qualitative. When the assessment procedure is based on rich data and human interpretation, we refer to it as *qualitative*. In such research, the researcher interprets information without necessarily decomposing it into its constituent parts or assigning numbers to the observations. Qualitative assessment procedures include *fieldwork*, where the researcher enters the environment where the assessment is to be made, possibly participating in actions and interaction, for an extended period of time; *interviews*, where the researcher obtains first-person or documented accounts of experiences either from an individual or as a group (e.g., focus group); *task-based* procedures, where the researcher observes and records the conduct of specified tasks; and *open observation*, where assessment is based on observation in a chosen context for a limited time.

4.2.2 Quantitative. When the assessment procedure is based on data on a numeric scale and analysis of statistical properties, we refer to it as *quantitative*. Such research includes *task performance* assessments, where participants perform a task and predefined measurements are obtained; *physiological*, where measurements are obtained from participants' bodies; *subjective ratings*, where raters, who may or may not be the participants themselves, rate artefacts using predefined scales; and *behavioural*, where participants' overt behaviour is coded into numeric data according to a predefined scheme, either automatically or manually by specially trained raters.

5 LITERATURE SURVEY RESULTS

Research on cognition in SE has seen attention for more than five decades. This body of literature covers several research approaches, SE knowledge areas, and study contexts, all of which are detailed in this section. We use our taxonomy (see Figure 1), together with the SE Body of Knowledge (SWEBOK) [8], to illustrate the foci of the last five decades of research on cognitive concepts in SE and provide answers to RQ2 and its sub-questions. At the same time, these results act as a validation of the taxonomy of cognitive concepts and assessment procedures. We show how the taxonomy is able to organise and situate the research in terms of concepts from and related to cognitive psychology. In the figures below, we omit the concepts *Cognitive control* and *Social cognition*, as the former includes only one publication from 1988, and the latter includes only seven publications between 2008 and 2020. However, we analyse these concepts and the related publications where applicable. We also omit the super-category *Errors* when the results only contain articles in the sub-category *Cognitive biases*.

5.1 Research Evolution, Context, and Type

We first examine the contexts in which research on cognitive concepts in SE has been performed (RQ2.1) and the evolution of that research over time (RQ2.2). Figure 2 shows the number of publications per year. The earliest paper found was from 1973, reporting on a psychological evaluation of two conditional constructs commonly found in many programming languages, *if-goto* and *if-then-else* [53]. The paper belongs to the long-standing family of studies in program comprehension. Interestingly, the authors are affiliated with a department of psychology, showing how that field has had an influence on the body of knowledge related to SE for decades. Nevertheless, as noted earlier, psychology cannot be said to have become part of the foundations of SE research.

We analysed the outlets that publish work on cognitive concepts, the context in which published research takes place, as well as the publication types (Figure 3). In particular, Figure 3(a) shows the number of publications in different types of scientific outlets. The publications can be mainly found in journals and conferences, as expected in this field. Workshop and magazine articles, together with book chapters, make up roughly 10% of the total set of primary studies. Different cognitive concepts are covered to a different extent in different types of outlets (see Figure 3(b)). *Cognitive load*, *Memory* and *Cognitive biases* are each covered in workshop papers, while the other areas are





Fig. 2. Distribution of publications per year. The data ranges from 1973 to May 2020.

not discussed in workshop papers. Papers regarding all the cognitive concepts, with the exception of *Attention* are published in similar numbers in journals and conference proceedings.

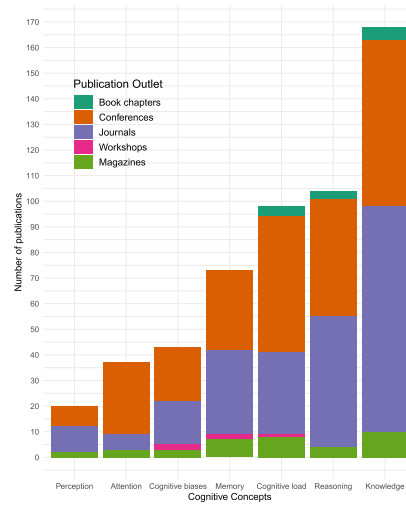
Figure 3(c) shows the number of publications according to the context in which the studies were conducted. Studies performed in an academic context, which typically means having student participants, are only slightly more common than studies in an industry context. Studies in which both an academic and industry context are utilised are the least common. There are also studies with no empirical context—these are theoretical, philosophical, or position papers that do not include an empirical component but can contribute to conceptual and theoretical development in the area.

Figure 3(d) shows which kind of papers, according to the Wieringa et al. SE publications taxonomy [60], cover the different cognitive concepts. The *Knowledge* concept is most addressed, through a large number of empirical studies. However, *Cognitive load*, the second-most addressed concept overall, had more solution proposal, philosophical, and opinion papers. The third-most addressed concept was *Cognitive biases*, while *Memory*, *Attention*, *Perception*, and *Reasoning* were the least addressed, in that order. A large majority of all papers were empirical studies. Solution proposals and philosophical papers were frequent but much less so than the empirical papers. Some opinion and experience papers were found; the former concerning all concepts, and the latter concerning *Knowledge*, *Cognitive load*, *Cognitive biases*, and *Attention*. Overall, this indicates that the research represented by the papers in our result set is grounded in empirical evidence, that there is some philosophical grounding in all the concepts, and that some solution proposals have been put forward.


Research on different cognitive concepts has evolved over time, as shown in Figure 4. Studies related to *Reasoning*, *Perception*, *Memory*, and *Knowledge* are published regularly since 1973, and are prevalent until the early 1990s while other areas have longer gaps between years of publication during the period we consider. Studies dealing with *Cognitive load* and *Cognitive biases* first appeared in the early 1980s but were in focus after the year 2000. In 2016, more than 75% of papers were published in these areas. Research on *Attention*, after being introduced in the literature in 1979, and with some occurrences in the first part of the 2000s, gained traction in 2012 and in 2015 almost one third of the published primary studies dealt with this concept. *Social cognition* (omitted



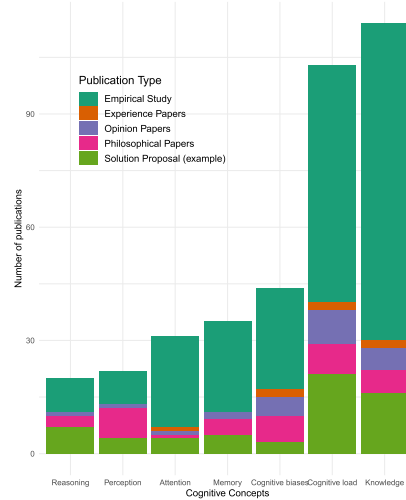
(a) Publications per scientific outlets.



(b) Publications per outlets and cognitive concepts.



(c) Publications per research context.



(d) Publications per study types and cognitive concept.

Fig. 3. Distribution of SE publications according to outlets, paper type, and cognitive concepts.

from the figure) has seen recent attention, but it is too early to say if this will lead to a longer-term research focus. *Cognitive control* (also omitted from the figure) appears not to have gained such focus, with only one publication from 1988.

5.2 Coverage of Concepts and SWEBOK Areas

We now turn to examining the SE knowledge areas that have been studied under the lens of cognitive concepts, addressing RQ2.3. Figure 5 shows the distribution of publications according to SWEBOK areas. The largest number of publications are found in Software Construction, Design, Requirements, and Maintenance, each with at least 50 publications.³ There are relatively few

³A paper can cover one or more SWEBOK area.




Fig. 4. Percentage of publications covering the different cognitive concepts for each year. The black line indicates the total number of publications. No publications for years 1974–1976 and 1978. A publication can cover more than one concept.

papers on Software Testing, Process, and Quality. There is no publication in the Computing and Mathematical foundation, and a single publication in the Engineering foundation.⁴

The concept coverage differs between SWEBOK areas (see Figure 6). Four central SE life-cycle areas—Requirements, Design, Construction, and Maintenance—are most often covered by the *Knowledge* concept and, to a lesser extent, by *Cognitive load*, *Reasoning*, *Memory*, and *Cognitive biases*. *Knowledge* is a broad concept and it is inevitable that it is covered in many primary studies. No concept is predominant in the area of Software Testing, for which *Reasoning*, *Perception*, and *Attention* are scarcely covered.

Part of this result can be attributed to the fact that there are, overall, more papers in the main life-cycle areas in SWEBOK than in areas such as SE Management or Economics. However, this does not fully explain the lack of research in Testing, Configuration Management, Processes, or Quality—all areas with large numbers of publications. It is also worth noting that SE Professional Practice, which explicitly covers human aspects such as individual cognition and problem-solving, has very few papers covering any cognitive concept at all. The lack of research in these areas can be partly explained by papers not being classified appropriately when their main focus is on a life-cycle area.

Examining the areas and concepts more closely, we see how the taxonomy organises the research. *Attention* stands out as being investigated mostly in Software Construction, Maintenance, and to a lesser extent in Professional Practice. This shows that the taxonomy is able to capture an important cognitive process that is relevant to performing technical, and often programming-related tasks. Also with *Reasoning*, we see a difference in how the cognitive concept is emphasised

⁴Due to the low number of publications, we do not report results about these areas in the rest of this paper.




Fig. 5. Number of publications according to SWEBoK areas.

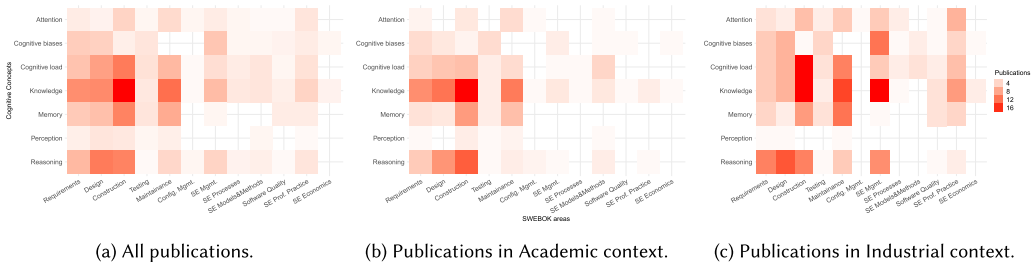


Fig. 6. Publications characterised by Cognitive Concepts in the different SWEBoK areas. Computing, Mathematical, and Engineering foundations are excluded as there is only one study published in these areas.

in different SE areas. While research on *Attention* focuses on construction and maintenance tasks, *Reasoning* covers more of the earlier life-cycles stages such as Requirements and Design. This is in line with the definition of the concepts, since *Reasoning* includes problem-solving and decision-making activities that are relevant for these SWEBoK areas. It may be more difficult, and therefore less common, to investigate *Attention* in such areas.

Figure 7 shows the distribution of publications per year according to the SWEBoK areas. The published work in Software Testing has increased slightly since 2010, indicating a possibility that cognition will be better covered in this area in the future compared to our findings. Research in Software Quality has also grown recently, but the number of publications in this area is still small and papers are sporadically published.

The research on cognitive concepts has had focal changes over the observed time period. We observed changes in different parts of the taxonomy. Focus on *Attention* has increased over the observation period. However, *Selective*, *Divided* and *Sustained Attention* are not commonly reported. For the latter, all research is quite recent. We believe this indicates a lack of conceptual clarity in research investigating *Attention*. The role of attention in several software development

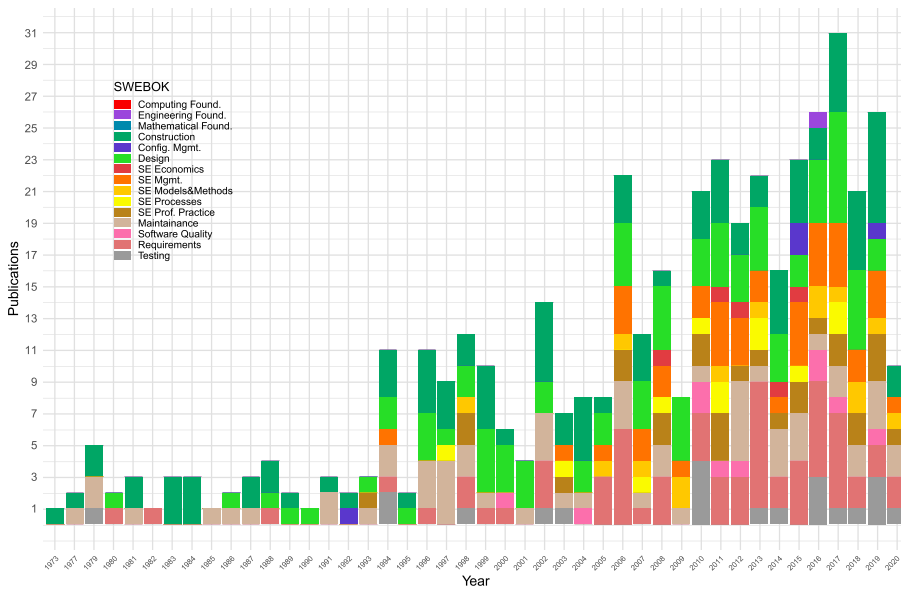


Fig. 7. Distribution of publications covering the SWEBOK areas per year.

tasks is important. Therefore, research in this area needs to be more grounded in established conceptualisations and specify more precisely what aspect of attention is investigated.

Memory has been researched more during 1973-1988 and after that the work in this area seems to have declined. *Working memory* has sporadic coverage over the whole area and more regular coverage in the last ten years, but there are still few primary studies on this concept. *Short-term memory* was covered more in the 1970s and 1980s, and again at the end of the 1990s and start of 2000s, but there have been no papers published on this concept in the last ten years. *Long-term memory* is covered in some papers, with a concentration at the end of the 1990s and start of 2000s, but seems to have declined since then.

Cognitive load has become more regularly addressed and the relative amount of research has grown. Cognitive load is a concept that feels intuitive to address, as practitioners and researchers alike can relate to the thought of being overloaded during a demanding task. However, we see room for improvement in the use of this concept. First, it is most commonly researched as a general concept without much regard to the underlying explanation of how cognitive load arises. This can be observed in the very sparse occurrences of research that explicitly examines the *Intrinsic* and *Extrinsic cognitive load* concepts. Thus, studies do not seem to conceptually differentiate whether the load is related to task complexity or evoked by the task design. Second, as cognitive load is related especially to working memory, we would expect to see many papers on cognitive load to also address *Memory*. However, that concept is in decline, and only 20% of cognitive load papers consider memory. Studies could gain explanatory power by considering the relationship between cognitive load and memory, and by building on related studies in cognitive psychology.

Perception is regularly addressed during the observation period. There are not many papers on this concept but the regularity of publications has increased to some extent. It can be useful to consider the human perceptual processes and the concepts around them more specifically in SE research.

Research on *Reasoning* covers the whole period, but only sporadically. The *Problem-solving* concept is frequently addressed and there is a relatively large number of publications, which seems to

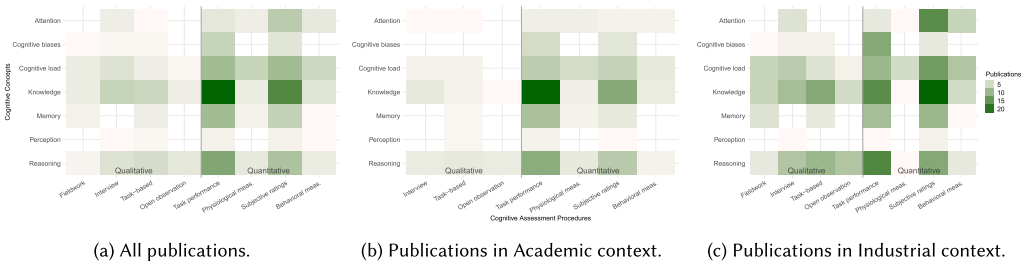


Fig. 8. Publications characterised by Cognitive Concepts according to different Assessment procedures.

be declining. *Decision-making* is more regularly addressed in recent years, but does not represent a particularly large portion of the studies.

The *Cognitive biases (Errors)* concept has been addressed more regularly since 2005. However, it seems too early to say whether it will grow to cover a significant portion of the research on cognition in SE.

Knowledge is regularly and frequently addressed over the whole period. Similarly to *Attention*, the more precise division into sub-concepts is less common. *Explicit knowledge* is sporadically addressed and seems to have decreased in recent years. *Tacit knowledge* is covered in some papers since 1998, but sporadically and in small numbers. *Technical tacit knowledge* has been covered sporadically throughout the period but has become more sparse between 2010 and 2020. *Cognitive tacit knowledge* has been addressed very sparsely.

5.3 Coverage of Assessment Procedures

RQ2.4 poses the question of how cognitive concepts in SE have been assessed. Figure 8 shows the number of primary studies covering the cognitive concepts according to different assessment procedures, both qualitative and quantitative. There are differences in how the concepts are assessed, and most assessments are quantitative. *Knowledge*, *Cognitive biases (Errors)*, *Memory*, *Reasoning*, and *Attention* are mainly assessed using quantitative procedures, especially *Task Performance* (e.g., a programming exercise) and *Subjective Ratings* (e.g., self-reported understanding of requirements). *Physiological* and *Behavioural* procedures are under-utilised, perhaps due to limited availability of equipment for the former, and the lack of methodological training and guidance for both.

Qualitative assessments make use of *Task-based* studies, *Interviews*, and, to a lesser extent, *Fieldwork* observations. Nonetheless, *Fieldwork* is utilised for several concepts (e.g., *Cognitive load*, *Knowledge*, *Reasoning*, and *Memory*), whereas current research on cognition in SE lacks *Open observations* assessments.

The methods used likely reflect the methods familiar to the SE research community, the training typically included in researcher education in the field, and the traditions related to different methodological approaches. These aspects are all likely to change over time, and that is in fact what we found. Figure 9 shows the number of publications in different years according to the type of research method. Throughout the observed time period, quantitative methods represent the largest part of the research. However, qualitative methods are also utilised and, for some years, represent half of the total published research. More recently, mixed (or multimethod) approaches have become more common. These refer to studies which have two (or more) separate approaches side-by-side to investigate a topic, but they also include mixed-methods designs that strategically integrate quantitative and qualitative methods based on an explicit research philosophy, use specifically chosen mixing points, and explicate a logic for deriving valid results from the combination. The latter still appears to be less common than simpler combinations, and we have not attempted to separate the two in this analysis.




Fig. 9. Research methods over time.

Figure 10 shows the distribution of publications covering different research methods in SWE-BOK areas. Generally, the areas with most research are also the areas where qualitative and mixed methods are used the most. In SE Professional Practice, only half of the publications utilise a quantitative approach, and in SE Economics and Engineering Foundations, qualitative approaches outweigh the quantitative.

Finally, we examined how different cognitive concepts have been investigated using different assessment procedures. Among the qualitative approaches, *Interviews* are used for almost all concepts, with *Knowledge* and *Cognitive load* being the two most common. Task-based procedures are also common, and most used for *Problem-solving* and *Knowledge*. *Fieldwork* is used sparingly and mostly for *Knowledge* and *Cognitive load*, while *Open observation* is also used sparingly and then mainly for *Reasoning*, *Problem-solving*, and *Knowledge*.

The most common quantitative assessment procedures are *Task performance* and *Subjective ratings*. Both have a remarkably similar profile in terms of what concepts they have been used in. *Knowledge*, *Cognitive load*, and *Problem-solving* are frequent targets of both. *Task performance* has been used more for specifically assessing *Intrinsic* and *Extrinsic cognitive load* as well as *Short-term memory*, but these are relatively infrequent uses. Among the other procedures, *Physiological measures* have been used most frequently for *Cognitive load* and to a slightly lesser extent for *Problem-solving* and *Knowledge*. *Behavioural measures* are most commonly used to assess *Cognitive load*, *Knowledge*, and *Attention*.

We have also extracted cognitive measures, if identifiable, from the primary studies. Table 6 illustrates a sample for each assessment procedure, indicating also which cognitive concept is measured. The complete list of measures is available in the online supplementary material [21].

5.4 Use and Value of Psychological Theories

We now turn to the role of theory in studies on cognition in SE, addressing RQ2.5. The examined papers mention a wide range of different theories, theoretical concepts, or conceptual frameworks as general inspiration. Some papers utilise theories or conceptual frameworks developed within




Fig. 10. Research methods in SWEBOK areas.

Table 6. Examples of Cognitive Measures

Assessment procedure	Cognitive measure
	Qualitative
Fieldwork	Experiences concerning knowledge-transfer, task performance (Cognitive load/Knowledge); The occurrence of phenomena in a software development project that can be explained by a variety of theories (Reasoning/Cognitive biases)
Interview	Experiences with job/project rotations (Knowledge); Merge conflict processes, barriers, and resolution strategies (Reasoning/Knowledge); Experiences concerning knowledge transfer (Knowledge)
Task-based	Description of how experts and novices perform software design (Reasoning/Knowledge); Type and frequency of utterances, such as inquiries, questions and conjectures (Reasoning); Comparison of code comprehension by practitioners with theoretical model (Knowledge)
Open observation	Shared mental models and teamwork characteristics (Reasoning/Knowledge); Identification and description of wasteful activities in software development (Cognitive load/Knowledge)
	Quantitative
Task performance	Influence of anchoring on estimation performance (Cognitive biases); Balance of domain-specific and programming-specific information in the mental representations participants had of a program at different points in time (Knowledge); Ability to identify and fix bugs in the source code (Cognitive load/Knowledge)
Physiological measures	Electrical signal emitted from muscle nerves (Reasoning/Knowledge); Pupillography (Cognitive load); Eye fixation (Reasoning)
Subjective ratings	Visual memory capacity (Memory); Vulnerability of SE tasks to interruptions, dependent on suspension length and number of nested interruptions (Attention); Perceived cognitive load (Cognitive load)
Behavioural measures	The proportional fixation time on each area of interest (Attention); Active time spent in IDE (Attention)

the field of computing, such as the cognitive dimensions of notations framework [28], “physics” of notations [36], and cognitive informatics [58]. However, the vast majority of papers do not use theory or theory-grounded conceptual frameworks to situate the research, to inform their research design, or to explicitly draw hypotheses. No papers use their results to clearly and explicitly contribute back to theory, i.e., by refuting, supporting, altering, or extending a theory.

Studies dealing with *Attention* use no psychological theories specific to attention. Most papers in this area appear in combination with other cognitive concepts, and many of the theories mentioned are related to cognitive load. Program comprehension is a frequent research topic among these papers. It is perhaps somewhat surprising that specific psychological theories of attention are not more specifically included as part of the theoretical palette in such studies.

Studies on *Memory* mostly do not mention any specific memory theory, but some papers do mention theoretical concepts such as chunking (e.g., [43]), aids for memory failures [38], and multi-store memory theory [12]. However, with the exception of chunking and other program comprehension theories relating to memory, these are mainly part of setting the scene for the research and not explicit theories that would be used for informing the research design or hypotheses.

Papers examining *Cognitive load* mostly mention some general form of cognitive load concept or theory, but usually without specifying in detail what the theory proposes or drawing any hypotheses specifically from the theory. *Cognitive load* is also examined from the perspective of code or program comprehension, with theories such as bottom-up and top-down comprehension being mentioned in some papers (e.g., [51]). Program comprehension studies have a long history during which theories specific to that field have been developed. This may in part explain why the theories are not among the ones found in cognitive psychology in general: the theories have already diverged from the more general theories in psychology. However, this cannot be said for most of the concepts examined in the papers.

There were no specific theories related to *Perception* or *Reasoning*. For the latter, some papers do mention broad theoretical ideas as background knowledge. Examples include problem space theory [31], the general idea that problem-solving involves searching a problem space, and Simon's theory on humans as information-processors that can be modelled similarly to computers [54]. Similarly, *Knowledge*, *Cognitive control*, and *Social cognition* papers did not include specific psychological theories.

While few of the papers make extensive use of theory, the cognitive concepts themselves play a central role in most of the papers. The specific cognitive concepts used in the vast majority of papers (which may be more detailed than the general categories in our taxonomy) really are central to the papers in their entirety, to the research approach, or to the arguments made. While they may not be well grounded in psychological theories, they do have specific intended meanings. For example, cognitive load is used in many papers to signify human limitations in information processing capacity. Several papers attempt to manipulate cognitive load in different ways to cause observable changes in task performance (e.g., error rates). Other papers use the cognitive concepts less explicitly: to provide a conceptual framework to explain what is studied, to inform the study design without providing an explicit theory, or to situate the research in the field of human factors in SE.

However, the use of cognitive concepts is somewhat imprecise in many papers, largely due to the lack of a connection to established definitions and theories in psychology and cognitive science. As an example of a particular pattern, a paper may purport to examine mental effort—a dynamic characteristic of an individual human—but actually end up assessing task difficulty—a characteristic of the task. Furthermore, some papers seem disconnected from the cognitive concept they claim to examine. Their use of the concept is more of a passing motivation resembling an inspiration or argumentative device to support the relevance of the study, or an assumption or conjecture on the authors' part. An example of this pattern would be a paper arguing that a specific cognitive process is the underlying cause for specific human behaviour that the study has observed in software development, but does not actually demonstrate that this is the case.

Thus, the value of the cognitive concepts appears to differ a lot between the papers. Authors whose own conceptual development is thorough succeed better at producing trustworthy results. This applies both in papers that seek to deepen the understanding of cognition in SE and in papers that seek to introduce new methods or tools to aid cognitive processes among software engineers.

6 DISCUSSION

How software engineers think has been studied for over five decades, leading to insights on how they work and how individuals' performance can be improved through means of tooling, work

organisation, and training. Software construction has been in the centre of attention (98 primary studies), followed by software design (84 primary studies), and software requirements (66 primary studies).

We described the construction of a taxonomy for cognitive concepts and assessment procedures, and used it in a literature survey on research on cognition in SE. The taxonomy is based on expertise from cognitive psychology and related areas, and provides a viewpoint on research in SE that relates to cognitive processes. Our application of the taxonomy illustrates how it can organise and improve research by better utilising concepts from psychology and cognitive science as reference fields. Our taxonomy also shows the breadth of assessment options that exist and that researchers should consider.

We asked how cognitive concepts relevant for SE can be classified and measured (RQ1). Our taxonomy, and the examples of papers covering the cognitive concepts and assessment procedures, provides an answer to this question. We propose that cognition in SE should be conceptualised by drawing from reference fields such as psychology and cognitive science, where decades of research have resulted in well-defined and empirically defensible conceptualisations. Trying to invent new concepts to investigate cognition in SE will likely result in conceptual confusion and in a set of concepts which are not suitable to investigate cognition.

We investigated what is the state of the art in cognitive concepts in SE (RQ2). The most developed strands of research seem to be in the areas of the four life-cycle stages *Software Requirements*, *Design*, *Construction*, and *Maintenance*. Most research is quantitative and focuses on *Knowledge*, *Cognitive load*, *Memory*, and *Reasoning*. Surprisingly, there is a clear lack of research on cognitive concepts in *Software Testing*. The state of the art overall appears fragmented when viewed from the perspective of cognition. There is a lack of use of cognitive concepts that would represent a coherent picture of the cognitive processes active in specific tasks. The research methods are focused mainly on approaches that are familiar to SE researchers and, while SE research seems to adopt new methods, the uptake of methods from fields like psychology seems slow. Nevertheless, among the questions and topics examined, there are good examples of research in cognition, as shown in this paper.

We asked what gaps and potential future directions exist in the field of cognition in SE (RQ3). Our literature survey shows that while cognitive concepts have been studied in SE for more than 50 years, there is no single consensus on how to view or study cognition in this field. The field is organised mainly in terms of the SE activities that practitioners perform rather than in terms of the cognitive processes that are active while they perform them. Thus, the state of the art is specific to each area within SE and based on the traditions of that area. The advantage is that researchers address problems that are understandable within a taxonomy of SE concepts (e.g., SWEBOK)—such as program comprehension during software construction. The drawback is that the state of the art does not present an in-depth explanation of the cognitive processes that are active and that can lead to different outcomes in performing the activities.

In the following sections, we (i) discuss the research gaps in each cognitive concept and provide recommendations for future research, (ii) compare and industrial context and provide recommendations on the latter, (iii) provide exemplary use cases for using the taxonomy.

6.1 Gaps and Future Research Directions

Perception is the least studied concept. This may be the case because it commonly deals with making sense of and organising visual information. However, the predominant medium in software development is text, both in natural language and code.

We showed that *Perception* is investigated in the context of the SE tasks that make use of visuals (e.g., UI sketches in requirements engineering or architectural diagrams in software

design). Nowadays, making sense of and communicating the large amount of data collected at every step of the development life-cycle necessitates visual abstractions, for which *Perception* plays an important role. In particular, the areas of *Software Processes* and *Software Quality* fit the study of *Perception* due to their reliance on graphical information—e.g., notations used for process modelling or metrics visualisation used for software quality improvement.

Perception is inherently difficult to assess using unobtrusive techniques, such as open observations or fieldwork. Accordingly, we did not find any publication investigating this concept using these procedures. More intrusive approaches—e.g., physiological measures obtained using biometrics—can fit the study of *Perception*. For example, biometrics are widely used in the area of code comprehension (e.g., [52]) and can be adapted to investigate visual information in SE. However, biometrics are under-utilised for the study of *Perception*.

Future research directions for the *Perception* concept

- ★ Investigate *Perception* in areas of SE where visuals are used to summarise large amounts of data.
- ★ Two areas that can particularly benefit from the study of *Perception* are *Software Processes* and *Software Quality*.
- ★ We recommend employing *Physiological measures* for the study of this cognitive concept.

Knowledge is among the most studied concepts in the literature, and it covers several areas of SE. As *Knowledge* deals with how individuals incorporate new information and experiences in their work, this concept should be inherently present in every software development activity. Nevertheless, we show that *Knowledge* has not been investigated in large sub-areas, such as *Software Quality* and *Configuration Management*.

Studies of *Knowledge* in areas such as *Software Design* focus on how developers are able to acquire and use new information obtained from models like UML [14]. In other words, the knowledge acquired by a study subject is used as an outcome to evaluate an intervention (e.g., the use of new UML symbols). The study of *Knowledge* is transferable to other areas. For example, in the area of *Configuration Management*, one can study how a release recommendation system impacts the knowledge that different stakeholders have of the system to be deployed—e.g., in terms of requirements included in the release.

We found that *Knowledge* is assessed based on individuals performing tasks (i.e., *Task performance* and *Task-based*) as well as using their *Subjective rating*—these procedures are a natural fit for this cognitive concept. Other assessments procedures, such as *Fieldwork*, *Open observations*, *Physiological* and *Behavioural measurements* can be used to infer an individual's *Knowledge*. These assessments combined offer an alternative to approaches that necessitate the active involvement of the individual participating in a study, and can mitigate threats such as apprehension bias [47].

Future research directions for the *Knowledge* concept

- ★ All areas of SE can benefit from studying the *Knowledge* concept.
- ★ We especially recommend measuring *Knowledge*, as the outcome of an intervention, in the areas of *Software Quality* and *Configuration Management*.
- ★ We recommend investigating assessment procedures that can measure *Knowledge* without requiring explicit actions from the participants.

We showed that *Reasoning* has also been intensively studied in SE. Similarly to *Knowledge*, *Reasoning* can be investigated throughout different SE areas as it relates to decision making and problem solving. Although *Reasoning* is an implicit process (as opposed to *Knowledge* which is analytical in nature) it can be manipulated, and its outcome observed. For example, LaToza et al. [32] compared developers who could employ their own decision making strategy for debugging a program to developers who had to follow a given one.

Software quality and *Testing* are areas involving problem solving skills, whereas *Configuration management* and *Software processes* involve decision making. Nevertheless, these areas are understudied from a *Reasoning* perspective.

This cognitive concept is investigated using more direct assessment procedures, such as *Task performance* and *Subjective ratings*, which are proxies for the *Reasoning* process. Similarly to *Knowledge*, procedures that do not require an intervention on the individual are not broadly utilised.

Future research directions for the *Reasoning* concept

- ★ Although *Reasoning* is hard to observe as a direct outcome of a study, it can be manipulated in several tasks across SE areas, and future studies should seek to develop such research designs.
- ★ We recommend studying the effects of manipulating *Reasoning* specifically in the areas of *Software Quality*, *Testing*, *Configuration Management*, and *Software processes*.
- ★ We recommend using procedures which can make the *Reasoning* process explicit and that do not require direct measures—e.g., *Fieldwork* and *Open observation* accompanied by think-aloud protocol and *Interviews*.

Attention is being studied more and more in the last 20 years. As opposed to other concepts, such as *Knowledge* and *Reasoning*, its study may not fit process-related areas which are involved with cross-cutting concerns—e.g., *SE Processes*, *Software Quality*, and *SE Models & Methods*. *Attention* is observable over a short time span, which makes its study particularly suited for individual activities such as software *Construction* and *Maintenance*. For example, *Attention* plays a role in code comprehension where focusing on lexical similarities between code elements is important [25]. We showed that areas where individual attention is important for a task—e.g., *Testing* and *Configuration management*—need to be better investigated.

In the literature, *Attention* is generally assessed through quantitative procedures, such as *Subjective rating*, *Behavioural* and *Physiological* performance. Due to its transient nature, it is difficult to assess this concept in purely observational settings (i.e., using procedures such as *Fieldwork* and *Open observation*). Yet, some procedures such as *Interview* and *Task-based* assessment can uncover the motivation behind *Attention* deficit or prolonged attentional states [37].

Future research directions for the *Attention* concept

- ★ *Attention* is a transient phenomenon best investigated in SE areas requiring an individual to manage their focus in detail-oriented tasks.
- ★ We recommend studying the effects of different states of *Attention* in the areas of *Testing* and release planning as part of *Configuration Management*.
- ★ We recommend studying the motivation for changes in *Attention* using qualitative approaches, if possible.

In contrast to *Attention*, the study of *Memory* in SE decreased in the last 20 years. *Memory* has an inherent duality associated to it. Short term memory fits SE areas that includes more transitory tasks, such as code comprehension (in the area of *Construction*) and code reviews (in the area of *Maintenance*). Long-term memory fits tasks which unfold over a longer span of time, such as maintaining traceability links in the area of *Configuration management*. While the study of *Memory* has covered short-term memory—for example, Baum [5] showed the relationship between working memory and code review effectiveness—we showed that long-term memory needs more investigation.

Memory is a fundamental component in many complex tasks, and holds information and supports mental models that individual software developers rely on in their work. Similarly to *Attention*, *Memory* can be assessed using a wide range of quantitative procedures, but lacks consideration from a qualitative perspective—e.g., to identify the reason behind memory lapses,

consequences of recalling incorrect information, or to understand the background for reasoning and decision-making.

Future research directions for the *Memory* concept

- ★ We recommend studying *Memory* in areas for which long-term *Memory* is important (e.g., *Software processes* and *Software Models & Methods*)
- ★ We recommend supporting quantitative procedures with qualitative ones, such as *Interviews* and *Cognitive Task Analysis* procedures.

Cognitive load is one of the concepts which has constantly received attention from the SE research community starting from the 90s, particularly regarding code comprehension in the areas of software *Construction* and *Maintenance*. In both areas, developers use several strategies to understand the use of, for example, APIs or the behaviour of existing programs written by others (e.g., for code reviews). For instance, Siegmund [52] showed that beacons (i.e., semantic cues) are used to locate relevant code, and are associated with specific activation patterns in the brain.

Activities in the areas of *Configuration management*, such as understanding, constructing, and maintaining traceability links, or in *Software quality*, such as understanding and characterising defects, have been neglected. As *Cognitive load* is difficult to observe, researchers used proxy measures with procedures involving active individual participation (e.g., *Task performance* and *Subjective ratings*). Other procedures aim to quantify *Cognitive load* using *Physiological* and *Behavioural* measures. Although under-utilised, qualitative approaches are complementing the study of *Cognitive load* to elicit indicators of cognitive overload.

Future research directions for the *Cognitive load* concept

- ★ We recommend studying *Cognitive load* in areas other than the ones strictly related to software construction, such as *Configuration management* and *Software Quality*.
- ★ We recommend studying the reasons and consequences related to, and experience of cognitive overload using qualitative procedures, such as *Interviews*.

The research interest on *Cognitive biases* has been sustained over the last 20 years. Similarly to other concepts, we showed that research on *Cognitive biases* focused on the different phases of the development life-cycle (i.e., from *Requirements* to *Testing*). For example, Salman [44] investigated how confirmation bias takes place in test case design, whereas Mohanani [35] examined how fixation bias impacts creativity of requirements. In general, studies of *Cognitive biases* in these areas elicit a specific bias and measure its impact on a task relevant to that area. The remaining areas are understudied with the exception of *SE management*.

There are large gaps in the assessment procedures for *Cognitive biases*. As biases usually result in sub-optimal decisions when performing a task, researchers used *Task performance* and *Subjective ratings* to measure them. Yet, quantitative assessment based on *Physiological* or *Behavioural* measures are missing from the SE literature. Similarly to other concepts, it is inherently difficult to observe *Cognitive biases* and pin-point their effects. Nevertheless, qualitative procedures can support researchers in uncovering the causes of some cognitive biases, as well as evaluating interventions aimed at reducing their detrimental effects.

Future research directions for the *Cognitive bias* concept

- ★ We recommend studying *Cognitive biases* in areas where decision-making has cross-cutting impact, such as *Software requirements*, *Software Design*, and *Software processes*.
- ★ We recommend applying the procedures currently used to investigate *Cognitive biases* during the initial phases of the life-cycle to *Maintenance*, which is currently understudied.
- ★ We recommend applying assessment procedures used for the study of other concepts, such as *Physiological* and *Behavioural* measures. Qualitative approaches should be investigated as well.

Our study focused on cognition as seen from a psychological angle, extending to social cognition that still situates cognitive processes in individuals. However, cognition can also be viewed differently, as distributed across individuals, embedded and situated in artefacts and environments [29]. Although our result set includes a few papers which use distributed cognition theory as a backdrop, our study does not answer questions about research in SE using this conceptualisation. Future research could address this and contrast the taxonomy in this paper against new concepts found, potentially extending the taxonomy.

6.2 Cognitive Concepts Studies in Industry vs. Academia

Although both industry and academia have received similar attention in the literature, the study of cognitive concepts over the different areas of SE varies between these contexts. In both contexts, the focus is on the different phases of the development life-cycle, from *Requirements* to *Maintenance*—specifically on tasks related to software *Construction*. Notably, the investigation of *Knowledge* in the area of software *Construction* presents the most published studies in both academia and industry.

For selecting which area to focus on, context matters. Industrial studies are in fact predominant in areas of cross-cutting concerns for SE—from *SE management* to *SE economics*—which are difficult to attain in a purely academic context. As pointed out in Section 6.1, cognitive concepts that are investigated over an extensive time span (e.g., long-term *Memory*) are better suited for SE areas akin to Industrial context (e.g., *Software professional practice*). Similar examples are *Attention*, *Cognitive biases*, and *Cognitive load*—i.e., concepts studied mainly in an industrial context, and in SE areas representing cross-cutting concerns. Exceptions are the areas of *Configuration management* and *SE Economics* which are understudied despite fitting industrial investigations.

Future research directions in Industrial Context (SE areas)

- ★ Cross-cutting concerns should be investigated in industry studies, as the conditions for them are difficult to replicate in academic settings.
- ★ To investigate cognitive concepts over long time-spans, researchers should consider SE areas that capture the complexities of the relevant industrial context.
- ★ Sharply focused aspects within SE areas and clearly demarcated cognitive concepts (e.g., limited in time) can be investigated in academic settings, but care must be taken when attempting to transfer findings to industry.
- ★ Researchers should prioritise studies of cognitive concepts in areas suited for industrial investigation, such as *Configuration Management* and *SE Economics*, which are currently understudied.

The assessment measures observed in the literature reporting studies in academic context heavily favour quantitative approaches. *Fieldwork*—an approach which requires extensive presence of the researcher in the environment—is never utilised to carry out academic studies on cognitive concepts. In this context, researchers study *Knowledge* and *Reasoning* using qualitative approaches. Conversely, qualitative approaches are never used to investigate *Cognitive biases*.

In industrial studies, there is more balance between qualitative and quantitative approaches. Qualitative studies, such as *Fieldwork* and *Task-based* observations are widespread in this setting as they are not disruptive of the day-to-day activities, for example, practitioners working in a company. To some extent, this applies to quantitative assessments, too. Low-disruption and non-intervention approaches, namely *Task performance* (i.e., for a task which can be measured offline, such as defects found in a review) and *Subjective ratings* (i.e., in the form of quick surveys or experience sampling [26]), are preferred. More invasive procedures, which require ad-hoc devices, such as *Physiological measures*, are still uncommon. However, less invasive technologies have shown great potential in enabling biofeedback-based procedures in industry [62].

Future research directions in Industrial Context (Assessment procedures)

- ★ Qualitative procedures, currently used in industrial studies, need to be used more in academic studies as the latter present a context that can be an adequate yet cheaper approximation of the former [22] (e.g., in terms of recruiting subjects).
- ★ Qualitative and quantitative procedures should be chosen according to research objectives. We recommend the use of advanced mixed-methods research designs when investigating issues cutting across multiple SE areas, with potentially more than one cognitive concept, and where there are multiple research objectives (e.g., in terms of description, explanation, and theory-building).
- ★ We recommend considering unobtrusive technologies to unlock the use of *Physiological measures* for investigating cognitive concepts such as *Perception* and *Cognitive biases* in industrial context.

In general, we call for the adoption of new methods in cognition in SE research. Adopting methods brings with it the need to understand them and their underlying theories, which can advance the state of the art in the field. Physiological measurements, but also more advanced qualitative and mixed-methods approaches can make the research on cognition in SE more detailed, rich, and meaningful. A broader repertoire of methods would allow investigating cognition from a multitude of perspectives, and could lead to development of novel methods or variants of methods within SE research itself.

6.3 Uses for the Taxonomy

We argue that despite decades of research, there are plenty of opportunities for fruitful work in the area. Our development and application of a taxonomy of cognitive concepts and assessment procedures shows that it is possible to start organising the SE field in terms of cognition. Researchers and educators can use the taxonomy as a support for the following tasks.

Filtering existing studies. Researchers can apply the taxonomy to the primary studies covered in our SMS to obtain a list of papers about one or more cognitive concept of interest. This use case fits teams of SE researchers wanting to build background knowledge in an area of cognition and cross-disciplinary teams, who already have background knowledge in psychology, wanting to understand the research gaps relative to a cognitive concept applied to SE.

In a more pragmatic use case, SE methodology experts, wanting to investigate the area of cognition, can use the taxonomy to filter primary studies based on one of the assessment procedures to (i) identify gaps in the literature to steer their research agenda, and (ii) understand the nuances of applying a given procedure when investigating cognition. For educators, the two strategies above can be used to select reading material for a specialised course or module on cognition in SE, or for a general course or module on SE research methods.

Classifying new studies. This taxonomy represents a reusable classification scheme for researchers interested in performing secondary studies in the area of cognition in SE and close research areas where cognitive concepts play an important role, such as Human-Computer Interaction. SE Researchers can use the terms introduced in this taxonomy to better position their work in terms of cognitive concepts and assessment procedures. The taxonomy provides a set of keywords to be used together with the existing SE-specific ones (e.g., ACM CCS⁵). Depending on how specific a study is in terms of the cognitive concepts it investigates, it can be placed in a higher

⁵<https://dl.acm.org/ccs>.

(less specific) or lower (more specific) category. Furthermore, the relationships between cognitive concepts can also be used in classifying new studies. Many phenomena of interest in SE involve complex combinations of cognitive concepts. Studies investigating more than one concept can accordingly be classified into multiple categories in the taxonomy. Such patterns of relationships between cognitive concepts can further detail the focus of a specific study and be used to classify it.

Supporting researchers preparing primary studies in the area. The taxonomy can guide researchers in structuring papers so that important information on cognition is easy to find. Explicit sections need to be dedicated to the description of cognitive concepts of interest, their role in the study, and how they connect to the SE-specific activities. Similarly, the procedures need to be explicitly connected to the cognitive process they aim to assess. Finally, a paper needs to discuss the results in light of the covered cognitive concepts. Our experiences with extracting information from primary studies has shown us that authors seldom connect back to the cognitive concepts they originally set out to investigate.

7 CONCLUSION

In this article, we provided a taxonomy, a literature survey, and a gap analysis of cognitive concepts in SE. The taxonomy comprises the top-level cognitive concepts of perception, attention, memory, cognitive load, reasoning, cognitive biases, knowledge, and social cognition as well as qualitative and quantitative cognitive assessment procedures. The taxonomy provides a useful means to filtering existing studies, classifying new studies, and supporting researchers preparing studies in the area.

In the literature survey, we systematically analysed 311 primary studies published since 1973. We investigate the state of the art of cognitive concepts in SE based on the developed taxonomy and the SE Body of Knowledge. We showed that the most developed strands of research seem to be in the areas of the four software life-cycle stages, namely requirements, design, construction, and maintenance. Surprisingly, there is a clear lack of research on cognitive concepts in software testing. Most research is quantitative, which is more familiar to most software engineering researchers, and focuses on knowledge, cognitive load, memory, and reasoning. The state of the art appears fragmented when viewed from the perspective of cognition. There is a lack of use of cognitive concepts that would represent a coherent picture of the cognitive processes active in specific tasks.

We highlighted future research directions for each cognitive concept in the taxonomy. In the future, SE practice where cognition is more centrally considered could see tools that augment the skills of individual developers, techniques that help reduce bias and errors, and a workforce that is more aware of their own cognition as the basis for successful software design and development. The taxonomy and research synthesis presented in this paper give an overview of important cognitive concepts and assessment approaches that can help both researchers and practitioners take steps towards such goals.

REFERENCES

- [1] Apostolos Ampatzoglou, Stamatia Bibi, Paris Avgeriou, and Alexander Chatzigeorgiou. 2020. *Guidelines for Managing Threats to Validity of Secondary Studies in Software Engineering*. Springer, Switzerland, 415–441.
- [2] Apostolos Ampatzoglou, Stamatia Bibi, Paris Avgeriou, Marijn Verbeek, and Alexander Chatzigeorgiou. 2019. Identifying, categorizing and mitigating threats to validity in software engineering secondary studies. *Information and Software Technology* 106 (2019), 201–230.
- [3] Vaibhav Anu, Wenhua Hu, Jeffrey C. Carver, Gursimran S. Walia, and Gary Bradshaw. 2018. Development of a human error taxonomy for software requirements: A systematic literature review. *Information and Software Technology* 103 (2018), 112–124.
- [4] M. H. Ashcraft. 2002. *Cognition* (3rd Ed.). Prentice Hall, Upper Saddle River, NJ.

- [5] Tobias Baum, Kurt Schneider, and Alberto Bacchelli. 2019. Associating working memory capacity and code change ordering with code review performance. *Empirical Software Engineering* 24, 4 (2019), 1762–1798.
- [6] Sarah Beecham, Nathan Baddoo, Tracy Hall, Hugh Robinson, and Helen Sharp. 2008. Motivation in software engineering: A systematic literature review. *Information and Software Technology* 50, 9 (2008), 860–878.
- [7] Alan F. Blackwell, Marian Petre, and Luke Church. 2019. Fifty years of the psychology of programming. *International Journal of Human-Computer Studies* 131 (Nov. 2019), 52–63.
- [8] Pierre Bourque, Richard E. Fairley, et al. 2014. *Guide to the Software Engineering Body of Knowledge (SWEBOK): Version 3.0*. IEEE Computer Society Press, New York, NY.
- [9] John T. Cacioppo and Jean Decety. 2011. Social neuroscience: Challenges and opportunities in the study of complex behavior. *Annals of the New York Academy of Sciences* 1224, 1 (2011), 162–173.
- [10] Rafael Capilla, Anton Jansen, Antony Tang, Paris Avgeriou, and Muhammad Ali Babar. 2016. 10 years of software architecture knowledge management: Practice and future. *Journal of Systems and Software* 116 (2016), 191–205.
- [11] Andrew M. Colman. 2009. *A Dictionary of Psychology*. Oxford University Press, Oxford, England.
- [12] Neil S. Coulter. 1983. Software science and cognitive psychology. *IEEE Transactions on Software Engineering* SE-9, 2 (1983), 166–171.
- [13] Shirley Cruz, Fabio Q. B. da Silva, and Luiz Fernando Capretz. 2015. Forty years of research on personality in software engineering: A mapping study. *Computers in Human Behavior* 46 (May 2015), 94–113.
- [14] José A. Cruz-Lemus, Marcela Genero, M. Esperanza Manso, Sandro Morasca, and Mario Piattini. 2009. Assessing the understandability of UML statechart diagrams with composite states—A family of empirical studies. *Empirical Software Engineering* 14, 6 (2009), 685–719.
- [15] José Adson O. G. Cunha, Hermano P. Moura, and Francisco J. S. Vasconcellos. 2016. Decision-making in software project management: A systematic literature review. *Procedia Computer Science* 100 (Jan. 2016), 947–954.
- [16] Bill Curtis. 1984. Fifteen years of psychology in software engineering: Individual differences and cognitive science. In *Proceedings of the 7th International Conference on Software Engineering*. IEEE Press, Orlando, Florida, USA, 97–106.
- [17] B. Curtis, E. M. Soloway, R. E. Brooks, J. B. Black, K. Ehrlich, and H. R. Ramsey. 1986. Software psychology: The need for an interdisciplinary program. *Proc. IEEE* 74, 8 (Aug. 1986), 1092–1106.
- [18] Françoise Détienne. 1990. Program understanding and knowledge organization: The influence of acquired schemata. In *Cognitive Ergonomics: Understanding, Learning and Designing Human-Computer Interaction*. Academic Press, 245–256.
- [19] Françoise Détienne. 1997. Assessing the cognitive consequences of the object-oriented approach: A survey of empirical research on object-oriented design by individuals and teams. *Interacting with Computers* 9, 1 (Aug. 1997), 47–72.
- [20] Tore Dybå, Neil Maiden, and Robert Glass. 2014. The reflective software engineer: Reflective practice. *IEEE Software* 31, 4 (2014), 32–36.
- [21] Fabian Fagerholm, Michael Felderer, Davide Fucci, Michael Unterkalmsteiner, Bogdan Marculescu, Markus Martini, Lars Göran Wallgren Tengberg, Robert Feldt, Bettina Lehtelä, Balázs Nagyvárad, and Jehan Khattak. 2021. Data Set and Supplementary Files for Cognition in Software Engineering: A Taxonomy and Survey of a Half-Century of Research. (Oct. 2021). <https://doi.org/10.5281/zenodo.5845968>
- [22] Davide Falessi, Natalia Juristo, Claes Wohlin, Burak Turhan, Jürgen Münch, Andreas Jedlitschka, and Markku Oivo. 2018. Empirical software engineering experts on the use of students and professionals in experiments. *Empirical Software Engineering* 23, 1 (2018), 452–489.
- [23] Joseph L. Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological Bulletin* 76, 5 (1971), 378.
- [24] A. César C. França, Tatiana B. Gouveia, Pedro C. F. Santos, Celio A. Santana, and Fabio Q. B. da Silva. 2011. Motivation in software engineering: A systematic review update. In *15th Annual Conference on Evaluation Assessment in Software Engineering (EASE’11)*. Institution of Engineering and Technology, New York, NY, 154–163.
- [25] Thomas Fritz, Andrew Begel, Sebastian C. Müller, Serap Yigit-Elliott, and Manuela Züger. 2014. Using psychophysiological measures to assess task difficulty in software development. In *Proceedings of the 36th International Conference on Software Engineering*. Association for Computing Machinery, New York, NY, USA, 402–413.
- [26] Daniela Girardi, Filippo Lanubile, Nicole Novielli, and Alexander Serebrenik. 2021. Emotions and perceived productivity of software developers at the workplace. *IEEE Transactions on Software Engineering* (2021).
- [27] Lucian Gonçalves, Kleinner Farias, Bruno da Silva, and Jonathan Fessler. 2019. Measuring the cognitive load of software developers: A systematic mapping study. In *2019 IEEE/ACM 27th International Conference on Program Comprehension (ICPC)*. IEEE, 42–52.
- [28] Thomas R. G. Green. 1989. Cognitive dimensions of notations. *People and Computers V* (1989), 443–460.
- [29] Edwin Hutchins. 1995. *Cognition in the Wild*. MIT Press, Cambridge (Mass.).
- [30] Daniel Kahneman. 2011. *Thinking, Fast and Slow*. Macmillan, London, UK.

- [31] Elaine Kant and Allen Newell. 1984. Problem solving techniques for the design of algorithms. *Information Processing & Management* 20, 1 (1984), 97–118.
- [32] Thomas D. LaToza, Maryam Arab, Dastyni Loksa, and Amy J. Ko. 2020. Explicit programming strategies. *Empirical Software Engineering* 25, 4 (Mar. 2020), 2416–2449.
- [33] Per Lenberg, Robert Feldt, and Lars Göran Wallgren. 2015. Behavioral software engineering: A definition and systematic literature review. *Journal of Systems and Software* 107 (Sept. 2015), 15–37.
- [34] Stanley Letovsky and Elliot Soloway. 1986. Delocalized plans and program comprehension. *IEEE Software* 3, 3 (1986), 41.
- [35] Rahul Mohanani, Iftaah Salman, Burak Turhan, Pilar Rodríguez, and Paul Ralph. 2018. Cognitive biases in software engineering: A systematic mapping study. *IEEE Transactions on Software Engineering* 46, 12 (2018), 1318–1339.
- [36] Daniel Moody. 2009. The “physics” of notations: Toward a scientific basis for constructing visual notations in software engineering. *IEEE Transactions on Software Engineering* 35, 6 (2009), 756–779.
- [37] Sebastian C. Müller and Thomas Fritz. 2015. Stuck and frustrated or in flow and happy: Sensing developers’ emotions and progress. In *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, Vol. 1. IEEE, 688–699.
- [38] Chris Parnin and Spencer Rugaber. 2012. Programmer information needs after memory failure. In *2012 20th IEEE International Conference on Program Comprehension (ICPC)*. IEEE, 123–132.
- [39] Päivi Parviainen and Maarit Tihinen. 2014. Knowledge-related challenges and solutions in GSD. *Expert Systems* 31, 3 (2014), 253–266.
- [40] Luiz Pessoa. 2008. On the relationship between emotion and cognition. *Nature Reviews Neuroscience* 9, 2 (2 2008), 148–58.
- [41] Paul Ralph. 2019. Toward methodological guidelines for process theories and taxonomies in software engineering. *IEEE Transactions on Software Engineering* 45, 7 (2019), 712–735.
- [42] Maryam Razavian, Antony Tang, Rafael Capilla, and Patricia Lago. 2016. Reflective approach for software design decision making. In *2016 Qualitative Reasoning about Software Architectures (QRASA)*. IEEE, 19–26.
- [43] Mary Beth Rosson and Sherman R. Alpert. 1990. The cognitive consequences of object-oriented design. *Human-Computer Interaction* 5, 4 (1990), 345.
- [44] Iftaah Salman, Burak Turhan, and Sira Vegas. 2018. A controlled experiment on time pressure and confirmation bias in functional software testing. *Empirical Software Engineering* 24, 4 (Dec. 2018), 1727–1761.
- [45] C. Daniel Salzman and Stefano Fusi. 2010. Emotion, cognition, and mental state representation in amygdala and prefrontal cortex. *Annual Review of Neuroscience* 33, 1 (2010), 173–202.
- [46] Mary Sánchez-Gordón and Ricardo Colomo-Palacios. 2019. Taking the emotional pulse of software engineering—A systematic literature review of empirical studies. *Information and Software Technology* 115 (2019), 23–43.
- [47] Philip Sedgwick. 2012. What is recall bias? *BMJ* 344 (2012).
- [48] Zohreh Sharafi, Zéphyrin Soh, and Yann-Gaël Guéhéneuc. 2015. A systematic literature review on the usage of eye-tracking in software engineering. *Information and Software Technology* 67 (Nov. 2015), 79–107.
- [49] Beau A. Sheil. 1981. The psychological study of programming. *ACM Computing Surveys (CSUR)* 13, 1 (1981), 101–120.
- [50] Martin Shepperd. 2014. Cost prediction and software project management. In *Software Project Management in a Changing World*, Günther Ruhe and Claes Wohlin (Eds.). Springer, Berlin, 51–71.
- [51] Ben Shneiderman. 1976. Exploratory experiments in programmer behavior. *International Journal of Computer & Information Sciences* 5, 2 (1976), 123–143.
- [52] Janet Siegmund, Norman Peitek, Chris Parnin, Sven Apel, Johannes Hofmeister, Christian Kästner, Andrew Begel, Anja Bethmann, and André Brechmann. 2017. Measuring neural efficiency of program comprehension. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering (ESEC/FSE’17)*. Association for Computing Machinery, New York, NY, USA, 140–150.
- [53] Max E. Sime, Thomas R. G. Green, and D. J. Guest. 1973. Psychological evaluation of two conditional constructions used in computer languages. *International Journal of Man-Machine Studies* 5, 1 (1973), 105–113.
- [54] Herbert A. Simon. 1979. Information processing models of cognition. *Annual Review of Psychology* 30, 1 (1979), 363.
- [55] Elliot Soloway, Kate Ehrlich, and Jeffrey Bonar. 1982. Tapping into tacit programming knowledge. In *Proceedings of the 1982 Conference on Human Factors in Computing Systems*. 52–57.
- [56] Webb Stacy and Jean MacMillan. 1995. Cognitive bias in software engineering. *Commun. ACM* 38, 6 (June 1995), 57–63.
- [57] Muhammad Usman, Ricardo Britto, Jürgen Börstler, and Emilia Mendes. 2017. Taxonomies in software engineering: A systematic mapping study and a revised taxonomy development method. *Information and Software Technology* 85 (2017), 43–59.
- [58] Yingxu Wang. 2003. Cognitive informatics: A new transdisciplinary research field. *Brain and Mind* 4, 2 (2003), 115–127.
- [59] Gerald M. Weinberg. 1971. *The Psychology of Computer Programming*. Vol. 29. Van Nostrand Reinhold, New York.

- [60] Roel Wieringa, Neil Maiden, Nancy Mead, and Colette Rolland. 2005. Requirements engineering paper classification and evaluation criteria: A proposal and a discussion. *Requirements Engineering* 11, 1 (Nov. 2005), 102–107.
- [61] Claes Wohlin. 2014. Guidelines for snowballing in systematic literature studies and a replication in software engineering. In *Proceedings 18th International Conference on Evaluation and Assessment in Software Engineering (EASE)*. ACM, London, UK, 1–10.
- [62] Manuela Züger, Christopher Corley, André N. Meyer, Boyang Li, Thomas Fritz, David Shepherd, Vinay Augustine, Patrick Francis, Nicholas Kraft, and Will Snipes. 2017. Reducing interruptions at work: A large-scale field study of flowlight. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 61–72.

Received February 2021; revised December 2021; accepted December 2021