TECHNISCHE UNIVERSITÄT DRESDEN

ACADEMIC SKILLS IN COMPUTER SCIENCE

Lecture 2: Research and Publishing in Computer Science

Markus Krötzsch
Knowledge-Based Systems

TU Dresden, 9th April 2019

Goals for today

Learning goals of this lecture:

(1) Understand the research process in computer science
(2) Learn why and how research is published
(3) Distinguish essential types of publications
(4) Gain insights into peer reviewing, the main quality control mechanism

What is research?

Typical researchers are easy to recognise:

And in computer science?

What is research?

An OECD publication[1] gives the following definition (emphasis added):

> "Research and experimental development (R&D) comprise creative and systematic work undertaken in order to increase the stock of knowledge – including knowledge of humankind, culture and society – and to devise new applications of available knowledge."

and further derives the following key characteristics:

> "The activity must be:
> - novel [aimed at new findings, not known yet]
> - creative [based on original, non-obvious concepts/hypotheses]
> - uncertain [outcome and/or successfulness unknown]
> - systematic [planned & consciously managed; rigorous]
> - transferable and/or reproducible [results could be reproduced]."

[1] Frascati Manual 2015: Guidelines for Collecting and Reporting Data on Research and Experimental Development; doi:10.1787/9789264239012-en

What is research?

The OECD definition is intentionally broad, and distinguishes several types of research based on their general goals and motivation:

- basic research (curiosity-driven research)
- applied research (application-driven research)
- experimental development (product-driven research)

Note 1: OECD's definitions have some (unavoidable) imprecision
Note 2: There are other ways to define and classify research.

The research process

How does systematic knowledge generation look like in computer science?

The answer depends on how close we look:

- **Individual research questions** are often resolved by a sequence of steps
- **Zooming in:** Each step may follow some distinct methodology (e.g., empirical evaluation, mathematical proof, etc.)
- **Zooming out:** Research programmes and whole research fields undergo long-term developments, characterised by asking many research questions and adjusting methodologies over time

We first consider the middle level of individual research questions. Hustadt[1] describes three sequential research models in Computer Science, which we slightly adapt here:

- Theoretical research
- Experimental research
- Research through design

[1] Ullrich Hustadt: Professional Skills in Computer Science (COMP110), Lecture 6: Computer Science Research. University of Liverpool, 2016. Slides available online (link)

Sequential research: Theory

This is the predominant research process in mathematics and theoretical computer science:

Theoretical research
1. Define objects of study
2. State a conjecture
3. Construct a proof
4. Discuss and interpret the results

Example 2.1: Is P the same as NP?
1. Define complexity classes and other relevant mathematical notions
2. Conjecture "$P = NP$"
3. Proof by applying known mathematical identities
4. Discuss assumptions made in the proof (e.g., that $N = 1$ or $P = 0$); discuss further consequences

Sequential research: Experiment

This is the predominant research process in natural and social sciences:

Experimental research
1. Construct an initial theory (hypothesis, model)
2. Make a prediction based on the initial theory
3. Design and carry out experiments to test the prediction
4. Analyse and compare outcome of the experiments with prediction

Example 2.2: Linux vs. Windows
1. Hypothesis: Linux is more user-friendly than Microsoft Windows
2. Predict that Linux users perform better on routine, everyday tasks
3. Ask groups of Linux and Windows users to configure and build kernel 5.0-rc7 on their system of choice; use think-aloud protocols and measure times
4. Discuss findings and possible biases in experimental setup

Sequential research: Design

This is the predominant research process in engineering and organisational sciences:

Design-oriented research

1. State or derive requirements
2. Design a system
3. Implement and evaluate the system
4. Analyse results and draw conclusions about system design

Example 2.3: Machine learning (ML) on mobile devices

1. To train ML models on smartphones, we need energy-efficient algorithms
2. Design for a novel neuro-morphic, agile, embedded, adaptive, IoT-enabled single-layer perceptron
3. Implementation and evaluation w.r.t. prediction quality and energy use
4. Discuss findings and consequences for the design of future ML architectures

Theory vs. Experiment vs. Design

Note: Real computer science research works will often combine aspects of several research processes

Example 2.4: This could be the content of a single research work:

- Design a new algorithm and build a system using it
- Hypothesise that this new approach is inherently faster than the old way of doing things
- Give mathematical proof of better worst-case complexity properties
- Design and carry out experiments to validate if those theoretical gains a relevant in practical settings
- Discuss findings and results

The design (of software, hardware, algorithms, mathematical theories, and other artefacts) is a part of most CS research, and can be combined with many methods.[1]

[1] See also Hevner, March, Park, & Ram: Design Science in Information Systems Research. MIS Quarterly Vol. 28 No. 1, pp. 75-105/March 2004

More general research processes

Strict sequential research models are simplified abstractions

- Useful as blueprint for highly focussed activities (and presentations!)
- Fixed, pre-determined sequence of steps may need to be modified
- Not capturing how research works in larger terms

More elaborate models have been proposed:[1]

- Generalised research process models: Replace sequence by directed acyclic graphs to allow for alternative paths
- Circulatory research process models: Cyclic schemes that model how new findings feed back into earlier stages (e.g., to inspire changes in the hypothesis)
- Evolutionary research process models: Abstract model that considers that the steps performed in the (cycle of) research may change over time (e.g., if new methodologies become standard)

In practice, this is relevant mostly for theorising about research, not for doing it.

[1] See Christian W. Dawson. Projects in Computing and Information Systems: A Student's Guide. Second Edition, Addison Wesley 2009.

The Academic Publication Process

Why (not) publish?

What makes people publish research results?

- Sharing knowledge with other researchers and society as a whole
- Seek exchange with other researchers and solicit feedback
- Fame and recognition (publication pressure)
- Influencing opinion (promotion of own ideas & products)

What makes people not publish research results?

- Failure to find an outlet that wants to publish the results
- Retaining knowledge advantage (industry and military research may be classified)
- Lack of interest (no sufficient personal gain)
- Negative effects for own agenda or business (as in the case of ExxonMobil's research on climate change (link))

Conclusion: Some published research might be misleading, biased, irrelevant, and wrong – and some highly significant and relevant research might not be published.

Types of publications

Today, anybody can publish anything.

So what "counts" as a research publication?

The following basic types of text publications should be distinguished:

(1) **Formal research venues:** articles in journals and proceedings of research conferences, with established academic standards and rigorous quality control

(2) **Informal research venues:** proceedings of workshops, meeting notes, etc.

(3) **Monographs and collections:** books, including textbooks, and edited collections of invited research articles

(4) **Theses:** Texts written for obtaining an academic degree

(5) **Technical reports:** self-published research papers that may not have undergone any quality control yet, but are usually archived and stable

(6) **Other online texts:** blog posts and other web pages

(7) **Fake publications:** Fraudulent or pseudo-scientific texts that try to look like research

Ensuring quality and stability

Two key questions: (1) How has the quality of a work been ensured?

(2) Is the work archived in some permanent way?

	Quality control	Archiving
Formal research venues	Peer review, copy-editing	Publisher (typically library-indexed)
Informal research venues	Peer review/none	Publisher
Monographs and collections	None/friendly reviews, copy-editing	Publisher (typically library-indexed)
Theses	University process (varies widely)	University, libraries
Technical reports	None	Publisher or scientific archive
Other online texts	None	Publication site, archive.org
Fake publications	None	Publisher

Peer review vs. copy-editing

The predominant quality control mechanism in research is peer review:

Peer review:

- Manuscripts are submitted to a publication venue
- An editor/program chair asks experts to review the submission
- Based on the experts' opinions, the editor/program chair decides if the submission can be accepted

This scientific quality control is different from the copy-editing done by some publishers:

Copy-editing:

- Accepted manuscripts are finalised and sent to the publisher
- A trained copy-editor checks language and formatting issues
- Comments are sent to the authors or implemented directly by the copy-editor (but the authors should always give final approval)

The publication process in computer science (1)

Most computer science research is published at (small and large) conferences and workshops.

The general publication process for conferences is as follows:

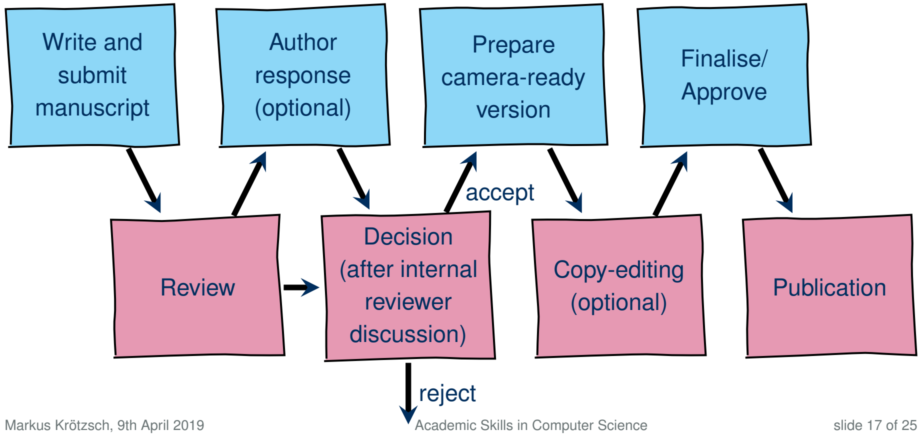

The publication process in computer science (2)

Journals adopt a slightly different process, since they publish more frequently and are therefore not bound to a firm timeline for acceptance.

The general publication process for journals is as follows (note the cycle):

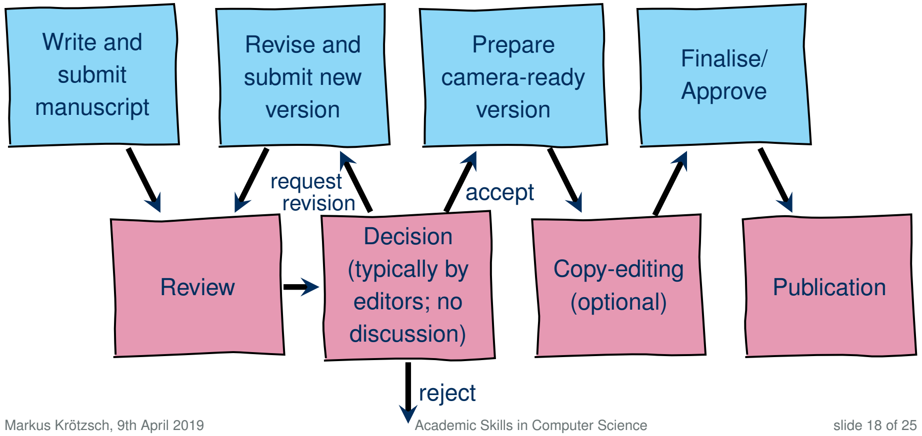

Conferences vs. journals in computer science

Conference proceedings

- Fixed timeline (paper has to be accepted before event happens)
- Length restrictions (to allow timely review)
- Short time to publication (typical: two months submission–acceptance)

Journal articles

- Open timeline (reviewers can ask for minor or major changes)
- Usually no length restrictions (to allow detailed, rigorous presentation)
- Possibly significant time to publication

Workshop, posters, short papers

- "Workshops" in computer science are often mini-conferences for preliminary works
- Some conferences also offer second-tier publication formats that do not get full articles in the proceedings (poster presentations, short papers, etc.)
- Emphasis of such events is on exchange; little relevance as publication venues

Attention: other academic fields have completely different publication cultures!

(Example: journal articles in the life-sciences are rather short and reviews are fast, similar to CS conference papers, but with the revision-based review process; life-science conferences are mostly for exchange and play little role as publication venues)

Peer review

The most widely used method of quality control across all of academia

Several basic models:

- Single-blind: reviewers are not known to the authors
  - Most common model today
  - Motive: protect reviewers from authors (who might be unhappy with verdict), prevent authors from trying to influence reviewers privately
- Double-blind: like single-blind, plus reviewers do not know author names either
  - Increasingly popular, but not always practical (esp. journal articles must discuss own prior works, thus revealing author identity)
  - Motive: increase chances for outsider authors, who might be reviewed more sceptically than the "big names" when presenting new ideas
- Non-blind (open): reviewers and authors know each other's identity
  - Implemented by some journals and conferences
  - Motive: make reviewers more accountable for their reviews; increase reviewing process transparency

How peer review works (1)

How are reviews organised?

- **Journals:** the managing editor invites experts to provide a review
- **Conferences:** the programme chair recruits a programme committee (PC) upfront; each PC member will be assigned several submissions to review
- **Large conferences:** sometimes use hierarchies of reviewers (extra roles include area chairs, track chairs, and senior PC members)
- **Book projects:** the editors organise reviews in whatever way they see fit

Note: reviewing is a community service and usually not paid for!

What are the results of a review?

- **Journals:** usually “accept”, “accept with minor revisions”, “request major revisions”, “reject with suggestion to resubmit”, “reject”
- **Conferences:** “accept” or “reject”, sometimes also “accept as short paper” or similar. Reviews may use more fine-grained scoring systems (example: “strong accept”, “accept”, “weak accept”, “borderline”, “weak reject”, “reject”, “strong reject”)

How peer review works (2)

How many reviews per submission?

Usually three or more; rarely just two or even just one; sometimes none (“desk reject”)

Who can be a reviewer?

- Any qualified expert
- Should have own research and publication experience (someone who never wrote a journal paper is not in the best position to tell others how to do it)
- Must not have a conflict of interest

What is a conflict of interest?

- Author is former Ph.D. supervisor, former Ph.D. student, family member, close friend
- Recently (past three years): collaborations with author, working in same organisational unit
- Conflicting commercial or academic interest
- Any other circumstance that prevents somebody from giving a fair and unbiased review

Reviewing criteria

Especially conferences often ask for reviews to evaluate several criteria for better comparability of submissions

Typical criteria include:

- Relevance: Does the contribution fit the conference/journal? Is it relevant to the research area it was submitted in?
- Significance: Are the results significant (big enough)? Does it advance our knowledge a lot?
- Originality: Is the work novel (new results, new methods, etc.)? Also compared to prior publications by the same authors.
- Correctness: Are the claims likely to be true? Are the proofs free of errors? Are the experimental designs sound? Are the conclusions valid?
- Presentation: Is the paper readable and clear?
- Related work: Does the work clarify how it compares to previous works in this area? Are all relevant references cited?

Reviewers are often asked to rate and comment on each dimension.

Summary

Several different research processes are common in computer science

Publications play a key role in making research results known

A rigorous and trustworthy quality control mechanism is essential to ensure that research publications are useful

Peer review is the most widely used quality mechanism

Publication is not a definite certificate of high quality – critical thinking is needed

What’s next?

- Gathering information in research
- Finding the most relevant literature and experts
- Reading academic papers

Image copyrights