



UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

What researchers need... ...and what they ask for

Liam Keegan, SSC

Talk Outline



- Introduction
- What researchers ask for
- What researchers need
- Examples
- Suggestions

Introduction

Who am I




I'm a Research Software Engineer at the Scientific Software Center of Heidelberg University. (→ see *talk by Inga Ulusoy*)

One of the things we do is offer a free consultation to researchers.

This means that researchers come to us from all kinds of disciplines, with all kinds of levels of technical background, and with all kinds of (more or less) software-related questions.

What this talk is about



Researchers often come to us with very specific technical questions.

And often, answering this specific question is the best way to help them...

But sometimes there is a different question they could have asked, with a much more helpful answer.

This talk is about finding that question.

Disclaimer



This is not an “expert” talk about established best practices.

The advice presented here is purely based on my own experiences.

And comments, questions, suggestions and alternative points of view are very welcome!

What researchers ask

A researchers journey



By the time they come to us with their technical question, they have probably made a lot of progress towards their goal.

In the process they have made a lot of decisions about technical aspects of their work, such as which programming language, which libraries and which algorithms to use.

The paths not taken



Some of these decisions they will make explicitly, after comparing different options.

But many decisions get made implicitly.

And sometimes they are not aware that there is even a decision to be made.

All this means that there are typically a lot of assumptions hidden behind their question.

The relevant missing context




It's generally not realistic to try to fully understand what they are trying to do, and all the choices they have made so far.

Even if it were, much of this is likely irrelevant to helping them with their issue.

But sometimes even a little additional context can lead to a much better solution.

What researchers need

What researchers need



Most of the time, they just need the answer to their question!

...but there are a few common scenarios where there is a better question waiting to be answered.

If all you have is a hammer



Sometimes they will use a tool because they are familiar with it, and you can point them towards a better alternative:

- A library that already does what they are trying to do
- Another programming language that is better suited for the task
- A more efficient algorithm, or one that is easier to parallelise

Seeing the forest beyond the trees



Sometimes they will focus on a particular solution, without realising that a different approach might be more effective:

- Move logic into the web frontend to avoid the need for a backend
- Provide Python bindings to users of your library instead of expecting them to use and compile c++

Examples


Example #1: Initial question

- Initial question:
 - *“How can we parallelise this Python code?”*
- Some possible answers:
 - Use multiprocessing
 - Use Numba
 - Use Dask
- But was there are better question?

Example #1: Improved question

- Initial question:
 - *“How can we parallelise this Python code?”*
- Searching for the missing context:
 - What is this code doing?
 - Is it part of something larger?
 - Why do you want to parallelise it?
- Improved question:
 - *“How can we do this analysis step faster?”*

Example #1: Result



- Ported the few critical lines of Python logic to C++
- Added a Python interface to fit into their existing workflow
- >50x single-core speed up
- + CPU/GPU parallelism

Example #2: Initial question

- Initial question:
 - *“Why is my Matlab function so slow?”*
- Possible answer:
 - What data do you run this function with?
 - Profile the CPU and RAM usage
 - It's running out of RAM and swapping to disk, making it slow
 - Could try to reduce the RAM usage, or just use a computer with more RAM
- But was there are better question?

Example #2: Improved question

- Initial question:
 - *“Why is my Matlab function so slow?”*
- Searching for the missing context:
 - What is this function doing?
 - How does it fit into the rest of the code?
 - Is this something you invented, or an implementation of an existing algorithm?
- Improved question:
 - *“How can I do XYZ efficiently in Matlab?”*

Example #2: Result

- There is a built-in Matlab function to do XYZ
- Replace entire function doing XYZ with a call to this built-in
- RAM and CPU usage drastically reduced

Example #3: Initial question

- Initial question:
 - *“Build and host a website with a database (urgent)?”*
- Some possible answers:
 - Quickly hack something together
 - Repurpose some existing similar project
 - Find an external service (that satisfies the data privacy requirements)
- But was there are better question?

Example #3: Improved question

- Initial question:
 - *“Build and host a website with a database (urgent)?”*
- Searching for the missing context:
 - What is this website doing?
 - What is the larger context, how does it fit into the application process?
 - What needs to happen to the data in the database?
- Improved question:
 - *“How can applicants fill in and submit a form?”*

Example #3: Result

- A simple static web page with a form applicants can fill in
- When finished they click download and get a file with their data
- They then upload this file along with others as part of their online application

Example #4: Initial question

- Initial question:
 - *“How to insert missing punctuation into sentences?”*
- Possible answers:
 - Find a NLP library that does text segmentation
 - Or find a machine learning model and fine tune it
- But was there are better question?

Example #4: Improved question

- Initial question:
 - *“How to insert missing punctuation into sentences?”*
- Searching for the missing context:
 - What language are these sentences in?
 - Where do these sentences come from?
 - What is the larger context, what is the source data and desired final output
- Improved question:
 - *“How can speech be translated and transcribed?”*

Example #4: Result

- (late 2022) Try this new open source model “whisper”
- Replace their entire pipeline with a single (much better) model

Suggestions

Suggestions

- Identify what it is they are trying to accomplish
 - Aim for a level of abstraction just above that of their question
- Identify what technical choices have already been made
 - In particular the implicit ones they may not be aware of making
- Look for possible better tools
 - Are they using a hammer for a task requiring a screwdriver?
- Look for possible better approaches
 - Are they stuck with a sub-optimal solution or approach?

Thank you



Comments or questions very welcome!

Image credits: generated using

- huggingface.co/stabilityai/stable-diffusion-3.5-large

with text prompts of the form

- *“a painting in the style of van Gogh of ...”*

With apologies to Vincent.