Investigating students' activities and approaches when solving Fermi problems using Fermi Graphs

Martin Purfürst and Jonas Bergman Ärlebäck
Department of Mathematics, Linköping University

This paper explores how a relatively new analytical tool called Fermi Graphs can be used to characterize and compare students' approaches to solving Fermi problems. Drawing on data from two groups of Swedish upper secondary students working on two Fermi problems, we examine the distribution and sequencing of students' activities throughout the solution process. The analysis shows that while both groups engaged in a wide range of activities, most frequently proposing assumptions and making guesstimations, their approaches differed structurally. One group worked sequentially through sub-problems, while the other adopted a strategy making assumptions before performing calculations. Fermi Graphs effectively captured these differences, offering a nuanced view of students' reasoning and problem decomposition strategies.

Introduction

In recent years, so-called *Fermi problems* have gained attention in mathematics education research for their potential to foster critical thinking, estimation skills, problem decomposition, and connection to mathematical modelling (Segura et al., 2025). As interest in their didactical potential grows, so does the need for robust analytical frameworks that can capture the nuances of students' problem-solving and modelling processes when working on Fermi problems. Existing frameworks such as *MADs—Modelling Activity Diagrams* (Ärlebäck, 2009) and *FPATs—Fermi Problem Activity Templates* (Albarracín & Ärlebäck, 2019) offer valuable perspectives; MAD emphasizes temporal sequencing of overarching modelling activities, while FPATs focus on task decomposition and data sourcing (Ärlebäck & Albarracín, 2019; 2024).

However, these frameworks do not necessarily fully capture the structural and logical complexity of students' solutions, especially when multiple solution paths and reasoning strategies are involved (Albarracín et al., 2025). In this context, a relatively new analytical tool called Fermi Graphs (Borys & Hartmann, 2022; Hartmann et al., 2020), offer a complementary approach by representing students' solutions as directed graphs that visualize the decomposition, recomposition, and flow of reasoning. This graphical representation allows for a more nuanced analysis of how students engage in various activities (such as proposing assumptions, validating results, and performing estimations) and how these activities are embedded within the overall structure of the solution (Hartmann et al., 2020).

This paper seeks to contribute with an empirical basis to enhance the theoretical discussion on analytical tools for analysing, understanding and mapping students' work on Fermi problems by exploring how Fermi Graphs can be used to characterize and compare students' solutions. Drawing on empirical data from two groups of Swedish upper secondary students solving two Fermi problems, we explore and examine the distribution and sequencing of the activities the students engage in. The research questions we address in the paper are: *What are the distributions of students' activities when solving the Fermi problems that can be discerned using a Fermi Graph analysis?* and *What characteristics of the problem solving approaches to the Fermi problems come to the fore using a Fermi Graph analysis?*

Theoretical considerations and selected prior research

Fermi problems, MADs and FPATs

Fermi problems are open-ended estimation tasks that require students to make assumptions, break down complex questions into simpler sub-problems, and use rough estimates to reach a plausible solution. Named after physicist Enrico Fermi, these problems typically lack precise data and encourage logical reasoning and quantitative thinking (Efthimiou & Llewellyn, 2006). In mathematics education research they have been defined as "open, non-standard problems requiring students to make assumptions about the problem situation and estimate relevant quantities before engaging in, often, simple calculations" (Ärlebäck, 2009, pp. 331–332). Their structure (decomposition, estimation, and synthesis) makes them similar to small-scale modelling tasks (Robinson, 2008), offering rich opportunities for mathematical thinking and discussion.

Several analytical tools have been developed to study students' engagement and work with Fermi problems. One such tool is the *Modelling Activity Diagram (MAD)*, used by Ärlebäck (2009) and Albarracín et al. (2019) to link students' work to the so-called modelling cycle (cf. Niss & Blum, 2020). MADs help identify which modelling activities students engage in during different stages of problem solving and how much time is spent on each. Unlike cyclical models, MAD presents the process linearly in terms of the six *modelling activities* Reading, Making Model, Estimating, Calculating, Validating, and Writing (Ärlebäck, 2009).

Another tool, the *Fermi problem Activity Templates (FPATs)*, has been used to analyse FP solutions and support task design (cf. Albarracín & Ärlebäck, 2024). FPATs highlight how students break down a problem, determine necessary data, and build toward a final solution. FPATs are particularly useful for comparing the final chosen solution strategies and examining how students generate numerical data. For the latter, Albarracín and Ärlebäck (2019) identify four methods for this: guesstimation, measurement and experimentation, information retrieval, and statistical data collection. Although MADs and FPATs provide useful information about students' overall working on, and solutions to, Fermi problems, the descriptions of students' work is rather coarse.

Fermi graphs

A less researched but more fine-grained description of the process is given by *Fermi Graphs* (see examples in Figure 2 and the Appendix), which visually represent a solution to a Fermi Problem, designed to emphasize the structure and complexity of the solution process in the form of a logical tree as a directed mathematical graph (Borys & Hartmann, 2022). Within this graph, numerical assumptions and intermediate results are noted along the edges, while the nodes capture how these data are processed, such as for example arithmetic operations when relevant (Jablonski et al., 2025). The processing of data at each node is also categorized using distinct *Fermi activities*, which describe the type of reasoning or calculation taking place (Hartmann et al., 2020). In some cases, a single node may involve more than one Fermi activity, reflecting the complexity of the students' reasoning. The solution process is read from left to right (giving the directional aspect of the graph), but unlike MADs, does not include an explicit timeline. Instead, it illustrates the structural flow of the solution process and Fermi activities engaged in, allowing for a multi-tiered analysis of how students decompose and reconstruct the problem more in line with a FPAT (Albarracín et al., 2025). In this paper we investigate the potential of Fermi Graphs to get a more detailed and nuanced description of students work as an analytical complement to MADs and FPATs.

Participants, data collection and analysis

We now turn to describe the participants of the study, the data and the data collection methods used, and how the data was analysed in the study [1].

Participants

Seven upper secondary school students from southern Sweden, divided into two groups based on their availability for data collection, participated in this study. Group A consisted of three male students (Alf, Ali, and Ari) and group B consisted of three male students (Bill, Ben, and Bo) and one female student (Bea). These names are pseudonyms to allow the students to remain anonymous. The students are all high performing math majors, but they had no explicit experience of solving Fermi problems prior to this study.

Methods and data collection

This study was conducted as a case study (Yin, 2014) in which two student groups' work on two Fermi problems was observed and analysed. The entire solution process of both groups was video recorded, and the students' utterances were transcribed verbatim. Non-verbal forms of communication, such as gestures or other forms of communication, were noted in the transcription using square brackets ("[]"), and instances when a student's utterance was unintelligible was marked as such using a regular parenthesis. A short break in a student's utterance, or instances when a student was trailing off (e.g. when a student was interrupted mid-sentence, or when a student paused to ponder what to say next), was transcribed using three dots ("…").

Two Fermi problems were given to the students to solve in this study: *The Google and ChatGPT problem*, and *The wrapping paper problem*. These two tasks are presented in Figure 1. The problems were presented in Swedish and were translations of two original Fermi Problems developed within the EU-funded Dim^{2}ension project [2] by Jonas Bergman Ärlebäck and Simon Barlovits respectively.

Many people like to wrap presents. Especially at Christmas, many gifts are wrapped in gift paper to decorate the gift. Wrapped gifts may look beautiful, but wrapping paper has two main problems.

Firstly, wrapping paper is a product that is usually used once and then thrown away after unwrapping. Secondly, wrapping paper is usually plastic-coated paper. This type of wrapping paper looks very decorative, but it cannot be recycled. Only wrapping paper without a plastic coating (i.e. "real" paper) or plastic film can be recycled.

These two problems lead to the following question: How many tonnes of non-recyclable wrapping paper is used in the European Union every year at Christmas?

A single Google search uses approximately 0.0003 kWh of energy. This is enough energy to power a 60-watt light bulb for about 17 seconds, or the equivalent to 0.2 grams of CO2.

Similarly, each ChatGPT query uses approximately 0.0005 kWh of energy.

How much energy is used by Google searches each day? By ChatGPT queries? What are the corresponding CO2 emissions?

Compare your answers with the facts that:

- the average household in the EU uses approximately 10 kWh of electricity per day; and that
- a flight between London and New York emits 1 590 kg of CO_2.

Figure 1. The two Fermi problems used in this study.

Both groups were first given *The Google and ChatGPT problem* to work on. They were instructed to “think aloud” and to provide a written solution accounting for all their assumptions, estimates, and calculations. The written solution was collected after the students were finished with the task. After a brief pause the students were then given *The wrapping paper problem* to work on in the same manner. The first author of this paper was present in the room as the students were working on the problems, but did not partake in the groups’ solution processes.

Analysis

Based on the prior research literature on Fermi Graphs, MADs and FPATs, a list of Fermi activities was compiled to be used as an analytical lens (cf. Purfürst, 2025). These activities are presented in Table 1 below:

Table 1. A short description of all eleven Fermi activities used in this study.

Fermi activity	**Description**
Common Knowledge (CK)	The students make use of information or knowledge that in context can be considered common knowledge.
Propose Assumption (PA)	The students propose an assumption or estimation.
Set Assumption (SA)	Previously proposed assumptions are set/confirmed and used throughout the rest of the solution process.
Guesstimation (Ge)	An assumption or estimate in the form of a *“qualified guess”*.
Collect/Split up (C/Sp)	Previous assumptions or sub-results are collected/split up for further use.

Validate (Va)	The students interpret or critically scrutinize an assumption or (sub-)result.
Unit Conversion (UC)	The students perform a simple unit conversion.
Get from Text (GT)	The students get, and use, information from the task description.
Calculate (Ca)	The students perform an exact calculation without rounding or using approximate values.
Rough Calculation (RC)	The students perform a rough/approximate calculation using rounding or approximations.
Adjust/Simplify (A/S)	The students adjust or simplify a (sub-)result or previously set assumption.

For each student group and task, a Fermi graph was created. An initial draft of each graph was based on the groups' written solution, in which different solutions' steps were identified and categorized as instantiations of distinct Fermi activities. Once the basic structure of each group's Fermi graph was established, the transcriptions of their work on the tasks were used to provide a more detailed representation of the solution process, gradually extending and revising the initial graphs.

The transcriptions not only allowed us to establish the sequential aspect of the solution (i.e. in what order the Fermi activities were carried out), but also provided insight into the assumptions and estimates the students did not account for in their written solutions. Different parts of the transcriptions, in which the groups were considered to progress in the solution, were categorized as different Fermi activities. To illustrate; in the following short excerpt, we see Ben proposing an assumption about how many times a day an average person uses Google based on a guesstimation (PA and Ge; #1). Bea and Bo confirm this assumption (SA; #2–4), then Ben and Bill propose an assumption about the Earth's population based on common knowledge (PA and CK; #5–#6), and the whole group confirms this assumption (SA; #7):

#1, Ben: I think it's at least 10… if you take an average.

#2, Bo: Should we… Average people Google 10 times a day, then we can just strike out a zero here. [Writes something] … (Unintelligible) times the amount of people.

#3, Bea: Yes, 10 feels pretty reasonable, I think.

#4, Bo: Yes, and it… How many people are we here on earth?

#5, Ben: 8–9?

#6, Bill: 8 billion. That was around a year ago so it should…

#7, Everyone: Yes…

From this short excerpt we can identify several Fermi activities and clearly establish in what order they were carried out in the solution process. This type of categorization was done for all instances in the transcribed material where the students could be considered to engage in activities or discussions related to a progression in their solution.

After incorporating the activities categorized in the transcriptions into the Fermi graphs, the distribution of each student group's activities related to each task was summarized in a table. A complete distribution of activities (i.e. adding all activities for each student group) was also created. Finally, the Fermi graphs for each task was compared from a structural perspective to identify qualitative differences in the solution strategies between the two groups. The focus of this part of the analysis was to highlight different approaches the students chose to use to solve the two problems. In particular, we focused on contrasting what type of Fermi activities the students mainly were engaged in at different times of the solution process, as well as identifying the types of sub-problems the students decomposed the problems into.

Results

Next, we present the results from the analysis. We first provide an overview of the two groups' activity distributions. After this we present a qualitative comparison of the groups' solution strategies for each task based on the Fermi graphs created.

The distribution of students enacted Fermi activities

Both student groups employed a wide range of Fermi activities when solving the two Fermi problems. The most frequently observed activity in both groups was *Propose Assumption (PA)*, followed by *Guesstimation (Ge)*. Group A also frequently engaged in the *Set Assumption activity (SA)*, whereas Group B did not to the same extent. However, Group B participated in the *Calculate* activity *(Ca)* more often than Group A. A complete and more detailed summary of the distribution of activities across both student groups and problems is presented in Table 2 below:

Table 2. A summary of the distribution of both student groups' Fermi activities when solving the two Fermi problems.

Task & group	**CK**	**PA**	**SA**	**Ge**	**C/Sp**	**Va**	**UC**	**GT**	**Ca**	**RC**	**A/S**
Google & ChatGPT, Group A	1	11	8	10	5	4	0	2	4	1	3
Wrapping paper, Group A	5	21	12	16	6	1	1	0	7	3	2
Total, Group A	**6**	**32**	**20**	**26**	**11**	**5**	**1**	**2**	**11**	**4**	**5**
Google & ChatGPT, Group B	1	10	3	9	3	7	2	7	10	1	1
Wrapping paper, Group B	3	18	6	15	4	2	1	0	5	2	2
Total, Group B	**4**	**28**	**9**	**24**	**7**	**9**	**3**	**7**	**15**	**3**	**3**

Comparing and contrasting the groups' work using Fermi graphs

We now turn to a comparison of the two student groups' work on the tasks, using the constructed Fermi graphs as the basis for our analysis. The primary focus here is on structural differences in how the groups approached the problems. Specifically, we examine the types of Fermi activities the students engaged in at various stages of the solution process, as well as the kinds of sub-problems they chose to decompose the tasks into. The Fermi graphs representing each group's work on the tasks are presented in Figure 2 below. A larger version of Fermi graph c) is available in the Appendix.

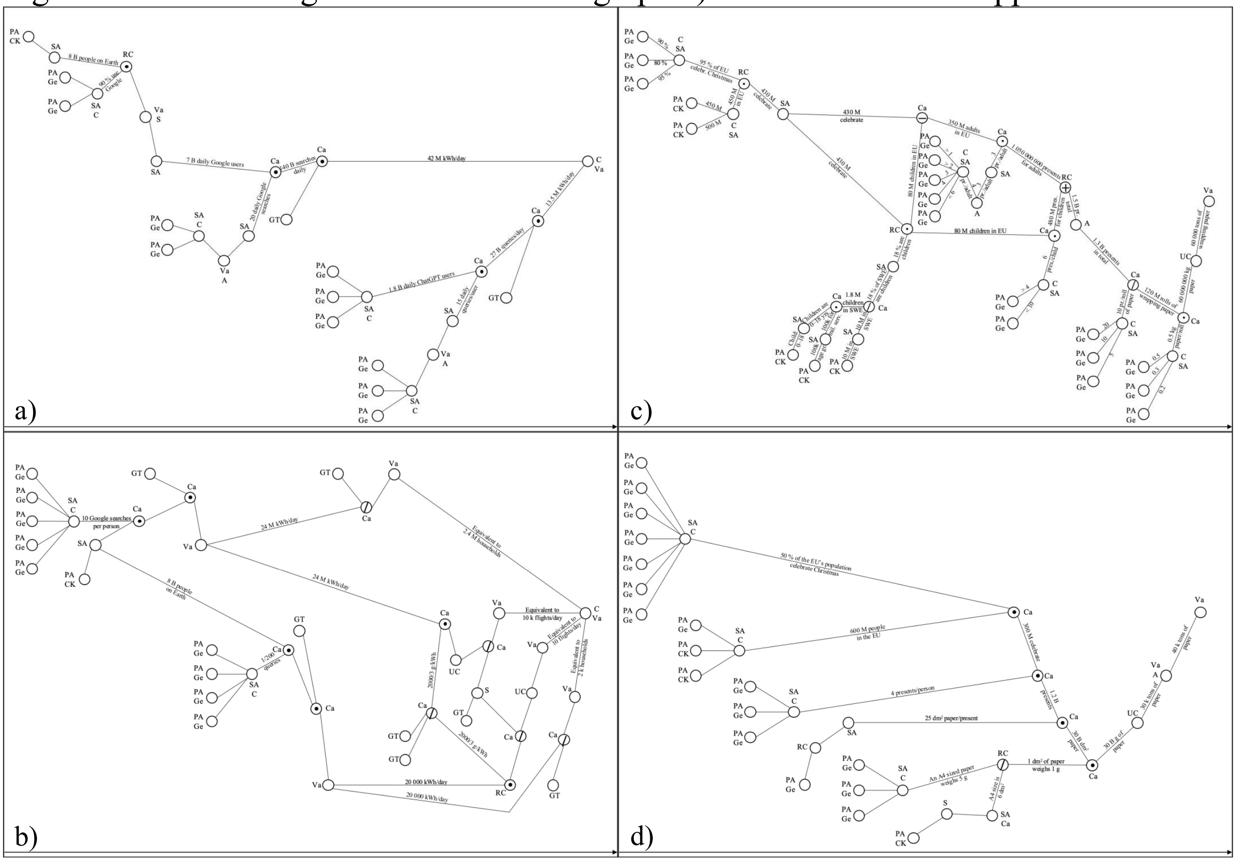

Figure 2. Fermi graphs of the groups' solution processes. a) and b) depicts group A and B's solutions to *The Google and ChatGPT problem* respectively, and c) and d) depicts the groups' solutions to *The wrapping paper problem*.

The overall structure of the Fermi graphs in Figure 2 illustrate that the two groups approached the tasks in distinctly different ways. By comparing the structure of the Fermi graphs, it is evident that group A (graphs a) and c)) generally solved one sub-problem at a time before arriving at a final solution. In contrast, Group B (graphs b) and d)) first proposed and set all assumptions they considered necessary, and then primarily engaged in calculating activities, effectively solving their (implicit) sub-problems.

These contrasting approaches are particularly evident in the groups' solutions to *The Wrapping Paper Problem* (graphs c) and d) in Figure 2). Group A displays a relatively even distribution of Fermi activities throughout the solution process, solving one sub-problem at a time, each involving a range of activities. In contrast, Group B primarily engaged in activities related to proposing and setting assumptions before transitioning to calculating activities, once all necessary assumptions had been established. A similar structural approach can be observed in the groups' work on *The Google and ChatGPT Problem*, although the comparison is less clear in this case, as Group A did not fully answer the posed questions.

Conclusion and discussion

The analysis presented shows that both student groups engaged in all Fermi activities listed in Table 1 when solving the two problems presented in this study, with Propose Assumption (PA) and Guesstimation (Ge) being the most frequently observed. Notably, although the overall distributions of activities were similar across the groups, their approaches to solving the Fermi problems differed structurally. Group A addressed the tasks by working through one sub-problem at a time, engaging in a variety of activities throughout the process. In contrast, Group B adopted *a mise en place* approach, first identifying and setting all necessary assumptions and estimations before proceeding with calculating activities based on those predefined assumptions.

The applied theoretical framework, Fermi graphs, proved effective in highlighting the distinct flows of reasoning within each student group. When considering the sequential aspects of the students' solution processes, the graphs clearly revealed contrasting approaches between the groups. As used in this study, the Fermi graphs provided insight into the overarching structure of each group's solution strategy and effectively illuminated the different types of sub-problems the students chose to address.

Contrasting a Fermi Graph's description of a solution process with representations from MAD or FPAT analysis, Fermi Graphs generalize the *modelling activities* of MADs by introducing more nuanced *Fermi activities*, and offer a clearer overview of problem decomposition strategies. In this sense, Fermi Graphs provide a hybrid representation that captures both the logical structure and flow of reasoning. Meanwhile, MADs and FPATs retain strengths in their accessibility and in representing the key temporal progression and actual solution of a Fermi problem. Together, these three tools offer complementary lenses for analysing students' engagement with Fermi problems: Fermi Graphs excel in mapping reasoning structure, MADs in highlighting activity sequences, and FPATs in clarifying (the final) decomposition of the Fermi problem and data sourcing. Future research could explore hybrid analytical models that integrate these perspectives for a more holistic understanding of students' modelling processes.

In addition, we observed that throughout the solution processes, both groups engaged in various validating activities that could not adequately be represented in the Fermi graphs. To fully capture the logic of a group's solution strategy the students' various validating activities should be taken into consideration. To accomplish this, future researchers could consider introducing additional and more nuanced Fermi activities connected to students' validating activities; or using other analytical tools (i.e. MADs or FPATs) as a complement to the Fermi graphs (cf. Albarracín et al., 2025).

Notes

1. This work originate from, and develops the ideas in, the student thesis by the first author (Purfürst, 2025) carried out under the supervision of the second author.

2. DIM^{2}ENSION–**DI**gital Supported **M**athematical **M**od**E**lli**N**g For **S**ustainable Development Goals **I**n Eur**O**pean Educatio**N**–is an EU-funded research project (2024-1-DE01-KA220-HED-000245297) aimed at integrating digital tools into teaching about (and with) mathematical modelling related to the Sustainable Development Goals. https://dimension-project.eu

References

Albarracín, L., & Ärlebäck, J. (2019). Characterising mathematical activities promoted by Fermi problems. *For the Learning of Mathematics*, *39*(3), 10–13. https://www.jstor.org/stable/26854427

Albarracin, L., Ärlebäck, J., Civil, E., & Gorgorio, N. (2019). Extending modelling activity diagrams as a tool to characterise mathematical modelling processes. *The Mathematics Enthusiast, 16*(1), 211–230. https://doi.org/10.54870/1551-3440.1455

Albarracín, L., Ärlebäck, J. B., Okamoto, H., Hartmann, M., & Borys, T. (2025). *Tools for analysing student behaviour when solving Fermi problems: Insights and future developments* [Manuscript submitted for publication].

Ärlebäck, J. B. (2009). On the use of realistic fermi problems for introducing mathematical modelling in school. *The Mathematics Enthusiast*, *6*(3), 331–364. https://doi.org/10.54870/1551-3440.1157

Ärlebäck, J. B., & Albarracín, L. (2019). The use and potential of Fermi problems in the STEM disciplines to support the development of twenty-first century competencies. *ZDM Mathematics Education, 51*(6), 979–990. https://doi.org/10.1007/s11858-019-01075-3

Ärlebäck, J. B., & Albarracín, L. (2024). Fermi problems as a hub for task design in mathematics and STEM education. *Teaching Mathematics and its Applications: An International Journal of the IMA*, *60*(7), 25–37. https://doi.org/10.1093/teamat/hrad002

Borys, T., & Hartmann, M. (2022). Der „Fermi-Graph". Illustrative Rekonstruktion der Lösungsprozesse von Fermi-Aufgaben. In T. Borys, M. Brandl, & A. Brinkmann (Eds.), *Mathevernetzt–Anregungen und Materialien für einen vernetzenden Mathematikunterricht* (pp. 26–32). Verlag Bücherbunt im MUED e.V.

Efthimiou, C. J., & Llewellyn, R. A. (2006). Cinema, Fermi problems and general education. *Physics Education*, *42*(3). https://doi.org/10.1088/0031-9120/42/3/003

Hartmann, M., Borys, T., Okamoto, H., & Kawasaki, T. (2020). *Der „Fermi-Tree" – Prozesse und Strukturen beim Lösen von Fermi-Fragen erfassen*. https://doi-org.e.bibl.liu.se/10.17877/DE290R-21348

Jablonski, S., Ehemann, J., Deng-Xin, K., & Ludwig M. (2025). One step at a time? –An empirical study on the use of graphical representations for solving Fermi problems. *Proceedings of CERME14*. hal-05201444

Niss, M., & Blum, W. (2020). *The teaching and learning of mathematical modelling.* Routledge.

Purfürst, M. (2025). *Fermiproblem för hållbar utveckling: En analys av svenska gymnasieelevers arbete med Fermiproblem kopplade till FN:s globala mål för hållbar utveckling* (Master's thesis, Linköping University). DiVA Portal. https://liu.diva-portal.org/smash/record.jsf?pid=diva2:1968712

Robinson, A. W. (2008). Don't just stand there—teach Fermi problems! *Physics Education*, *43*(1), 83-87. https://doi.org/10.1088/0031-9120/43/01/009

Segura, C., Gallart, C., & Ferrando, I. (2025). Influence of pre-service primary school teachers' prior knowledge of measurement and measurement estimation in solving modelling problems. *Journal of Mathematics Teacher Educ*ation. http://doi.org/10.1007/s10857-025-09685-3

Yin, R. K. (2014). *Case Study Research Design and Methods* (5th ed.). Sage.

Appendix – Fermi graph of group A's work on *The wrapping paper problem*