

# Reliance on Artificial Intelligence Tools May Displace Research Skills Acquisition Within Engineering Doctoral Programmes: Examples and Implications

Yevhenii Mormul, Jan Przybyszewski, Andrew Nakoud and Paul Cuffe, *Senior Member, IEEE*

**Abstract**—The escalation in capabilities of Large Language Models has triggered urgent discussions about their implications for tertiary education, particularly regarding how they might facilitate academic misconduct in graded engineering coursework. However, graduate research education — where a student works closely with a supervisor over years to develop both implicit and explicit research skills — has received comparatively less attention in this discussion. This paper seeks to develop this discourse by presenting targeted case studies that explore the opportunities and threats posed by artificial intelligence to engineering doctoral education. For instance, using a specimen exercise from a PhD-level research skills module, we demonstrate how artificial intelligence tools can now deeply penetrate research workflows in technical computing and scripting. We likewise investigate the capabilities of chatbot tools to assist engineering PhD candidates with the broader research skills central to their training and development. These include writing and proofreading theses and research papers, producing data visualizations, simulating peer review processes, and preparing scientific diagrams. By evaluating the capabilities and limitations of extant artificial intelligence in these areas, we can discuss both the potential benefits and ethical concerns of doctoral students engaging with such assistance.

## I. INTRODUCTION

The emergence of consumer chatbot tools with sophisticated Artificial Intelligence (AI) capabilities [1] has provoked considerable debate about their implications for assessment integrity in higher education [2], [3]. Much of the extant literature focuses on undergraduate education [4], [5].

What, though, does this ongoing technical upheaval portend for the research skills that students learn, explicitly and implicitly, while enrolled in graduate research programmes [6], [7]? For instance, one recent report claims to have built an AI tool that “*generates novel research ideas, writes code, executes experiments, visualizes results, describes its findings by writing a full scientific paper, and then runs a simulated review process for evaluation*” [8]. Experimental work in [9] found that “*college-educated professionals performing midlevel professional writing tasks substantially increased their productivity when given access to ChatGPT.*” Likewise for computer programming, one AI assistant’s suggestions were deemed “*comparable to humans’ contributions in terms of quality*” [10]. Less encouragingly, other recent reports note that AI-generated manuscripts, of dubious quality and provenance, are already flowing into peer reviewed outlets [11].

---

Y. Mormul, J. Przybyszewski and P. Cuffe (paul.cuffe@ucd.ie) are with the School of Electrical and Electronic Engineering, University College Dublin. A. Nakoud is with the University Of California, San Diego. This work is funded by the Irish Exchequer, via a grant to UCD by the Higher Education Authority’s Strategic Alignment of Teaching and Learning Enhancement scheme.

It is timely to question if an over-reliance on powerful AI tools might impede the acquisition of research skills by doctoral students in engineering. How should doctoral programmes in engineering evolve [12] now that the conduct and dissemination of primary research itself can increasingly be automated using sophisticated software tools [13]?

In this paper we investigate the abilities of a prominent Large Language Model in supporting doctoral candidates, and discuss the educational and ethical implications of it doing so. Specifically, we gauge this general purpose AI chatbot’s proficiency in assisting doctoral students with:

- computer scripting to visualise technical data [14]–[16]
- constructing scientific diagrams [17]
- typesetting numerical tables [18]
- reviewing literature and summarising primary research [19], [20]
- simulating manuscript peer review and conducting structured proofreading [21]

If each of these specific scholarly tasks can now be offloaded to a chatbot tool, how can we remain assured that doctoral students will actually master granular and transferable research skills over the course of their programme?

## II. METHODOLOGY

To investigate the potential impact of AI chatbot tools on research skills acquisition within engineering doctoral programmes, a series of tasks mirroring typical research workflows activities has been designed. These tasks will be presented to *Claude 3.5 Sonnet*, an advanced AI model selected for its logical reasoning and visual comprehension abilities [22]. The methodology aims to loosely simulate the research process a PhD student might undertake, feeding into an initial subjective evaluation of the AI’s performance across the breadth of granular skills embodied in graduate-level research.

For each task, a qualitative, interactive dialogue with the AI will be engaged in, mirroring the iterative nature of research work and the flexible, freeform work styles permitted to doctoral students.

### A. Gauging Capabilities of Artificial Intelligence

1) *Scripting to visualise technical data*: A structured data visualization and computer scripting task will be provided to the AI, extracted from an existing induction module designed for PhD students. This module explicitly teaches and assesses research skills related to data visualization and analysis. The task utilizes a real-world dataset from a technical domain, mirroring the type of challenge doctoral students encounter in their research. The AI will be prompted to write a full computer script to meet the task. This exercise will show the AI’s capability to:

- Translating complex natural language requirements to code
- Produce publication-ready graphics

2) *Constructing scientific diagrams*: The AI will be tasked with creating a complex diagram to illustrate a theoretical concept in engineering. This task will illustrate the AI’s ability to:

- Interpret verbal descriptions of complex concepts
- Suggest appropriate diagramming approaches
- Generate code (e.g., TikZ or Graphviz) to produce such diagrams

3) *Typesetting numerical tables*: Academic typesetting, particularly using  $\LaTeX$  is often a time-consuming and distracting task for doctoral students [23], [24]. To assess the AI’s potential to alleviate this burden, a raster image of a complex numerical table will be presented to the AI, along with a request to transcribe it into a publication-quality  $\LaTeX$  format. The AI’s ability will be assessed to:

- Accurately interpret and transcribe the visual data from the image
- Propose suitable  $\LaTeX$  table structures
- Generate  $\LaTeX$  code for the table
- Apply appropriate formatting for readability and adherence to typical journal standards

4) *Literature review and knowledge summary*: To assess the AI’s capability in comprehending and synthesizing complex academic content, a task crucial for doctoral students, multiple full academic papers from a specific engineering domain will be presented to the AI. This task simulates a key aspect of the literature review process, where students must efficiently extract and synthesize information from numerous sources. This will clarify the AI’s ability to:

- Provide accurate and concise summaries of technical paper, extracting key findings and methodologies
- Synthesize information to provide a coherent overview of the state of research in a field
- Highlight potential gaps or areas for future research based on the collective content of the papers

5) *Structured proofreading and prognostic simulation of peer review*: Developing the ability to anticipate and address potential critiques is a crucial skill for doctoral students in engineering. The prognostic “*simulation*” of how a manuscript may fare in peer review represents an important mental model for research students to internalize, enabling them to refine their work proactively and enhance their chances of successful publication. This skill traditionally develops through experience and mentorship over the course of a doctoral program.

To evaluate the AI’s capability in simulating this critical thought process, a complete academic paper in an engineering domain will be presented to the AI. The AI will be tasked with:

- Providing a comprehensive critical analysis of the paper, including its strengths and weaknesses
- Suggesting potential questions or concerns that peer reviewers might raise, covering aspects such as methodology, results interpretation, and significance of findings
- Offering strategies to address these hypothetical concerns, including potential revisions or additional experiments
- Assessing the overall likelihood of the paper being accepted for publication in its current form

### B. Holistic Assessment of AI Capabilities in Doctoral Research

This multifaceted methodology aims to gauge the extent to which AI tools can replicate—or potentially surpass—research skills traditionally honed during doctoral studies. By subjecting an advanced AI to tasks spanning data visualization, academic writing, peer review simulation, and literature synthesis, we seek to illuminate the potential reshaping of engineering doctoral education. Our analysis will weigh AI’s promise as an assistive tool against the risk of skill displacement, contributing to the ongoing dialogue on integrating AI in higher education. Ultimately, this investigation aspires to offer crucial insights for adapting doctoral programs to an AI-augmented academic landscape, balancing technological assistance with the imperative of fostering robust, independent research capabilities in future engineering scholars.

## III. RESULTS

Full transcripts of the prompting interactions are archived at a persistent online repository [25]. Due to space limitations, only selected extracts of Claude’s outputs can be reproduced within this manuscript.

### A. Data Visualisation and Technical Computing

Creating programming scripts for handling and visualizing data is a routine task within technical research workflows. To explore the ability of AI tools to displace human effort in this ambit, we utilize a pre-existing exercise from a PhD research skills induction module as a representative case study (the assignment handout is also available at [25].)

This assignment “*is about developing a mature workflow for serious numerical analysis tasks, where the steps of data ingestion, processing, and visualisation are performed in a repeatable and transparent manner within robust scripts.*” It requires students to write scripts to handle meteorological data<sup>1</sup>, create publication-quality graphics, and integrate these into a  $\LaTeX$  document formatted for a specific academic journal. This multistep process mirrors the real-world research workflows that engineering PhD students encounter, from raw data manipulation to the final stages of manuscript preparation and research dissemination.

We investigate how effectively an advanced LLM can interpret the assignment requirements, generate appropriate code for data processing and visualization, and even assist in the preparation of  $\LaTeX$  documents. This approach allows us to assess the potential of AI tools to support — or potentially supplant — aspects of the traditional apprenticeship model of doctoral education in engineering.

At first Claude is prompted to simply visualise the data by producing a simple Python script. Afterwards, it is asked to add more detail to the plots without any explicit instructions on how to do this.

After several cycles of iterative prompting, Claude provided the final Python script given in Fig. 1. The images in Table I documents the outputs of this script, in both its initial and refined form.

It can be seen that Claude is capable of understanding the requirements of this graded exercise and translating this to appropriate Python code. This code runs and produces attractive graphics; these can be refined and tweaked with natural language prompts.

<sup>1</sup><https://data.gov.ie/dataset/casement-hourly-data/resource/61a17144-91c2-4a1e-a05b-b03a660e27d7>

```

import pandas as pd
import matplotlib.pyplot as plt
import seaborn as sns
from scipy import stats
import numpy as np

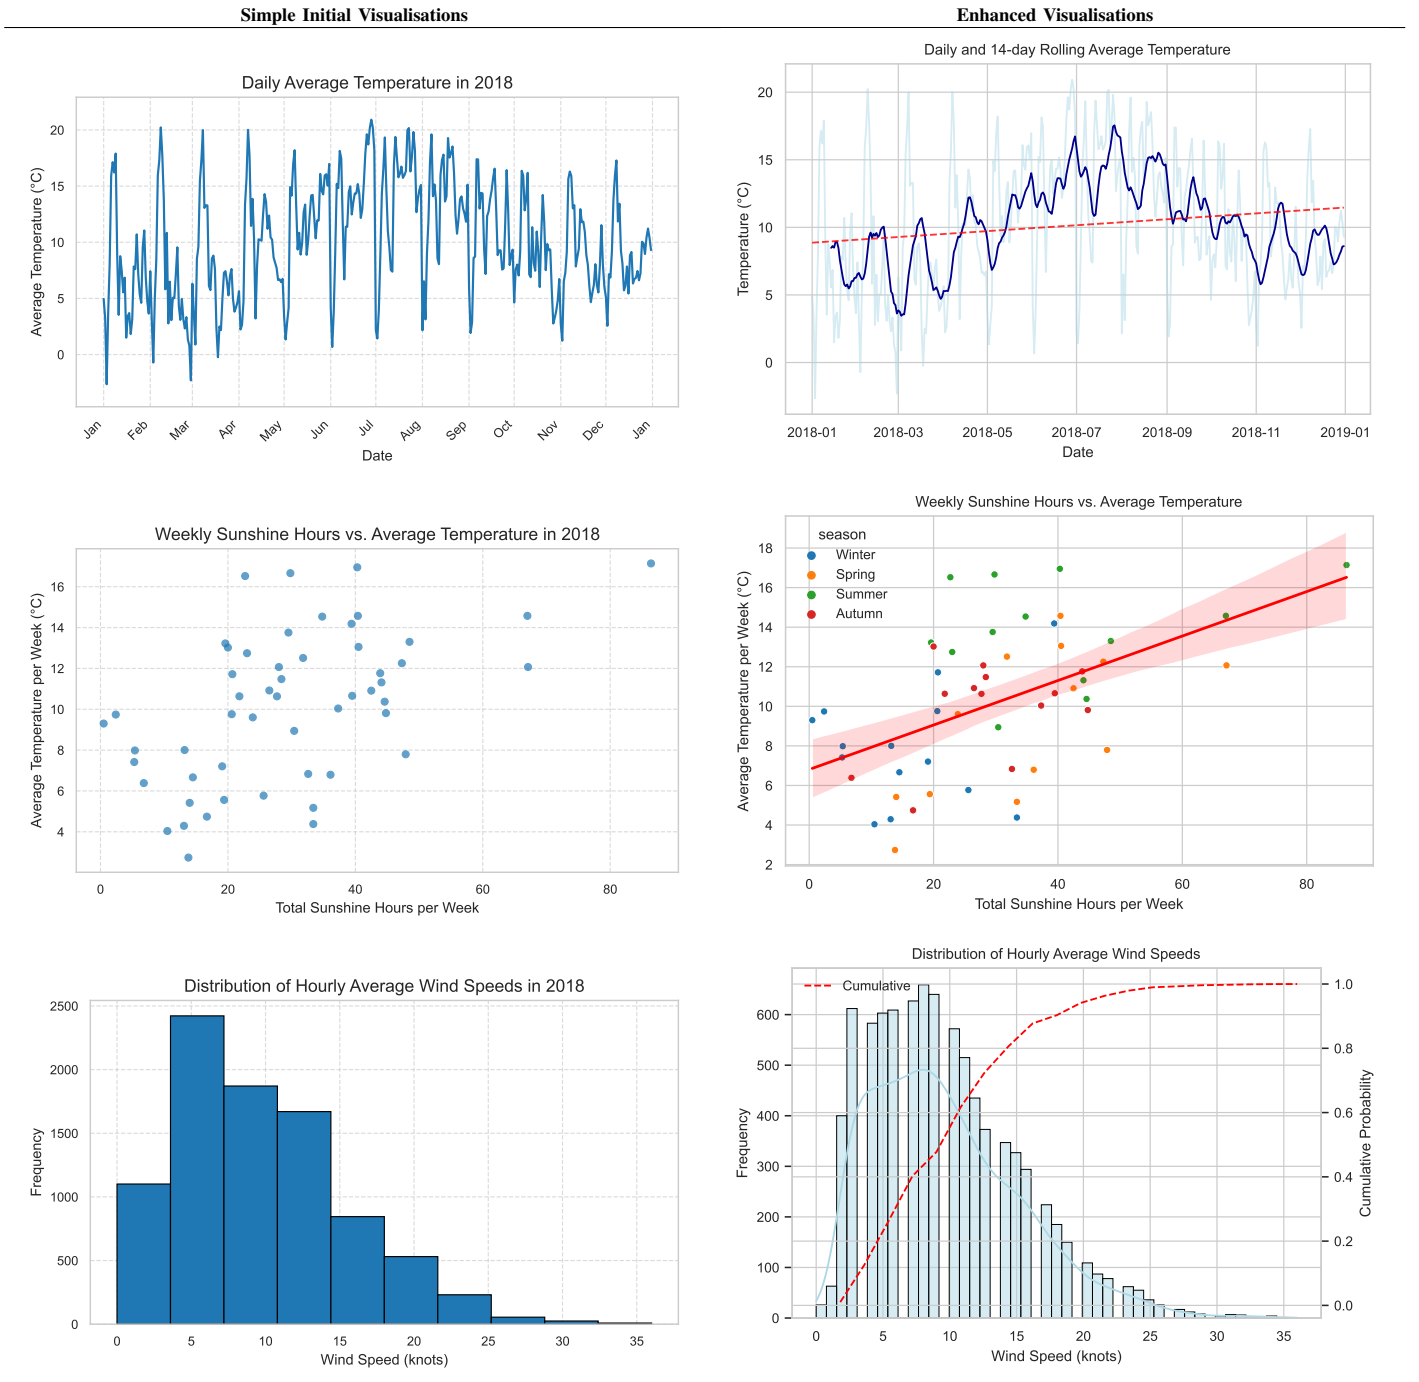
# Load the data
df = pd.read_csv('processed_data_20240718_120000.csv', parse_dates=['date'], index_col='date')
# Set up the plot style
plt.style.use('seaborn-whitegrid')
sns.set_style("whitegrid")
sns.set_context("paper")
# 1. Temperature plot
fig, ax = plt.subplots(figsize=(6, 4))
daily_temp = df['temp'].resample('D').mean()
rolling_temp = daily_temp.rolling(window=14).mean()
ax.plot(daily_temp.index, daily_temp.values, alpha=0.5, color='lightblue')
ax.plot(rolling_temp.index, rolling_temp.values, color='darkblue')
z = np.polyfit(range(len(daily_temp)), daily_temp.values, 1)
p = np.poly1d(z)
ax.plot(daily_temp.index, p(range(len(daily_temp))), "r--", alpha=0.8)
ax.set_title('Daily and 14-day Rolling Average Temperature')
ax.set_xlabel('Date')
ax.set_ylabel('Temperature (°C)')
plt.tight_layout()
plt.savefig('temperature_plot_2018.pdf')
plt.close()
# 2. Rainfall plot
fig, ax = plt.subplots(figsize=(6, 4))
daily_rain = df['rain'].resample('D').sum()
rolling_rain = daily_rain.rolling(window=14).mean()
ax.bar(daily_rain.index, daily_rain.values, alpha=0.5, color='lightblue')
ax.plot(rolling_rain.index, rolling_rain.values, color='darkblue')
ax.set_title('Daily Rainfall and 14-day Rolling Average')
ax.set_xlabel('Date')
ax.set_ylabel('Rainfall (mm)')
plt.tight_layout()
plt.savefig('rainfall_plot_2018.pdf')
plt.close()
# 3. Sunshine vs Temperature scatter plot
fig, ax = plt.subplots(figsize=(6, 4))
weekly_data = df.resample('W').agg({'sun': 'sum', 'temp': 'mean'})
weekly_data['season'] = weekly_data.index.month.map({1: 'Winter', 2: 'Winter', 3: 'Spring',
4: 'Spring', 5: 'Spring', 6: 'Summer',
7: 'Summer', 8: 'Summer', 9: 'Autumn',
10: 'Autumn', 11: 'Autumn', 12: 'Winter'})
sns.scatterplot(data=weekly_data, x='sun', y='temp', hue='season', ax=ax)
sns.regplot(data=weekly_data, x='sun', y='temp', scatter=False, ax=ax, color='red')
ax.set_title('Weekly Sunshine Hours vs. Average Temperature')
ax.set_xlabel('Total Sunshine Hours per Week')
ax.set_ylabel('Average Temperature per Week (°C)')
plt.tight_layout()
plt.savefig('sunshine_temp_plot_2018.pdf')
plt.close()
# 4. Wind speed histogram and KDE
fig, ax = plt.subplots(figsize=(6, 4))
sns.histplot(df['wdsp'], kde=True, ax=ax, color='lightblue', edgecolor='black')
ax.set_title('Distribution of Hourly Average Wind Speeds')
ax.set_xlabel('Wind Speed (knots)')
ax.set_ylabel('Frequency')
# Add cumulative distribution
counts, bin_edges = np.histogram(df['wdsp'], bins=20)
cdf = np.cumsum(counts)
ax2 = ax.twinx()
ax2.plot(bin_edges[1:], cdf/cdf[-1], color='red', ls='--', label='Cumulative')
ax2.set_ylabel('Cumulative Probability')
ax2.legend(loc='upper left')

plt.tight_layout()
plt.savefig('wind_speed_plot_2018.pdf')
plt.close()

```

Figure 1. This code listing shows the sophisticated data ingestion and visualisation Python script written by Claude after a short session of natural language prompting

Table I.  
VISUALISATION OF DATA USING CLAUDE 3.5



## B. Creating Scientific Diagrams

The task here is to provide an attractive diagrams that illustrates the well-known *binary search algorithm* [26] from a bare step-by-step textual description. Claude is prompted to provide code to produce a suitable diagram in the Graphviz language, which can be pasted into the  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$  environment to procedurally generate graphics and illustration.

Claude is successful at this task; the correct Graphviz code it provided is given in Fig. 3. The corresponding graphics is shown in Fig. III-B; an earlier iteration is given in Fig. III-B. While the relative aesthetics merits between these two graphics can be argued, this iterative refinement again shows that (a)

Claude can provide sophisticated technical code in response to demanding natural language inputs and (b) can refine these quite flexibly to meet the user’s needs.

## C. Building Tables in $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$

Quite often it is desired to produce high quality tables or to copy tables from other literature resources, and this can take a lot of unnecessary time for doctoral students.

In order to assess the performance of Claude in this domain, firstly it is provided with a raster screenshot of a human-made numerical presentation: see Table II. This table was selected

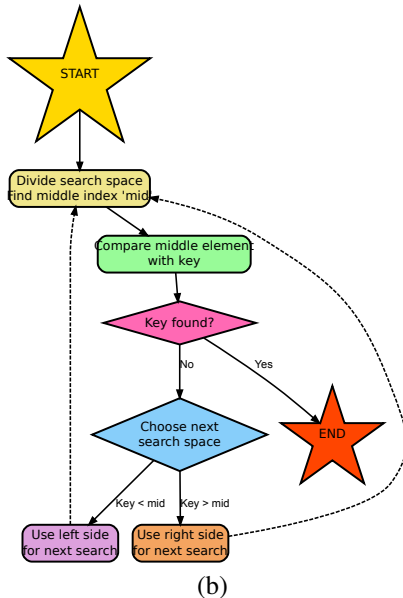
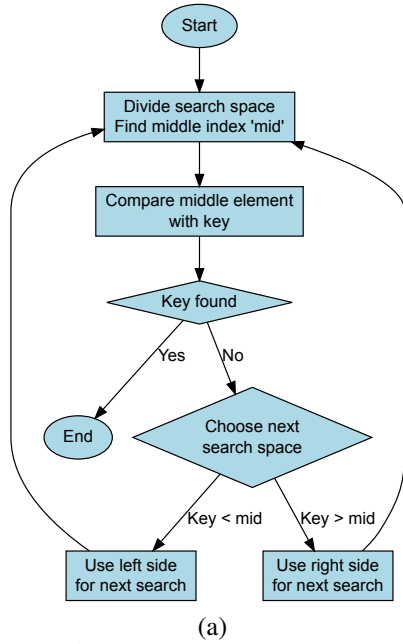


Figure 2. Illustrations of the binary search algorithm, envisioned and implemented in Graphviz by Claude with (a) the result from a single initial prompt (b) a subsequent refined diagram.

arbitrarily from earlier work published by the present senior author:

- A. Beiranvand and P. Cuffe, “Negative results on deploying distributed series reactance devices to improve power system robustness against cascading failures”, *IEEE Transactions on Power Systems*, vol. 36, no. 6, pp. 5210–5221, 2021.

Even working with just an unstructured screenshot, after a few prompts Claude is able not only to extract the textual and numerical content, but also to provide correct  $\text{\LaTeX}$  markup code that can recreate the same table format. This markup code is presented in Fig. 4; the resulting compiled table appears as Table III.

Again, this examples shows how, at a minimum, modern AI tools offer to save doctoral students time and frustration in working through routine data processing and manuscript preparation tasks.

Table II.  
THE ORIGINAL TABLE WHOSE SCREENSHOT WAS PROVIDED TO CLAUDE

Sample grid	Average load shed		Differential	
	$\partial = 0\%$	$\partial = 5\%$	$\partial = 10\%$	$\partial = 25\%$
nesta_case29_ieee	11.68 %	+2.21 pp	-0.04 pp	-2.63 pp
nesta_case30_as	2.46 %	-0.27 pp	-0.27 pp	+0.65 pp
nesta_case30_ieee	7.1 %	0 pp	0 pp	-0.56 pp
nesta_case39_ieee	11.85 %	0 pp	-0.01 pp	+0.21 pp
nesta_case57_ieee	12.25 %	0 pp	+0.04 pp	0 pp
nesta_case73_ieee_rts	7.26 %	-0.47 pp	+0.3 pp	-0.29 pp
nesta_case118_ieee	8.45 %	-0.72 pp	+0.99 pp	-0.87 pp
nesta_case189_ieee	31.46 %	0 pp	-1.9 pp	0 pp

Table III.  
THE TABLE GENERATED BY THE  $\text{\LaTeX}$  MARKUP WRITTEN BY CLAUDE

Sample grid	Average load shed		Differential	
	$\partial = 0\%$	$\partial = 5\%$	$\partial = 10\%$	$\partial = 25\%$
nesta_case29_ieee	11.68 %	+2.21 pp	-0.04 pp	-2.63 pp
nesta_case30_as	2.46 %	-0.27 pp	-0.27 pp	+0.65 pp
nesta_case30_ieee	7.1 %	0 pp	0 pp	-0.56 pp
nesta_case39_ieee	11.85 %	0 pp	-0.01 pp	+0.21 pp
nesta_case57_ieee	12.25 %	0 pp	+0.04 pp	0 pp
nesta_case73_ieee_rts	7.26 %	-0.47 pp	+0.3 pp	-0.29 pp
nesta_case118_ieee	8.45 %	-0.72 pp	+0.99 pp	-0.87 pp
nesta_case189_ieee	31.46 %	0 pp	-1.9 pp	0 pp

#### D. Literature review and knowledge summary

Due to good visual and logic capabilities and a large context window of 200K tokens [22], Claude 3.5 Sonnet should be able to summarise the data from written documents very well. This could mean that LLMs have the potential to act as a comprehensive literature review assistants. To interrogate this capability, the below recent paper, selected arbitrarily, is provided to Claude <sup>2</sup>:

- B. Walther, L. Polzin, M. van Delden, and T. Musch, “An ultra-wideband reference frequency chirp generator utilizing fractional frequency divider with high linearity”, *IEEE Open Journal of Circuits and Systems*, 2024

First, Claude is asked to summarise this paper and weigh up its strengths and weakness. This summary is presented in Fig. 5. As shown there, key performance metrics are extracted and itemised into a compact form. Moving from this general summary, a more granular question is posed to Claude: “can you summarise the operation of the system in high detail?”. The response to this probing question is provided in Fig. 6

From these listing, we can see that the summary and system operation description received from Claude were rather well written. They provide a clear review of the main points within the paper, possibly saving sizeable amounts of time for a reader who wishes to skim the literature, without introducing any noticeable logical fallacies. This speaks to the potential dangers of relying on tools like Claude, as outcomes as enticing as the ones shown in the figures may build an unhealthy reliance on such software for PhD students.

#### E. Proofreading and Prognostic Peer-review Capabilities

In order to improve the quality of draft manuscripts before submission a doctoral students could use an AI tool to perform a structured proofreading review. To frame such an exercise we adapt the manuscript preparation advice in [29] and in [30]. For expository purposes, a paper previously published by the present senior author is used as though it were a draft manuscript:

<sup>2</sup>The paper being critiqued here is licensed under CC BY-NC-ND 4.0, and with the authors’ agreement

```

digraph binary_search {
  rankdir=TB;
  bgcolor="#FFFFFF";
  node [fontname="Arial Bold", fontsize=12, penwidth=2.5, shape=box, style="filled,rounded", margin=0.2];
  edge [fontname="Arial", fontsize=10, penwidth=2, color="#000000"];

  start [label="START", shape=star, style="filled,bold", fillcolor="#FFD700", color="#000000", width=1, height=1];
  divide [label="Divide search space\nFind middle index 'mid'", fillcolor="#F0E68C", color="#000000"];
  compare [label="Compare middle\nelement with key", fillcolor="#98FB98", color="#000000"];
  found [label="Key found?", shape=diamond, style="filled,bold", fillcolor="#FF69B4", color="#000000"];
  choose [label="Choose next\nsearch space", shape=diamond, style="filled,bold", fillcolor="#87CEFA", color="#000000"];
  left [label="Use left side\nfor next search", fillcolor="#DDA0DD", color="#000000"];
  right [label="Use right side\nfor next search", fillcolor="#F4A460", color="#000000"];
  end [label="END", shape=star, style="filled,bold", fillcolor="#FF4500", color="#000000", width=1, height=1];

  // Invisible nodes for guiding the left arrow
  invisible1 [shape=point, width=0, height=0, style=invis];
  invisible2 [shape=point, width=0, height=0, style=invis];

  start -> divide [dir=forward];
  divide -> compare [dir=forward];
  compare -> found [dir=forward];
  found -> end [label="Yes", dir=forward];
  found -> choose [label="No", dir=forward];
  choose -> left [label="Key < mid", dir=forward];
  choose -> right [label="Key > mid", dir=forward];
  right -> divide [dir=forward, constraint=false, style=dashed];

  // Left wraparound arrow
  left -> invisible1 [dir=forward, constraint=false, style=dashed];
  invisible1 -> invisible2 [dir=forward, constraint=false, style=dashed];
  invisible2 -> divide [dir=forward, constraint=false, style=dashed];

  // Positioning
  {rank=same; start; invisible1}
  {rank=same; divide; invisible2}
  edge [style=invis];
  start -> invisible1 [weight=100];
  divide -> invisible2 [weight=100];
}

```

Figure 3. This code listing shows the Graphviz code produced by Claude to illustrate the binary search algorithm

```

\begin{table}
\centering
\caption{The Effects of DSR Penetration on the Average Load-Shedding for the Clustering Methodology}
\label{table:4}
\begin{tabular}{lcccc}
\hline
Sample grid & Average load shed & \multicolumn{3}{c}{Differential} \\
\cline{3-5}
& $\partial$ = 0\% & $\partial$ = 5\% & $\partial$ = 10\% & $\partial$ = 25\% \\
\hline
neste\_case29\_ieee & 11.68 \% & \cellcolor{pink!30}+2.21 pp & -0.04 pp & \cellcolor{green!30}-2.63 pp \\
neste\_case30\_as & 2.46 \% & \cellcolor{green!30}-0.27 pp & -0.27 pp & \cellcolor{pink!30}+0.65 pp \\
neste\_case30\_ieee & 7.1 \% & 0 pp & 0 pp & \cellcolor{green!30}-0.56 pp \\
neste\_case39\_ieee & 11.85 \% & 0 pp & -0.01 pp & \cellcolor{pink!30}+0.21 pp \\
neste\_case57\_ieee & 12.25 \% & 0 pp & \cellcolor{pink!30}+0.04 pp & 0 pp \\
neste\_case73\_ieee\_rts & 7.26 \% & \cellcolor{green!30}-0.47 pp & \cellcolor{pink!30}+0.3 pp & \cellcolor{green!30}-0.29 pp \\
neste\_case118\_ieee & 8.45 \% & \cellcolor{green!30}-0.72 pp & \cellcolor{pink!30}+0.99 pp & \cellcolor{green!30}-0.87 pp \\
neste\_case189\_ieee & 31.46 \% & 0 pp & \cellcolor{green!30}-1.9 pp & 0 pp \\
\hline
\end{tabular}
\end{table}
}

```

Figure 4. This code listing shows the L<sup>A</sup>T<sub>E</sub>X markup produced by Claude to reproduce a table from just a raster screenshot provided to it.

Prompt:

I need you to become a circuit analysis expert. I am including a paper from IEEE explore, can you review it. I need you to explain the main idea, describe the results, identify strong and weak sides of the proposed design. Feel free to refer to the original figures and tables in the paper.

Response from Claude 3.5 Sonnet:

Certainly, I'd be happy to review and analyze this paper for you as a circuit analysis expert.

**Main Idea:**

The paper presents a novel ultra-wideband reference frequency chirp generator for MIMO radar systems. The key innovation is using a fractional frequency divider in a feed-forward architecture to generate highly linear frequency chirps with low phase noise and high bandwidth. This approach aims to address limitations in existing chirp generation methods for distributed MIMO systems.

**Key Components:**

- Custom x64 Phase-Locked Loop (PLL)
- Reference Frequency Divider (RFD)
- FPGA with high-speed transceivers
- DDR4 memory for storing pre-calculated division ratios

**Results:**

The system achieves impressive performance:

- Output frequencies up to 2 GHz with adjustable bandwidth up to 1.75 GHz
- Low phase noise: -140 dBc/Hz at 1 MHz offset from carrier
- High linearity: RMS linearity error of 241 Hz (0.67 ppm) for a 1 ms chirp with 362 MHz bandwidth
- Integrated phase noise of -37.7 dBc (45 fs jitter) for the x64 PLL

These results are summarized in Table 1, which compares this work to other state-of-the-art systems. The proposed design shows superior performance in terms of linearity, phase noise, and bandwidth.

**Strong Points:**

- High Flexibility: The feed-forward architecture allows for reprogrammable frequency range, chirp function, and modulation type.
- Superior Linearity: Achieves the lowest RMS linearity error (0.67 ppm) among compared systems.
- Low Phase Noise: Demonstrates excellent phase noise performance, crucial for radar applications.
- Wide Bandwidth: Offers up to 1.75 GHz of adjustable bandwidth, among the highest in the comparison.
- Novel Synchronization: Uses a unique method to synchronize the FPGA and RFD without feedback, enabling high-speed operation.

**Weak Points:**

- Power Consumption: At 16.3W total (4.8W analog, 11.5W digital), it's higher than fully integrated solutions.
- Size: The discrete implementation results in a larger overall size compared to integrated designs.
- Maximum Chirp Slope: While not a direct limitation of the chirp generator, the maximum demonstrated chirp slope is lower than some compared systems due to limitations in the subsequent TX-PLL.

The paper provides detailed explanations and visualizations of the system's performance. Figure 16 shows excellent agreement between simulated and measured phase noise for the x64 PLL. Figure 20 demonstrates the clean, linear chirp generated by the system, both at the RFD output and after processing by a TX-PLL. In conclusion, this chirp generator offers state-of-the-art performance in terms of linearity, phase noise, and bandwidth, making it well-suited for high-precision MIMO radar applications. Its main trade-offs are increased power consumption and size compared to fully integrated solutions, but it offers superior flexibility and performance that may justify these drawbacks for many applications.

Figure 5. This interaction extract shows Claude's abilities at summarising and synthesising information from advanced, technical publications.

- P. Cuffe and A. Keane, "Visualizing the electrical structure of power systems", *IEEE Systems Journal*, vol. 11, no. 3, pp. 1810–1821, 2015

This paper, which suggests a new way to visualise power systems, is used to evaluate the ability of Claude to perform simulated peer review and to deeply examine the quality of technical writing. The results of this trial are presented in Figures 7 and 8.

From these listing, we see that Claude can rapidly pinpoint the (apparent) weak points of the paper [31], and provide tips on how to improve these. It should be noted that when Claude is tasked to be a "critical paper reviewer" it does sometimes over-exaggerate the scale of the problems. Overall, the comments about the text part of the paper seem mostly valid and appear to provide ideas that might improve the paper.

This capability could allow a student to simulate a review from the most critical peers and get the feedback about their work before even submitting their paper for a journal or conference. Students and researchers using this approach may attain a competitive advantage in publishing over others.

#### IV. DISCUSSIONS

The notional examples presented in this paper seem to indicate that LLMs are quite capable at automating the regular tasks that PhD students do day-to-day during their education process.

There are a number of concerns associated with this development:

It has been shown that the information returned by LLMs is not always an accurate representation of the underlying data [32]. Therefore, anyone using LLMs for these purposes cannot rely on the answers blindly and should ideally double-check everything themselves.

Additionally, those skills of data analysis, processing and presentation, that PhD students are often valued for [33], now have a different import under the influence of LLMs. On the one hand, the convenience of automating tedious and routine tasks should increase the true research productivity of doctoral students. On the other hand, this could have negative implications for fundamental programming skills of students, as they delegate more and more of these coding tasks to AI. While historically most engineering doctoral students will have had to learn advanced programming skills *en route* to their

Prompt:

Can you describe the operation of the system in high detail as well?

Response from Claude 3.5 Sonnet:

Certainly. I'll provide a detailed description of the system's operation, referencing the key components and stages as presented in the paper.

- 1) Reference Signal Generation: The system starts with a 1 GHz oven-controlled crystal oscillator (OCXO) as the primary reference source. This high-frequency, low-noise reference is crucial for achieving overall system performance.
- 2) x64 Phase-Locked Loop (PLL): The 1 GHz reference is fed into a custom x64 PLL. This PLL, implemented as a custom MMIC, multiplies the input frequency to 64 GHz. The low division ratio (64) in the feedback path helps maintain low phase noise. The output of this PLL serves as the high-frequency input to the Reference Frequency Divider (RFD).
- 3) Reference Frequency Divider (RFD): The RFD is another custom MMIC that can divide its input frequency by programmable ratios between 12 and 259. It's capable of handling input frequencies up to 94 GHz and produces output frequencies from 247 MHz to 5.3 GHz.
- 4) FPGA and DDR4 Memory: An FPGA (Xilinx Ultrascale+) with DDR4 memory is used to control the RFD. The DDR4 memory stores pre-calculated division ratios for the desired chirp pattern. These ratios are computed using a third-order MASH  $\Delta - \Sigma$  modulator to achieve fractional division ratios and high linearity.
- 5) Synchronization and Control: The FPGA uses high-speed transceivers (6.4 Gbps) to send division ratios to the RFD. A key innovation is the feed-forward design: the FPGA operates on the fixed 100 MHz reference (derived from the 1 GHz OCXO) rather than the modulated output of the RFD. This allows for higher speeds and better synchronization.
- 6) Chirp Generation Process: The DDR4 memory is preloaded with the required division ratios for the desired chirp. Upon receiving a trigger, the FPGA begins reading these ratios from memory. The FPGA's transceivers send the division ratios to the RFD at high speed (6.4 Gbps). The RFD applies these changing division ratios to its 64 GHz input, producing chirped output. This process continues for the duration of the chirp (e.g., 1 ms in the example given).
- 7) Output and Further Processing: The output of the RFD is a chirped signal in the range of 247 MHz to 5.3 GHz (up to 2 GHz demonstrated). This signal can then be used as a reference for further frequency multiplication in subsequent stages (e.g., a TX-PLL for mm-wave output).
- 8)  $\Delta - \Sigma$  Modulation and Noise Shaping: The use of a  $\Delta - \Sigma$  modulator in calculating the division ratios allows for fractional division and helps shape quantization noise to higher frequencies. This noise is then filtered out by subsequent stages (e.g., the loop filter in a TX-PLL).
- 9) Linearity and Phase Noise Performance: The high input frequency to the RFD (64 GHz) and the sophisticated control scheme allow for very fine frequency steps, contributing to the system's high linearity. The use of a high-quality OCXO and custom PLLs contributes to the excellent phase noise performance.
- 10) Flexibility and Programmability: The FPGA-based control allows for easy reprogramming of chirp parameters (bandwidth, duration, shape) by simply changing the pre-calculated division ratios stored in DDR4 memory.

This architecture combines high-speed digital control (via FPGA and high-speed transceivers) with precision analog/RF components (custom MMICs for PLL and RFD) to achieve its impressive performance. The feed-forward design and sophisticated synchronization scheme are key innovations that enable the system to generate highly linear, low-noise chirps over a wide bandwidth.

Figure 6. This interaction extract shows Claude's abilities at extracting certain granular information from a rich source like a peer-reviewed engineering publication.

PhD thesis, this may no longer be a sound assumption.

Quite often universities have special training courses [34] where students develop their fundamental knowledge about these routine data-related tasks. Should such modules now include instruction on disciplined workflows to prompt and fact-check AI tools [32]?

Lastly, the ethical implications of using LLMs for helping with teaching duties have to be noted [19]. Thus, an approach similar to the scientific research method itself may have to be taken, where each fact is scrutinised and corrected without over-reliance on LLM agents prone to hallucinating or presenting biased data [32].

## V. CONCLUSIONS

This qualitative set of case studies presents an initial view into the capabilities of a state-of-the-art commercial artificial intelligence tool when applied to the tasks PhD students and researchers face routinely: analysing, preparing and presenting technical information. It has been shown that *Claude 3.5 Sonnet* is already quite capable in performing these tasks to a reasonably good standard. As such, many of the ancillary research skills that doctoral students have hitherto found necessary to master during their research journeys can increasingly be offloaded to artificial intelligence tools. The ultimate rubric for a PhD thesis is that it makes a fundamental contribution to human knowledge; for the time

being it appears that artificial intelligences can give ample assistance in achieving this, but cannot themselves drive the process. As such, frontier models like Claude seem a greater threat to the ancillary skills developed during a PhD rather than the core meaning of the credential itself. The day may soon come when artificial intelligence tools can conjure a credible-seeming doctoral thesis *de novo*: what then for PhD education in engineering?



Prompt:

You must become a critical paper reviewer. Use 2 papers: "Ten simple rules for structuring papers" and "Codifying Systematic Manuscript Preparation Checklists as a Training and Productivity Resource for Research Students" attached to evaluate the quality of presentation of information in the "Visualizing the Electrical Structure of Power Systems" paper. Additionally perform a peer review of the "Visualizing the Electrical Structure of Power Systems" paper according to IEEE Systems Journal paper standards.

Response from Claude 3.5 Sonnet:

I will provide a critical review of the "Visualizing the Electrical Structure of Power Systems" paper, evaluating its presentation quality based on the guidelines in the two reference papers, and conduct a peer review according to IEEE Systems Journal standards. Evaluation of Presentation Quality: **Structure and Organization:** The paper generally follows a logical structure as recommended in "Ten simple rules for structuring papers". It has a clear introduction, methodology, results, and conclusion sections. However, the paper could benefit from a more explicit statement of the central contribution in the title and abstract, as emphasized in Rule 1.

**Abstract:** The abstract provides a good overview but could be improved by following the context-content-conclusion (C-C-C) structure more closely, as suggested in Rule 5. It lacks a clear statement of the specific gap this paper fills and the broader significance of the work.

**Introduction:** The introduction does a good job of setting up the context and motivation for the work. However, it could be strengthened by more clearly articulating the specific gap in current knowledge that this paper addresses, as recommended in Rule 6.

**Results Presentation:** The paper makes extensive use of figures to present results, which is commendable. However, some figures (e.g., Fig. 7) are difficult to read due to small text and overlapping elements. This goes against the recommendations in the "Data visualization checklist" mentioned in the "Codifying Systematic Manuscript Preparation Checklists" paper.

**Conclusion:** The conclusion effectively summarizes the main findings and their implications. However, it could be improved by more explicitly stating how the paper's contribution moves the field forward, as suggested in Rule 8.

**Overall Narrative:** The paper maintains a reasonably clear narrative thread throughout, but could benefit from more explicit signposting to guide the reader through the logical flow of the argument.

**Peer Review according to IEEE Systems Journal Standards:**

- 1) Originality and Significance: 8/10 The paper presents a novel approach to visualizing power systems, offering new insights into their electrical structure. The use of multidimensional scaling to project electrical distances into 2D space is innovative and potentially impactful for the field.
- 2) Technical Content and Methodology: 7/10 The methodology is sound and well-explained overall. However, some technical terms could be better defined, and the explanation of equations (1)-(11) could be more thorough.
- 3) Presentation and Clarity: 6/10 While generally well-written, the paper has room for improvement in clarity:
  - The abstract could more clearly state the specific contribution and its broader significance.
  - Some sections, particularly in the methodology, could benefit from more detailed explanations.
  - The narrative flow could be enhanced with more explicit signposting.
- 4) Relevance to IEEE Systems Journal: 9/10 The paper addresses important issues in power system analysis and visualization, aligning well with the journal's scope.
- 5) Literature Review: 7/10 The literature review is comprehensive, but could more clearly articulate how this work builds on or differs from previous approaches.
- 6) Figures and Tables: 6/10 While the figures are informative and numerous, some (particularly Figs. 7-11) are cluttered and difficult to read. Improving the clarity of these visualizations would significantly enhance the paper.
- 7) Structure and Organization: 8/10 The paper generally follows a logical structure with clear introduction, methodology, results, and conclusion sections. However, it could benefit from a more explicit statement of the central contribution in the title and abstract.
- 8) Conclusion and Future Work: 7/10 The conclusion effectively summarizes the main findings, but could be improved by more explicitly stating how the paper's contribution moves the field forward. A brief section on limitations and future work would be beneficial.
- 9) Adherence to Journal Format: 9/10 The paper appears to adhere well to the journal's format requirements in terms of length, section organization, and referencing style.
- 10) Overall Impact: 8/10 The paper presents valuable and novel research that has the potential to significantly impact how power systems are visualized and understood.

**Overall Score:** 75/100

Figure 7. This extract shows how Claude can act as a structured proofreader and simulated peer reviewer: a doctoral student might use such capabilities to iteratively refine a manuscript before submission. Impressively, here Claude is reading and applying the contents of two guidance paper to the manuscript under consideration.

**Recommendations:** Revise the abstract to more clearly state the specific contribution and its broader significance. Improve the clarity of technical explanations, particularly around equations and methodologies. Enhance the readability of figures, particularly 7-11, by reducing clutter and increasing font sizes. Strengthen the conclusion by more explicitly stating how this work advances the field. Consider adding a brief section on limitations and future work. Improve the narrative flow with more explicit signposting throughout the paper.

In conclusion, this paper presents valuable and novel research, but could benefit from revisions to improve clarity and impact. With these improvements, it would make a strong contribution to the IEEE Systems Journal.

Figure 8. This extract shows the didactic potential of Claude, here acting as an academic writing coach.

## REFERENCES

- [1] A. Gokul, "LLMs and AI: Understanding its reach and impact", *Preprints*, May 2023. DOI: 10.20944/preprints202305.0195.v1. [Online]. Available: <https://doi.org/10.20944/preprints202305.0195.v1>.
- [2] J. Qadir, "Engineering education in the era of ChatGPT: Promise and pitfalls of generative AI for education", in *2023 IEEE Global Engineering Education Conference (EDUCON)*, IEEE, 2023, pp. 1–9.
- [3] A. Hickey, C. Ó'Faoláin, J. Healy, K. Nolan, E. Doheny, and P. Cuffe, "A threat assessment framework for screening the integrity of university assessments in the era of large language models", in *IEEE Research and Technologies for Society and Industry 2024*, 2024.
- [4] H. Crompton and D. Burke, "Artificial intelligence in higher education: The state of the field", *International Journal of Educational Technology in Higher Education*, vol. 20, no. 1, p. 22, 2023.
- [5] S. Paek and N. Kim, "Analysis of worldwide research trends on the impact of artificial intelligence in education", *Sustainability*, vol. 13, no. 14, p. 7941, 2021.
- [6] R. Gilbert, J. Balatti, P. Turner, and H. Whitehouse, "The generic skills debate in research higher degrees", *Higher education research & development*, vol. 23, no. 3, pp. 375–388, 2004.
- [7] S. Mowbray and C. Halse, "The purpose of the PhD: Theorising the skills acquired by students", *Higher Education Research & Development*, vol. 29, no. 6, pp. 653–664, 2010.
- [8] C. Lu, C. Lu, R. T. Lange, J. Foerster, J. Clune, and D. Ha, *The AI scientist: Towards fully automated open-ended scientific discovery*, 2024. arXiv: 2408.06292 [cs.AI]. [Online]. Available: <https://arxiv.org/abs/2408.06292>.
- [9] S. Noy and W. Zhang, "Experimental evidence on the productivity effects of generative artificial intelligence", *Science*, vol. 381, no. 6654, pp. 187–192, 2023.
- [10] A. M. Dakhel, V. Majdinasab, A. Nikanjam, F. Khomh, M. C. Desmarais, and Z. M. J. Jiang, "Github Copilot AI pair programmer: Asset or liability?", *Journal of Systems and Software*, vol. 203, p. 111734, 2023.
- [11] M. Májovský, M. Černý, M. Kasal, M. Komarc, and D. Netuka, "Artificial intelligence can generate fraudulent but authentic-looking scientific medical articles: Pandora's box has been opened", *Journal of Medical Internet Research*, vol. 25, e46924, May 2023, ISSN: 1438-8871. DOI: 10.2196/46924.
- [12] J. Jung, "Imagining doctoral education in the fourth industrial revolution: Driving technology or being driven by technology", *Minerva*, vol. 60, no. 4, pp. 615–632, 2022.
- [13] H. Wang, T. Fu, Y. Du, W. Gao, K. Huang, Z. Liu, P. Chandak, S. Liu, P. Van Katwyk, A. Deac, *et al.*, "Scientific discovery in the age of artificial intelligence", *Nature*, vol. 620, no. 7972, pp. 47–60, 2023.
- [14] S. Yu, C. Fang, Y. Ling, C. Wu, and Z. Chen, "LLM for test script generation and migration: Challenges, capabilities, and opportunities", in *2023 IEEE 23rd International Conference on Software Quality, Reliability, and Security (QRS)*, IEEE, 2023, pp. 206–217.
- [15] Z. Yang, Z. Zhou, S. Wang, X. Cong, X. Han, Y. Yan, Z. Liu, Z. Tan, P. Liu, D. Yu, *et al.*, "Matplotagent: Method and evaluation for LLM-based agentic scientific data visualization", *arXiv preprint arXiv:2402.11453*, 2024.
- [16] B. J. Lear, "Using ChatGPT-4 to teach the design of data visualizations", *Journal of Chemical Education*, 2024.
- [17] Q. Xing, C. Zheng, N. Zhu, and D. Yip, "AI-generated content for academic visualization and communication in maker education", in *2023 3rd International Conference on Educational Technology (ICET)*, IEEE, 2023, pp. 52–56.
- [18] S. Schreiber, S. Agne, I. Wolf, A. Dengel, and S. Ahmed, "Deepdesrt: Deep learning for detection and structure recognition of tables in document images", in *2017 14th IAPR international conference on document analysis and recognition (ICDAR)*, IEEE, vol. 1, 2017, pp. 1162–1167.
- [19] Y. Hicke, A. Agarwal, Q. Ma, and P. Denny, "ChaTA: Towards an intelligent question-answer teaching assistant using open-source LLMs", *arXiv preprint arXiv:2311.02775*, 2023.
- [20] H. Maheshwari, S. Bandyopadhyay, A. Garimella, and A. Nataraajan, "Presentations are not always linear! GNN meets LLM for document-to-presentation transformation with attribution", *arXiv preprint arXiv:2405.13095*, 2024.
- [21] A. Checco, L. Bracciale, P. Loreti, S. Pinfield, and G. Bianchi, "AI-assisted peer review", *Humanities and Social Sciences Communications*, vol. 8, no. 1, pp. 1–11, 2021.
- [22] Z. Huang, Z. Wang, S. Xia, and P. Liu, "OlympicArena medal ranks: Who is the most intelligent AI so far?", *arXiv preprint arXiv:2406.16772*, 2024.
- [23] A. Mandow, J. L. Martínez, and A. García-Cerezo, "Project-based learning of scientific writing and communication skills for postgraduate students", in *2014 IEEE Frontiers in Education Conference (FIE) Proceedings*, 2014, pp. 1–4. DOI: 10.1109/FIE.2014.7044022.
- [24] N. Guizani and H. E. Rodriguez-Simmonds, "Developing personal and community graduate student growth through the implementation of a LaTeX workshop", in *2016 ASEE Annual Conference & Exposition*, 2016.
- [25] Y. Mormul, J. Przybyszewski, A. Nakoud, and P. Cuffe, "Prompts and interactions with the Claude AI tool that underpin the manuscript 'Reliance on Artificial Intelligence Tools May Displace Research Skills Acquisition Within Engineering Doctoral Programmes: Examples and Implications'", Aug. 2024. DOI: 10.6084/m9.figshare.26880901.v1.
- [26] L. F. Williams Jr, "A modification to the half-interval search (binary search) method", in *ACM Proceedings of the 14th Annual Southeast Regional Conference*, 1976, pp. 95–101.
- [27] A. Beiranvand and P. Cuffe, "Negative results on deploying distributed series reactance devices to improve power system robustness against cascading failures", *IEEE Transactions on Power Systems*, vol. 36, no. 6, pp. 5210–5221, 2021.
- [28] B. Walther, L. Polzin, M. van Delden, and T. Musch, "An ultra-wideband reference frequency chirp generator utilizing fractional frequency divider with high linearity", *IEEE Open Journal of Circuits and Systems*, 2024.
- [29] B. Mensh and K. Kording, "Ten simple rules for structuring papers", *PLOS Computational Biology*, vol. 13, no. 9, pp. 1–9, Sep. 2017. DOI: 10.1371/journal.pcbi.1005619. [Online]. Available: <https://doi.org/10.1371/journal.pcbi.1005619>.
- [30] P. Cuffe, "Codifying systematic manuscript preparation checklists as a training and productivity resource for research students", in *2022 IEEE Global Engineering Education Conference (EDUCON)*, IEEE, 2022, pp. 458–465.
- [31] P. Cuffe and A. Keane, "Visualizing the electrical structure of power systems", *IEEE Systems Journal*, vol. 11, no. 3, pp. 1810–1821, 2015.
- [32] D. Baidoo-Anu and L. O. Ansah, "Education in the era of generative artificial intelligence (AI): Understanding the potential benefits of ChatGPT in promoting teaching and learning", *Journal of AI*, vol. 7, no. 1, pp. 52–62, 2023.
- [33] S. Mowbray and C. Halse, "The purpose of the PhD: Theorising the skills acquired by students", *Higher Education Research & Development*, vol. 29, no. 6, pp. 653–664, 2010.
- [34] L. Panchenko, H. Korzhov, T. Kolomiets, and M. Yenin, "PhD student training: Principles and implementation", in *Journal of Physics: Conference Series*, IOP Publishing, vol. 1840, 2021, p. 012056.