CHIWORK '25 Adjunct: Adjunct Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work
June 23 - 25, 2025
Amsterdam, Netherlands

WORK IN PROGRESS

Investigating the Impact of AI-Assisted Tools on Software Practitioner Well-Being

FAIRUZ NAWER MEEM, George Mason University, Fairfax, VA, United States

BRITTANY JOHNSON, George Mason University, Fairfax, VA, United States

CHIWORK '25 Adjunct: Adjunct Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work (June 2025)
https://doi.org/10.1145/3707640.3731915
ISBN: 9798400713972

Investigating the Impact of AI-Assisted Tools on Software Practitioner Well-Being

Fairuz Nawer Meem
George Mason University
Fairfax, USA
fmeem@gmu.edu

Brittany Johnson
George Mason University
Fairfax, USA
johnsonb@gmu.edu

Abstract

The increasing adoption of AI-assisted tools, like ChatGPT, in software development presents both opportunities and challenges. While these tools enhance productivity and streamline tasks, they also introduce new job demands, such as cognitive overload, that can impact practitioners' well-being. Well-being in this context includes mental, emotional, and physical health factors, such as job satisfaction, stress, burnout, and engagement. This paper presents preliminary findings from a study grounded in the Job Demands-Resources (JD-R) model aimed at exploring how AI-assisted tools impact practitioners' well-being. We designed a comprehensive survey to investigate key factors such as job demands, organizational and social resources, and their interplay. Our findings thus far suggest that while AI-assisted tools can increase productivity and improve focus, it can also cause mental fatigue, cognitive strain, and blurred work-life boundaries. Insights from our ongoing efforts provide a critical foundation for responsible integration of AI-assisted tools in software development.

CCS Concepts

• **Human-centered computing → Empirical studies in HCI**; • **Social and professional topics** → Sustainability; • **General and reference → Empirical studies**.

Keywords

AI-Assisted Tools, Well-Being, Mental Health, Software Development

ACM Reference Format:
Fairuz Nawer Meem and Brittany Johnson. 2025. Investigating the Impact of AI-Assisted Tools on Software Practitioner Well-Being. In *CHIWORK '25 Adjunct: Adjunct Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work (CHIWORK '25 Adjunct), June 23–25, 2025, Amsterdam, Netherlands.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3707640.3731915

CHIWORK '25 Adjunct, Amsterdam, Netherlands

ACM ISBN 979-8-4007-1397-2/25/06
https://doi.org/10.1145/3707640.3731915

1 Introduction and Motivation

The integration of artificial intelligence (AI), particularly large language models (LLMs), into natural language tool support, like ChatGPT and Gemini, has transformed workflows across diverse professional domains [16]. These tools offer capabilities such as automating routine tasks, providing real-time suggestions, and supporting problem-solving through conversational interactions [2, 9]. While originally designed as general purpose models, LLMs have found widespread application in software development, with software practitioners leveraging them for coding, debugging, documentation, and planning [5, 8, 18]. Their flexible and adaptable nature has led to increased productivity and task efficiency, positioning them as indispensable assets in modern workplaces [4, 19].

However, the adoption of these tools is not without challenges. Practitioners often report experiencing cognitive overload due to the volume of generated outputs, as well as emotional stress from managing inaccuracies or unexpected tool behavior [6, 10]. These challenges, compounded by growing reliance on AI tools, can lead to new kinds of job demands, such as the pressure to meet heightened productivity expectations or validate tool outputs [14]. At the same time, these tools offer benefits, such as reducing mundane tasks, enhancing collaboration, and providing learning opportunities, which can positively impact well-being when balanced appropriately [11].

Given the potential for both positive and negative impact beyond productivity, we designed a study focused on understanding the dual-edged nature of AI-assisted tools, exploring their potential to both alleviate job demands and pose risks to practitioners' mental, emotional, and physical health. Grounded in the Job Demands-Resources (JD-R) model [12], our research examines the interplay between job demands, resources, and well-being outcomes in the context of LLM-driven tool usage in the software development workplace. In this paper, we outline our study design, initial findings and how our future efforts, and the broader research community, can contribute to the growing body of knowledge on the impact of LLMs and inform the responsible and sustainable integration of these tools into professional workflows.

2 The JD-R Theoretical Framework

The Job Demands-Resources (JD-R) model, shown in Figure 1, provides a robust framework for examining and evaluating the interplay between *job demands*, *resources*, and their impact on *well-being outcomes* [12]. *Job Demands* refer to the cognitive, emotional, and physical efforts required by practitioners in their work environment, often leading to strain or stress when demands exceed available resources [12]. In our study, we explore job demands along several dimensions, which we outline in Section 3.1. *Job Resources* encompass the physical, psychological, and organizational supports that enable practitioners to manage job demands effectively [12]. The JD-R model also considers how these demands and resources impact *well-being outcomes*. This include both positive outcomes, such as increased productivity and motivation, and negative outcomes, such as stress and burnout.

According to this model, job demands – such as workload, time pressure, or cognitive strain – can contribute to burnout and stress if not balanced by sufficient resources. Conversely, job resources – such as supportive leadership, effective tools, and clear communication – can mitigate the negative effects of demands while enhancing positive outcomes. In the context of software development, AI-assisted tools introduce both resources and demands. On the one hand, these tools can be used as *resources* to alleviate job demands by automating routine tasks, streamlining workflows, and offering real-time support for debugging and problem-solving [17]. On the other hand, they may introduce new *job demands* that may yet have adequate resources to support, such as cognitive overload from excessive suggestions, validation fatigue due to inaccuracies, and pressure to meet heightened productivity expectations [13].

3 Methodology

Our study leverages the Job Demands-Resources (JD-R) model to investigate the impact of AI-assisted tools on software practitioners' well-being. The audience of interest for our research is software practitioners actively engaged in using AI-assisted tools like ChatGPT and GitHub Copilot for their development tasks. We designed our study to understand *what the potential short and long term effects of AI-assisted tool use on software practitioner well-being* (**RQ1**) and *what considerations (individual or organizational) can increase the potential for responsible and healthy use of AI-assisted tools* (**RQ2**).

3.1 Study Design

For our research, we adopted a mixed methodology to explore the interaction between job demands, resources, and well-being outcomes in software development settings where AI-assisted tools are widely used. In this paper, we report preliminary insights from our survey, which will inform future interviews. Our survey, which is publicly available [1], is designed to explore how AI-assisted tools impact software practitioners' well-being, with a strong foundation in the Job Demands-Resources (JD-R) model. This research is approved by our university's Institutional Review Board (IRB) [2].

Each section of the survey focuses on a specific dimension of the framework, with carefully crafted questions to capture meaningful insights into practitioners' experiences. To cover *job demands*, we asked questions related to cognitive strain (for assessing the mental effort required to use AI-assisted tools effectively, e.g., Q6, Q8), emotional stress (for exploring the frustration or stress caused by AI tools, e.g., Q7, Q9), time pressure (for evaluating whether tool usage accelerates task deadlines or creates urgency, e.g., Q10) etc. To explore *job resources*, we added questions related to social resources (to explore team collaboration and support for AI tool adoption, e.g., Q19, Q20), organizational resources (to capture support provided by the organization, e.g., Q31, Q32), skill alignment (to assess how well practitioners' skills align with tool usage, e.g., Q27, Q29) etc. Lastly, to investigate *well-being outcomes*, we added questions related to burnout and stress (to measure the extent to which practitioners feel overworked or stressed, e.g., Q49, Q51). Additionally, to assess positive outcomes, there were questions regarding engagement and motivation (e.g., Q43, Q50), job satisfaction (e.g., Q54), well-being (e.g., Q9, Q53, Q55) etc. We also added questions outside the JD-R framework to address the unique aspects of AI tool usage and software development.

3.2 Data Collection

Our audience of interest includes individuals from diverse professional backgrounds and industries that use AI-assisted tools in software development work. We are intentional about seeking and engaging practitioners with diverse needs to ensure we are considering those who may be most susceptible to negative impact. To engage this audience, we disseminated the survey disseminated through various online channels, including professional networks such as LinkedIn. We also plan to share it in developer communities like GitHub and Stack Overflow. These platforms provide access to practitioners with varying levels of experience and roles. We also leverage our professional and community networks. To encourage participation, those who complete the survey have the opportunity to enter a gift card raffle. This approach, despite the risks [7], not only helps maximize response rates but also ensures that practitioners feel valued for their time and contributions [1].

4 Preliminary Findings

Thus far, we have received 41 responses from software practitioners. Our preliminary analysis reveals a nuanced landscape of engagement with AI-assisted tools. Among our 41 respondents, the most widely used AI tool was ChatGPT (96%), followed by Deepseek (37%), Gemini (37%), and GitHub Copilot (33%). Others like Claude AI (19%), Cursor AI (7%), and various other tools (4%) saw limited use. These tools supported a range of software development tasks—most commonly writing code (78%), debugging (67%), and documenting (48%). Other supported tasks included learning about codebases (44%), project planning (30%), testing (26%), deployment (15%), and even collaboration (4%).

4.1 Potential Short & Long Term Effects on Well-Being (RQ1)

When reporting on **cognitive load and mental fatigue effects**, many respondents (65%) reported occasional feelings of being overwhelmed by the volume of information generated, with some (13%) experiencing frequent cognitive overload. Almost half of the participants (48%) reported experiencing mental fatigue. Most users (57%) found working with AI tools to be mentally demanding, with a subset (21%) describing the experience as highly or extremely demanding. We noticed similar patterns in open-ended responses, where one participant mentioned dealing with "*overwhelming content to learn in a short amount of time*". Respondents also attributed their mental fatigue to the effort involved in "*trying to find the right question to ask*." Despite the risk of mental fatigue indicated in our responses, several respondents noted that AI-assisted tools can sometimes relieve pressure and even enhance motivation. As stated by one respondent, "*finding solutions for some silly issues on*

[1] https://go.gmu.edu/SoftwarePractitionerWellBeing
[2] RAMP ID number: STUDY00000342

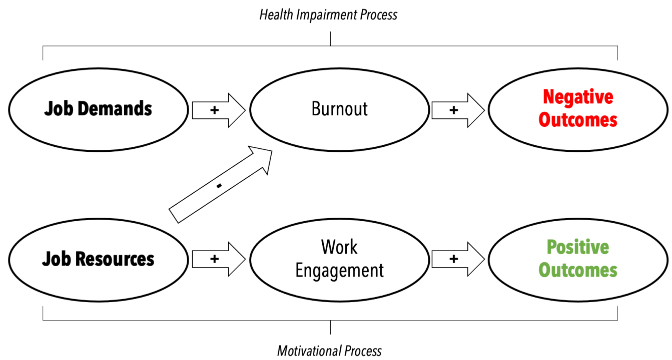

Figure 1: The Job Demands-Resources (JD-R) Framework

Google was hard, but AI made it easier now. So mental stress is no issue here. Rather, it is motivating me."

As for **physical and ergonomic effects**, physically, most practitioners (65%) did not experience discomfort, but some (30%) reported issues like eye strain and headaches, and a small fraction (4%) reported more noticeable physical discomfort requiring breaks, indicating the importance of ergonomic awareness in AI tool usage. When reporting on **impact on work-life balance**, a majority of respondents (52%) admitted to occasionally working beyond regular hours due to AI tool usage, with a smaller yet noteworthy group (17%) frequently extending their workdays. In fact, a majority of respondents (65%) noted finding the pace of development with AI-assisted tooling overwhelming with half of respondents feeling differences in the pressure to produce (55%).

When it comes to **emotional well-being and stress**, responses reflected a mix of optimism and concern. While 38% of respondents reported always feeling energized, the majority of respondents reported either occasionally feeling worn out (44%) or frequent to constant burnout (19%). Almost half of the respondents reported experiencing occasional stress or anxiety when using AI assistants, with a small group (6%) reporting frequent stress. One participant expressed this feeling as "*too much expectation within shortest amount of time*". Most respondents (53%) reported no change in their mental health from using AI tools, while 33% noted some improvement. However, 14% reported experiencing a decline in their mental health. The majority of participants reported an **increase in confidence** due to AI tool usage, with a small subset reporting a decline in confidence (7%). This could be connected to the fact that 80% of respondents reported that AI-assisted tools help them manage job-related stress, including one participant stating, "*they help me complete my tasks faster which lowers my stress.*". However, for a handful of respondents, AI-assisted tools do not help in this regard with some reporting negative or mixed experiences.

4.2 Individual and Organizational Considerations for Responsible Integration (RQ2)

Organizational support for AI usage varied considerably. While a majority of respondents (70%) indicated that their workplaces were supportive, a nontrivial portion reported either lacking support (19%) or being uncertain (11%) about organizational readiness. We also explored **social dynamics** surrounding AI tool use. While 53% of participants reported harmonious interactions with colleagues regarding AI usage, 29% reported experiencing occasional tensions, and 18% regularly dealt with interpersonal conflicts, highlighting the social complexity that can accompany collaborative AI adoption. Most respondents (58%) described their co-workers as generally or very supportive , with an additional 35% reporting occasional support. A small proportion (6%) felt their coworkers were rarely helpful. A reassuring 71% of respondents reported never experiencing harassment related to AI tool usage, though 24% noted rare instances, and 6% experienced it regularly, highlighting the need for organizations to remain vigilant about emerging social dynamics tied to AI adoption. While 53% of respondents reported having **clearly defined roles** with no conflicts, nearly half reported experiencing some level of role tension, pointing to a need for better role alignment when integrating AI tools into workflows.

We also ask respondents questions regarding **workflow disruptions and bureaucratic barriers** that may come with using AI-assisted tools. More than half of our respondents thus far (55%) indicated they always have enough work, but 30% experienced occasional slow periods, and 15% frequently ran out of tasks when using AI-assisted tools. While some respondents felt that changes to AI tools or their usage policies usually led to improvements, the majority reported occasional disruptions (59%), with some experiencing consistent negative impacts (18%). Similarly 35% of respondents felt that administrative procedures led to occasional complications and nearly a quarter (24%) reported frequent slowdowns or constant red tape, indicating that bureaucracy can undermine the efficiency gains offered by AI tools.

With respect to **leadership and team-level integration**, 35% of respondents received strong support, while the majority of the remaining respondents reported infrequent (35%) to occasional (24%) support. While 58% of respondents viewed their teams as generally or very effective at integrating AI tools, 41% indicated partial success, achieving some goals but falling short on others, suggesting room for improvement in team-level adoption and coordination. Experiences surrounding **skill utilization and policy** were generally positive, with the majority of respondents (77%) indicating that AI assistants met or exceeded their expectations. The majority of respondents also reported using many or all of their skills effectively when working with AI assistants. We also found generally positive sentiments regarding **career outlook and adaptability**, where most respondents (82%) believed that AI-assisted tools had a positive or somewhat positive impact on their career prospects. When it comes to rapid changes in AI technologies, most respondents (94%) indicated they are flexible to the changes.

When asked how we can improve integration in practice, several respondents mentioned the role of the organizations they work for. For example, respondents mentioned the importance of "*policy limiting the use of AI*" as well as a need for "*training and support on how to effectively integrate and use AI-assisted tools within their workflows*". Some respondents also noted the importance of "*the ability to differentiate the real and AI generated work*".

5 Research Limitations

We anticipate some limitation in the completion of this research. First is the curation of representative experiences and outcomes. We will be intentional in our future recruitment efforts to find and engage a diverse group of practitioners to increase the likelihood that our insights will scale. We also have contingency plans for recruitment, such as snowball sampling and recruiting for interviews outside the survey sample. The reliance on survey responses may introduce biases such as social desirability or recall bias. Triangulating survey data with other methods, such as interviews or workplace observations, will help strengthen the findings. Hence our plans to conduct a follow-up interviews with survey respondents to gain better and detailed insights about their responses.

Given the dynamic and fast evolving nature of AI-assisted tools, features and use cases that are relevant at the time of the study may become obsolete or transformed in the near future. While this is difficult to mitigate, we intend to continue this line of research which would involve regular "temperature checks" on the impact of AI-assisted tool use in practice. Lastly, while the JD-R model provides a comprehensive framework, it may not capture all aspects of the practitioner experience. Our future efforts may bring to light other complementary theories or models could provide additional perspectives.

6 Discussion and Future Work

Our findings emphasize the dual-edged nature of AI-assisted tools in software development work: while they offer significant productivity gains and can even boost practitioner confidence, they also introduce new demands and stressors, especially when not supported by thoughtful individual or organizational strategies. For example, we found that perceived advantages may also be accompanied by disruptive changes in workflows and expectations, mental fatigue, emotional exhaustion, and the blurring of work-life boundaries. This duality reflects the fundamental tension outlined in the Job Demands-Resources (JD-R) model - while tools like ChatGPT and Copilot can reduce routine burden, they may also raise expectations and introduce complexity, especially when organizational support is lacking.

Our insights thus far calls for a shift in how we think about both tool design and organizational implementation. AI tools should be designed not just for productivity, but also with user well-being in mind. This includes reducing cognitive overload through clearer outputs, minimizing repetitive validations, and potentially offering features that encourage breaks or support task management. Tools with more adaptive or emotionally aware interfaces could reduce frustration and improve the user experience [15]. On the organizational side, policies should support responsible use. This means providing training, encouraging healthy use boundaries, and offering mental health support or clear escalation paths when tool use creates stress or friction on teams. Many participants emphasized the importance of transparent, fair AI policies, something respondents in our study already find helpful, but which must be sustained as these tools evolve [17].

In this work, we applied the JD-R model to AI-assisted tool use in software development. Our findings thus far support its relevance to software development and highlights the need to treat AI integration as both a technical and human challenge. While our study focuses on software practitioners, we believe these findings are transferable to other knowledge-intensive fields such as education, healthcare, and creative work [3]. In future work, we plan to conduct follow-up interviews to explore themes such as emotional strain, adaptability, and tool misalignment in more depth. These interviews will help us capture lived experiences and refine design or policy recommendations. Additionally, we aim to engage more practitioners from diverse backgrounds to better understand the needs of underrepresented or more vulnerable populations, including software practitioners with clinically diagnosed mental health conditions. By centering well-being in AI integration and use, we hope to shift the narrative from efficiency alone to one that includes sustainability, inclusivity, and long-term human flourishing in software development workplaces.

References

[1] Basel Abdelazeem, Kirellos Said Abbas, Mostafa Atef Amin, Nahla Ahmed El-Shahat, Bilal Malik, Atefeh Kalantary, and Mostafa Eltobgy. 2022. The effectiveness of incentives for research participation: A systematic review and meta-analysis of randomized controlled trials. *PLoS One* 17, 4 (2022), e0267534.

[2] Erik Brynjolfsson, Danielle Li, and Lindsey R Raymond. 2023. *Generative AI at work*. Technical Report. National Bureau of Economic Research.

[3] Cecilia Ka Yuk Chan and Wenjie Hu. 2023. Students' voices on generative AI: Perceptions, benefits, and challenges in higher education. *International Journal of Educational Technology in Higher Education* 20, 1 (2023), 43.

[4] Mariana Coutinho, Lorena Marques, Anderson Santos, Marcio Dahia, Cesar França, and Ronnie de Souza Santos. 2024. The role of generative ai in software development productivity: A pilot case study. In *Proceedings of the 1st ACM International Conference on AI-Powered Software*. 131–138.

[5] Angela Fan, Beliz Gokkaya, Mark Harman, Mitya Lyubarskiy, Shubho Sengupta, Shin Yoo, and Jie M Zhang. 2023. Large language models for software engineering: Survey and open problems. In *2023 IEEE/ACM International Conference on Software Engineering: Future of Software Engineering (ICSE-FoSE)*. IEEE, 31–53.

[6] Michael Gerlich. 2025. AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking. *Societies* 15, 1 (2025), 6.

[7] Marybec Griffin, Richard J Martino, Caleb LoSchiavo, Camilla Comer-Carruthers, Kristen D Krause, Christopher B Stults, and Perry N Halkitis. 2021. Ensuring survey research data integrity in the era of internet bots. *Quality & quantity* (2021), 1–12.

[8] Fairuz Nawer Meem, Justin Smith, and Brittany Johnson. 2024. Exploring Experiences with Automated Program Repair in Practice. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*. 1–11.

[9] Alena Moiseeva, Dietrich Trautmann, Michael Heimann, and Hinrich Schütze. 2020. Multipurpose intelligent process automation via conversational assistant. *arXiv preprint arXiv:2001.02284* (2020).

[10] Jan-Peter Ostberg, Daniel Graziotin, Stefan Wagner, and Birgit Derntl. 2017. Towards the assessment of stress and emotional responses of a salutogenesis-enhanced software tool using psychophysiological measurements. In *2017 IEEE/ACM 2nd International Workshop on Emotion Awareness in Software Engineering (SEmotion)*. IEEE, 22–25.

[11] Lena Ivannova Ruiz-Rojas, Luis Salvador-Ullauri, and Patricia Acosta-Vargas. 2024. Collaborative working and critical thinking: Adoption of generative artificial intelligence tools in higher education. *Sustainability* 16, 13 (2024), 5367.

[12] Wilmar B Schaufeli. 2017. Applying the Job Demands-Resources model: A 'how to'guide to measuring and tackling work engagement and burnout. *Organizational dynamics* 46, 2 (2017), 120–132.

[13] Johanna Schmidhuber, Stephan Schlögl, and Christian Ploder. 2021. Cognitive load and productivity implications in human-chatbot interaction. In *2021 IEEE 2nd International Conference on Human-Machine Systems (ICHMS)*. IEEE, 1–6.

[14] Auste Simkute, Lev Tankelevitch, Viktor Kewenig, Ava Elizabeth Scott, Abigail Sellen, and Sean Rintel. 2024. Ironies of Generative AI: Understanding and Mitigating Productivity Loss in Human-AI Interaction. *International Journal of Human–Computer Interaction* (2024), 1–22.

[15] Auste Simkute, Lev Tankelevitch, Viktor Kewenig, Ava Elizabeth Scott, Abigail Sellen, and Sean Rintel. 2025. Ironies of generative AI: understanding and mitigating productivity loss in Human-AI interaction. *International Journal of Human–Computer Interaction* 41, 5 (2025), 2898–2919.

[16] Lukas Teufelberger, Xintong Liu, Zhipeng Li, Max Moebus, and Christian Holz. 2024. LLM-for-X: Application-agnostic Integration of Large Language Models to Support Personal Writing Workflows. *arXiv preprint arXiv:2407.21593* (2024).

[17] Thiago S Vaillant, Felipe Deveza de Almeida, Paulo Anselmo Neto, Cuiyun Gao, Jan Bosch, and Eduardo Santana de Almeida. 2024. Developers' Perceptions on the Impact of ChatGPT in Software Development: A Survey. *arXiv preprint arXiv:2405.12195* (2024).

[18] Man-Fai Wong, Shangxin Guo, Ching-Nam Hang, Siu-Wai Ho, and Chee-Wei Tan. 2023. Natural language generation and understanding of big code for AI-assisted programming: A review. *Entropy* 25, 6 (2023), 888.

[19] Shun Zhang, Zhenfang Chen, Yikang Shen, Mingyu Ding, Joshua B Tenenbaum, and Chuang Gan. 2023. Planning with large language models for code generation. *arXiv preprint arXiv:2303.05510* (2023).