

Automating knowledge work while preserving human judgement: Chat Everywhere and AI-mediated learning platforms in the automated economy

Joao Portugal  

Abstract

This article uses Chat Everywhere as a case through which to examine how generative artificial intelligence platforms can automate parts of knowledge work while keeping consequential judgement under human responsibility. In this AI co-created article, the human author provided the case material, interview-based information, article objective, journal category, literature direction and critical editorial decisions. The AI system generated the main analytical interpretation, suggested conceptual connections and drafted the central discussion, after which the human author reviewed, edited and selected the final content. The case suggests that AI-supported learning platforms are not only educational tools. They are also examples of automated knowledge-work systems that redistribute repetitive cognitive tasks, such as first feedback, diagnostic analysis, meeting documentation and practice generation. The article argues that the value of these systems depends on three conditions: bounded automation, human review and organizational oversight. It also develops design-by-incentives as a practice principle for embedding verification inside the workflow. The article contributes to discussions on the automated economy by showing how an AI-supported learning platform can be interpreted as a site where automation, human judgement and governance are negotiated in practice.

Keywords: automated economy; AI-mediated work; human judgement; generative AI; Chat Everywhere; workflow automation; AI co-creation

Type: AI-co-created Article

Citation: Portugal, J. (2026). Automating knowledge work while preserving human judgement: Chat Everywhere and AI-mediated learning platforms in the automated economy. *ROBONOMICS: The Journal of the Automated Economy*, 7, 103. <https://doi.org/10.67723/rj.v7i.130>

¹ Instituto Superior Manuel Teixeira Gomes, Portugal; email: p6907@ismat.pt; jpvport@gmail.com

 Corresponding author

Publication history

Received: 21/06/2026; *Revised:* 15/09/2026; 21/09/2026; *Accepted:* 21/09/2026; *Published online:* 22/09/2026;
Volume date: 31/12/2026



© 2026 The Author(s)

This work is licensed under the Creative Commons Attribution 4.0 International (CC BY 4.0).

To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>

1. Introduction

Artificial intelligence is increasingly used in educational and organizational settings to produce text, summarize information, grade simple tasks, generate learning materials and support decision-making (Kasneci et al., 2023). These uses do not only change how students learn or how teachers prepare classes (Tlili et al., 2023). They also change the organization of knowledge work (Jarrahi et al., 2023). Tasks that previously required human time can now be partly delegated to AI systems, while humans remain responsible for interpretation, correction and final judgement. This redistribution is increasingly understood as a reconfiguration of knowledge work rather than a simple process of technological substitution. AI can perform or support specific cognitive tasks, while contextual understanding, interpretation and judgement remain distributed across human and technological actors (Jarrahi et al., 2023; Kintzel et al., 2026).

The distinction between automation and augmentation is therefore central to understanding AI-mediated work. Automation transfers task execution from humans to AI, whereas augmentation combines the capabilities of both actors. However, these two processes are not necessarily alternatives. Raisch and Krakowski (2021) conceptualize them as an automation–augmentation paradox, in which organizations may simultaneously automate some activities while increasing the importance of human contributions in others. This perspective is particularly relevant to knowledge-intensive work because AI offers advantages in processing information, identifying patterns and handling complexity, while humans retain comparative strengths in dealing with uncertainty, equivocality and contextual judgement (Jarrahi, 2018). Human–AI collaboration can consequently be understood as a problem of division of labour: the relevant organizational question becomes not simply whether a task can be automated, but how tasks and decision rights should be allocated between humans and AI systems (Puranam, 2021).

Recent research further suggests that this division of labour is dynamic and task-specific. Different configurations of AI characteristics can produce different forms of knowledge-work augmentation, meaning that the consequences of AI depend on how technological capabilities are combined with human roles rather than on the presence of AI alone (Kintzel et al., 2026). At the organizational level, effective AI adoption similarly requires complementary changes in people, infrastructures and processes (Jarrahi et al., 2023). Human expertise remains particularly important because knowledgeable users are better positioned to interrogate algorithmic outputs, identify errors and retain agency when working with AI systems (Hillebrand et al., 2025).

This article examines this shift through the case of Chat Everywhere (chateverywhere.app), an AI-supported platform that began as a tool for access to generative AI and later developed educational and workplace functions. The case was selected partly because the author's proximity to members of Chat Everywhere's management team facilitated access to detailed information about the platform and its development. The original case material describes a platform that includes classroom access, AI-graded quizzes, diagnostic feedback, pronunciation practice, meeting summaries and organization-level controls. This makes Chat Everywhere useful not only as an educational technology case but also as an example of AI-mediated knowledge-work automation. The case is particularly informative because these functions place different cognitive activities—generation, assessment, diagnosis, documentation and review—within the same technological environment, while retaining different forms of human involvement and oversight. It therefore provides an opportunity to examine how cognitive tasks are automated, reviewed and governed at the workflow level rather than treating AI adoption as a single organizational decision.

This AI co-created article uses the case to explore a broader question: how can organizations automate repetitive cognitive tasks without removing human judgement from consequential decisions? In this sense, the article moves from a narrow discussion of educational technology to a wider discussion of the automated economy. The focus is not on whether AI replaces human work in general, but on how work is redistributed between AI systems, human users and organizations. This approach also recognizes that human oversight is not

achieved merely by keeping a person formally “in the loop”. Meaningful oversight requires humans to retain the capacity to evaluate AI outputs and intervene when necessary, while organizational arrangements define where authority and accountability ultimately reside (Laux, 2024; Zhu et al., 2026). AI governance is therefore part of work design: organizations must determine not only what AI is permitted to do, but also how its outputs are reviewed, challenged and incorporated into consequential decisions (Almeida & dos Santos Júnior, 2025; Papagiannidis et al., 2025).

Recent literature shows that AI can support explanation, tutoring, feedback and learning processes (Kasneci et al., 2023; Tlili et al., 2023). Other work stresses the importance of responsible, human-centred AI practices in education (Fu & Weng, 2024). In the field of learning and development, Dastane et al. (2024) show that emerging technologies are already changing how organizations approach training and development. Generative AI extends this transformation because it automates not only environments or simulations, but also parts of writing, assessment, analysis and communication. Seen from this broader perspective, AI-supported learning platforms are also organizational systems for allocating cognitive work. Their significance lies not only in what AI can generate or assess, but in how automation, human judgement, verification and organizational control are combined within everyday workflows. This provides the theoretical basis for examining Chat Everywhere as a case of AI-mediated knowledge-work automation within the automated economy.

2. Methodology: human-AI co-creation process

The human author supplied the case material, including a structured written questionnaire completed by a senior representative of Chat Everywhere, platform information, aggregated use indicators and prior article drafts. The questionnaire, entitled Chat Everywhere Case Study Questionnaire, was not generated by the AI system. It was designed to collect primary case information and was completed in writing by the company representative. The completed questionnaire was subsequently provided to the AI system as source material for the case analysis. The questionnaire is provided as Supplementary Material 2 for transparency. The human author also defined the target journal category, the main objective of the Robonomics version and the need to distinguish this article from a practice paper focused on organizational learning.

The AI system was used to generate an analytical reinterpretation of the case. The task given to the AI was to move the argument from educational technology to automated knowledge work, while keeping the case evidence grounded in the original interview material. The AI was asked to identify how specific platform functions illustrate the redistribution of cognitive tasks, human judgement, verification and governance. The AI also generated draft sections, refined the structure and proposed the use of design-by-incentives as a concept for describing verification mechanisms within workflows. The initial AI co-creation process was conducted between April and May 2026 using ChatGPT (OpenAI), with the GPT-5.5 model. The subsequent editorial revision of the manuscript was conducted using ChatGPT with the GPT-5.6 Sol model. The exact prompts used in the development and revision of the article are reproduced verbatim in the Supplementary Material 1.

The human author reviewed the generated content, selected the final argument, checked the use of references, removed claims that were not supported by the case material and ensured that the article did not present the case as an impact evaluation. The prompts used in the AI co-creation process are reproduced in full in Supplementary Material 1 and are explicitly linked to the AI-generated content presented in Section 3. This follows the logic of AI co-created articles in Robonomics, where the AI process, prompts and human reflection are part of the scholarly contribution (Tarbit et al., 2023).

The AI-generated content is mainly presented in Section 3. The Discussion and Conclusion then provide a human critical reflection on the generated interpretation. The AI system is not listed as an author because responsibility for the article, its claims and its final form remains with the human author.

3. AI-generated content: Chat Everywhere as a case of AI-mediated knowledge-work automation

The entire text in this section was generated by the AI system, without human rewriting, in response to the prompts reproduced verbatim in the Supplementary Material I. The human author did not rewrite or edit the AI-generated text presented in this section.

Chat Everywhere can be interpreted as a platform that reorganizes knowledge work around a set of repeatable AI-mediated tasks. In its educational use, the platform allows teachers to create classroom activities, deploy AI-graded quizzes and obtain first-level diagnostic information. In its organizational use, it can summarize meetings, produce transcripts and support documentation. These functions appear different at first sight, but they share the same logic. In each case, AI is used to process information, generate a first output and reduce the amount of repetitive work required from humans.

This is a common pattern in the automated economy. Automation does not always mean full replacement of human labour. In many knowledge-work settings, automation changes the sequence of work. A human defines the task, the system produces an initial output, and the human reviews, corrects or uses that output to make a decision. Raisch and Krakowski (2021) describe this as part of the tension between automation and augmentation in AI. Jarrahi (2018) similarly argues that human-AI collaboration often works through complementary strengths: machines process large amounts of information, while humans provide judgement, context and interpretation.

The Chat Everywhere case illustrates this division of labour. The platform can grade a quiz, but the teacher decides what the result means. It can produce feedback, but the teacher judges whether the feedback is useful for the student. It can identify weak areas in a class, but the teacher decides whether the class needs revision, more practice or a different explanation. It can summarize a meeting, but the organization decides whether the summary is accurate, appropriate and usable. The automated system produces information; the human actor gives that information meaning.

This boundary is important because knowledge work often involves consequential judgement. An incorrect summary may lead to a misunderstanding. A weak piece of feedback may shape a learner's confidence. A grade may affect how a student understands performance. An institutional dashboard may influence how managers view adoption or engagement. For this reason, AI-mediated systems need mechanisms that make review normal rather than exceptional. Automation creates value when its outputs are easy to inspect, question and revise.

The most useful concept generated from the case is design-by-incentives. The expression refers to workflows in which users have a practical reason to check AI outputs. In a quiz, a student may dispute an AI-generated grade because the grade matters. The student therefore has a reason to read the feedback and examine the reasoning. The teacher then has a reason to review the dispute and identify whether the system, the student or the task caused the problem. In this design, verification is not simply recommended. It is built into the activity. Design-by-incentives is relevant for the automated economy because many organizations rely on policies that tell employees to use AI responsibly. Policies are necessary, but they can remain external to daily work. A platform workflow can create a more direct form of responsibility when checking becomes useful for the person performing the task. This is different from a general warning that AI may be wrong. It gives users a reason to test the output because their own work, grade, decision or organizational process depends on it.

The case also shows that AI-mediated platforms create new governance needs. The Organization Portal in Chat Everywhere is not only an administrative function. It represents a layer of organizational control over accounts, permissions, visibility and data use. In automated work environments, governance cannot be added only after the system is deployed. It has to be part of the design. This includes deciding who can create activities, who can view outputs, who can correct them and who is responsible when AI-generated content is used.

The broader implication is that AI-supported platforms should be understood as organizational infrastructures rather than isolated tools. They automate tasks, but they also structure relationships between users, managers, teachers, learners and institutions. They make some activities faster and more visible, but they also require new routines of checking, correction and responsibility. In the automated economy, the relevant question is not only what AI can do. The relevant question is how work should be organized around AI outputs.

In this case, bounded automation appears to be the safest and most practical starting point. The strongest tasks are those that are repetitive, clearly defined and reviewable: first grading, diagnostic feedback, practice generation, pronunciation practice, transcript production and meeting summaries. These tasks do not remove the need for human judgement. Instead, they move human effort from producing first drafts or first evaluations to reviewing, interpreting and improving outputs.

The AI-generated interpretation therefore suggests three principles. First, organizations should begin with tasks where the expected output is clear and where errors can be detected. Second, consequential judgement should remain with humans who understand the context. Third, verification should be part of the workflow rather than an optional afterthought. These principles are not specific to Chat Everywhere. They may apply to many AI-mediated platforms in education, training and organizational development.

4. Human reflection on the AI-generated interpretation

The AI-generated content was useful in shifting the case away from a narrow description of an educational platform and towards a broader analysis of knowledge-work automation. This reframing is consistent with a growing view of AI at work as a reconfiguration of tasks and human–technology relationships rather than simply the replacement of human activity (van den Broek, 2025). It identified a clear connection between classroom functions, meeting documentation and organizational governance. Viewed in this way, these are not simply separate product features, but different configurations of human–AI collaboration in which cognitive tasks are distributed between technological and human actors (Puranam, 2021). Repetitive cognitive work can be delegated to AI-supported systems, while interpretation, exception handling and consequential decisions can remain under human control.

At the same time, the AI-generated interpretation required human review. The first drafts tended to present the case as more generalizable than the evidence allowed. They also moved quickly from platform description to broad claims about the automated economy. This experience illustrates a broader limitation of treating human involvement as a sufficient safeguard in AI-mediated work. Human oversight depends on the ability to question, evaluate and, when necessary, reject AI-generated outputs rather than merely approving them (Laux, 2024; Zhu et al., 2026). Human editing was therefore used to keep the case grounded in the questionnaire material and to present design-by-incentives as an interpretive practice principle rather than as a validated theory.

The co-creation process itself reflects the argument of the article. Rather than representing either complete automation or purely human authorship, it involved a division of cognitive labour between the AI system and the human author. The AI system produced structure, language and conceptual connections. The human author selected the purpose, checked the argument, controlled the final scope and accepted responsibility for the output. This resembles the automation–augmentation relationship described by Raisch and Krakowski (2021): automating particular cognitive activities does not necessarily eliminate human contribution but can redistribute where that contribution occurs. In that sense, the writing process reproduced the same pattern described in the case: AI generated an initial knowledge product, while human judgement was used to interpret, correct and finalize it.

This also shows one limitation of AI co-created scholarship. AI can help reframe a case and produce fluent text, but it can also smooth over uncertainty, exaggerate the coherence of a case and create conceptual language faster than evidence can support. Such risks reinforce the importance of retaining human evaluative agency when AI is used for knowledge production, particularly where outputs require contextual interpretation rather than routine execution (Jarrahi et al., 2023). For this reason, AI co-created articles should not only display machine-generated content. They should also reflect on how that content was produced, checked and bounded. The relevant issue is therefore not simply whether a human participated in the process, but what role human judgement played and where responsibility for the resulting knowledge remained.

5. Implications for Robonomics and the automated economy

The case has three implications for research on the automated economy. First, AI-mediated automation should be studied at the workflow level. The important unit is not the isolated tool but the sequence of task definition, AI output, human review and organizational use. This perspective extends task-centred interpretations of automation by emphasizing the interdependencies created when AI becomes embedded across sequences of knowledge work. Hillebrand et al. (2025) similarly argue for moving beyond isolated human–AI interactions towards a systems perspective that considers tasks, organizational context, collective agency and multilevel outcomes. From this perspective, the consequences of automation cannot be inferred solely from whether AI performs a particular task; they also depend on how that output enters subsequent human and organizational processes. Chat Everywhere illustrates how the same platform logic can operate across education and organizational documentation. AI-graded quizzes, diagnostic feedback and meeting summaries automate different activities, but each produces an output that can subsequently be interpreted, challenged or used by another actor. This supports van den Broek's (2025) argument that AI at work involves interconnected forms of data, knowledge and values work rather than a linear movement from data to automated decisions. The case therefore suggests that workflow configuration is an important level of analysis for Robonomics because automation in one part of a process can redistribute rather than eliminate work elsewhere in the system.

Second, human judgement remains a central resource in automated systems. The case does not support a simple replacement narrative. Instead, it illustrates the interdependence between automation and augmentation identified by Raisch and Krakowski (2021). Delegating bounded cognitive activities to AI can reduce direct human involvement in task execution while simultaneously increasing the importance of human evaluation at points where outputs require contextual interpretation. The case suggests that automation is most useful when the task is bounded and when the output can be checked by someone with contextual knowledge. This division of cognitive labour is consistent with Puranam's (2021) treatment of human–AI collaboration as an organization-design problem involving task allocation, interdependence and the integration of human and algorithmic contributions. It is also compatible with Jarrahi's (2018) account of human–AI symbiosis, in which computational capabilities complement rather than eliminate human strengths in dealing with uncertainty and equivocality. This supports the view that AI changes the distribution of work rather than only reducing the amount of work. For the automated economy, the relevant question may therefore be less whether human judgement remains involved and more where that judgement is positioned within an increasingly automated workflow.

The case adds a further dimension to this argument through verification. Human involvement has limited value if users routinely accept AI outputs without evaluating them. Research on meaningful human oversight similarly indicates that placing a human formally in or on the loop does not by itself guarantee effective control; users must retain the competence and opportunity to evaluate, contest and redirect automated outputs (Laux, 2024; Zhu et al., 2026). Chat Everywhere suggests an additional mechanism that can be described as design-by-incentives. In some workflows, users have practical reasons to verify AI outputs because those outputs affect their immediate interests. For example, a student may challenge an AI-generated grade, creating a reason for the teacher to review the assessment. Verification is therefore encouraged not only through warnings or formal oversight requirements but also through the structure of the interaction itself. This interpretation finds adjacent

support in Luong et al. (2026), who show experimentally that organizational incentive alignment can influence human–AI collaboration. However, the present case does not establish design-by-incentives as a validated mechanism. It identifies it as an emergent practice principle that requires direct empirical testing.

Third, governance should be treated as part of automation design. Permissions, visibility, dispute mechanisms and data rules are not secondary features. They shape how people trust, contest and use AI outputs. This interpretation is consistent with responsible AI governance research that moves beyond abstract ethical principles towards the organizational structures and processes through which AI is actually governed. Papagiannidis et al. (2025), for example, distinguish structural, procedural and relational governance practices, including the allocation of responsibilities and decision rights, organizational processes and the development of relevant human capabilities. Similarly, Almeida & dos Santos Júnior (2025) show that AI governance is operationalized through practices distributed across strategic, tactical and operational levels. Automated systems therefore need organizational arrangements that define responsibility before errors or conflicts appear. In the Chat Everywhere case, organization-level controls can consequently be interpreted not merely as administrative functionality but as part of the architecture through which automated work is bounded and supervised.

Taken together, these implications suggest that automation in the knowledge economy is not adequately described as a transfer of tasks from humans to machines. The case instead points to a reconfiguration of cognitive work across AI execution, human verification and organizational governance. Its theoretical contribution is therefore not to demonstrate that one configuration is universally superior, which the evidence cannot establish, but to identify three interconnected dimensions through which AI-mediated automation can be examined: workflow configuration, the positioning of human judgement and governance architecture. This perspective also connects Robonomics with the broader human–AI collaboration literature by treating human judgement and governance not as residual safeguards after automation, but as constitutive elements of how automated knowledge work is organized.

6. Conclusion

This article used Chat Everywhere as a human-AI co-created case analysis of knowledge-work automation. The case shows how an AI-supported learning platform can be interpreted as part of the automated economy. It automates first outputs in tasks such as feedback, grading, diagnostics, practice generation and meeting documentation, while keeping human judgement important for interpretation and final decisions. This configuration reinforces the view that automation and augmentation can coexist within the same workflow, with AI assuming particular cognitive activities while human contribution shifts towards evaluation, interpretation and decision-making (Raisch & Krakowski, 2021).

The main contribution is the idea that AI-mediated platforms should not be evaluated only by how much work they automate. They should also be evaluated by how they organize checking, responsibility and institutional control. Design-by-incentives captures this point by describing workflows in which users have a practical reason to inspect AI outputs. This complements research showing that meaningful human oversight depends on the capacity and opportunity to evaluate automated outputs rather than simply placing a human formally in the loop (Laux, 2024; Zhu et al., 2026). In such systems, responsible automation is not only a policy statement. It becomes part of how the work is done. However, design-by-incentives should be understood here as an emergent interpretive principle derived from the case, not as a validated theoretical mechanism.

The article also shows that AI co-created writing can be used as a method for conceptual case analysis, provided that the process is transparent and that human review remains central. The writing process and the case both point to the same conclusion: in the automated economy, the value of AI depends less on replacing humans and more on designing systems where humans and AI perform different parts of the work responsibly. This shifts

attention from the amount of automation achieved towards the configuration of human–AI work and the organizational arrangements through which AI outputs are reviewed and governed.

Several limitations qualify these conclusions. First, the analysis is based on a single illustrative case and should not be generalized to other AI-supported platforms or organizational contexts. Second, the evidence consists primarily of structured case material, platform documentation and aggregated usage indicators. These indicators demonstrate adoption and use but do not establish learning effectiveness, productivity gains or other causal outcomes. Third, the author's proximity to members of Chat Everywhere's management team facilitated access to detailed case information but may also introduce interpretive proximity that should be acknowledged when considering the analysis. Finally, because AI contributed to the conceptual interpretation and drafting of the article, the analysis is partly shaped by the same technology it examines. Human review was used to constrain this process, but it does not eliminate the methodological implications of AI-assisted interpretation.

Future research should therefore test the arguments developed here beyond this single case. Comparative and multi-case studies could examine whether similar configurations of automation, human judgement and governance emerge across educational platforms, workplaces and other knowledge-intensive settings. Empirical studies could also test design-by-incentives directly by examining whether workflow incentives actually increase verification, reduce inappropriate reliance on AI outputs or improve the quality of human–AI decisions. Longitudinal research could investigate how task allocation and human oversight change as users gain experience with AI systems, while experimental studies could compare different levels and forms of human intervention. Finally, future research on AI co-created scholarship could examine how different configurations of human and machine contribution affect analytical quality, transparency and reproducibility. These directions would allow the propositions emerging from this case to move from conceptual interpretation towards empirical validation.

References

- Almeida, P. G. R. de, & dos Santos J[unior], C. D. (2025). Artificial intelligence governance: Understanding how public organizations implement it. *Government Information Quarterly*, 42(1), 102003. <https://www.doi.org/10.1016/j.giq.2024.102003>
- Dastane, O., Rafiq, M., & Turner, J. J. (2024). Implications of metaverse, virtual reality, and extended reality for development and learning in organizations. *Development and Learning in Organizations: An International Journal*, 38(5), 27–32. <https://www.doi.org/10.1108/DLO-09-2023-0196>
- Fu, Y., & Weng, Z. (2024). Navigating the ethical terrain of AI in education: A systematic review on framing responsible human-centered AI practices. *Computers and Education: Artificial Intelligence*, 7, 100306. <https://www.doi.org/10.1016/j.caeai.2024.100306>
- Hillebrand, L., Raisch, S., & Schad, J. (2025). Managing with artificial intelligence: An integrative framework. *Academy of Management Annals*, 19(1), 343–375. <https://www.doi.org/10.5465/annals.2022.0072>
- Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human–AI symbiosis in organizational decision making. *Business Horizons*, 61(4), 577–586. <https://www.doi.org/10.1016/j.bushor.2018.03.007>
- Jarrahi, M. H., Askay, D., Eshraghi, A., & Smith, P. (2023). Artificial intelligence and knowledge management: A partnership between human and AI. *Business Horizons*, 66(1), 87–99. <https://www.doi.org/10.1016/j.bushor.2022.03.002>
- Kasneci, E., Sessler, K., K[ühnemann], S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., G[ün]demann, S., H[il]lmeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., ... Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274. <https://www.doi.org/10.1016/j.lindif.2023.102274>
- Kintzel, K., Clausen, S., & Stieglitz, S. (2026). Knowledge work augmentation with AI: A taxonomy of AI-based technology characteristics and their implications for meaningful work. *Behaviour & Information Technology*. Advance online publication. <https://www.doi.org/10.1080/0144929X.2026.2679599>
- Laux, J. (2024). Institutionalised distrust and human oversight of artificial intelligence: Towards a democratic design of AI governance under the European Union AI Act. *AI & Society*, 39(6), 2853–2866. <https://www.doi.org/10.1007/s00146-023-01777-z>
- Luong, A., Kumar, N., & Lang, K. R. (2026). Human–machine collaboration and algorithmic decision-making: The role of organizational incentives and AI capabilities. *Information & Management*, 63(6), 104384. <https://www.doi.org/10.1016/j.im.2026.104384>
- Papagiannidis, E., Mikalef, P., & Conboy, K. (2025). Responsible artificial intelligence governance: A review and research framework. *The Journal of Strategic Information Systems*, 34(2), 101885. <https://www.doi.org/10.1016/j.jsis.2024.101885>
- Puranam, P. (2021). Human–AI collaborative decision-making as an organization design problem. *Journal of Organization Design*, 10(2), 75–80. <https://www.doi.org/10.1007/s41469-021-00095-2>

- Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review*, 46(1), 192–210. <https://www.doi.org/10.5465/amr.2018.0072>
- Tarbit, J., Wirtz, J., Kunz, W., & Hartley, N. (2023). Interpretation of corporate digital responsibility risks and concerns by automated service technologies: An AI co-created article. *ROBONOMICS: The Journal of the Automated Economy*, 4, 52.
- Tili, A., Shehata, B., Adarkwah, M. A., Bozkurt, A., Hickey, D. T., Huang, R., & Agyemang, B. (2023). What if the devil is my guardian angel: ChatGPT as a case study of using chatbots in education. *Smart Learning Environments*, 10, 15. <https://www.doi.org/10.1186/s40561-023-00237-x>
- van den Broek, E. (2025). Unpacking AI at work: Data work, knowledge work, and values work. *Information and Organization*, 35(3), 100584. <https://www.doi.org/10.1016/j.infoandorg.2025.100584>
- Zhu, L., Lu, Q., Ding, M., Lee, S. U., & Wang, C. (2026). Designing meaningful human oversight in AI. *AI and Ethics*, 6, 286. <https://www.doi.org/10.1007/s43681-026-01147-7>

Supplementary Material I

Prompts and procedural documentation for the Robonomics AI co-created article

This supplementary material documents the prompts and instructions used during the AI co-creation and editorial revision of the article on Chat Everywhere and AI-mediated knowledge-work automation. For transparency, the prompts submitted to the AI system are reproduced separately from the human author's procedural notes.

A. Prompts submitted to the AI system

Prompt 1. General researcher persona

Assume the role of a senior researcher, academic editor and scientific writer specialised in artificial intelligence, automated economy, organizational learning, educational technology, human-AI collaboration, knowledge-work automation and responsible innovation. Write in clear academic English, with conceptual precision but without unnecessary jargon. Treat the human author as the final decision-maker. Do not invent empirical evidence. Use only the provided case material, structured interview information, platform documentation and verified literature. The goal is to transform case evidence into a publishable scholarly article while preserving transparency about the use of AI.

Prompt 2. Article-specific adaptation

Task: Develop a final manuscript for ROBONOMICS: The Journal of the Automated Economy, to be submitted under the section AI co-created article. The manuscript must use Chat Everywhere as an illustrative case of AI-mediated knowledge-work automation. Do not frame the article primarily as an educational technology practice paper. Reframe the case around automated economy, cognitive automation, human judgement, workflow reconfiguration, verification and governance.

Source material: use the structured interview with a senior Chat Everywhere representative, platform documentation, aggregated platform use data and the previous long-form article on Chat Everywhere. The case includes classroom access, AI-graded quizzes, diagnostic feedback, personalised practice, pronunciation practice, meeting summarization and organizational controls.

Analytical objective: explain how AI-supported platforms redistribute repetitive cognitive tasks between AI systems, human users and organizations. Show that AI produces first outputs while humans retain contextual interpretation and consequential judgement. Develop design-by-incentives as an interpretive practice principle for workflows in which users have a practical reason to check AI outputs.

Required structure: Introduction; Methodology: human-AI co-creation process; AI-generated content; Human reflection on the AI-generated interpretation; Implications for Robonomics and the automated economy; Conclusion; References.

Transparency requirements: make clear that the article was co-created with AI. Explain the role of the human author and the AI system. State that the AI generated the main analytical interpretation and that the human author reviewed, selected, corrected and finalized the manuscript. Do not list AI as an author.

Style requirements: academic but readable; concise paragraphs; no promotional language; no unsupported claims about learning impact or platform effectiveness; no title containing the phrase 'an AI co-created article'. Include the article type in the manuscript metadata instead.

*Reference requirement: include and use the following reference where conceptually relevant: Dastane, O., Rafiq, M. and Turner, J.J. (2024), 'Implications of metaverse, virtual reality, and extended reality for development and learning in organizations', *Development and Learning in Organizations: An International Journal*, Vol. 38 No. 5, pp. 27-32, doi: 10.1108/DLO-09-2023-0196.*

Prompt 3. Diagnostic reframing

Read the Chat Everywhere case material and identify which parts relate to education, which parts relate to organizational learning, and which parts can be reinterpreted as knowledge-work automation.

Prompt 4. Robonomics fit

Reframe the case for ROBONOMICS by focusing on automated economy, human-AI task distribution, cognitive automation and governance. Identify a distinct contribution that avoids duplicating a practice paper on education.

Prompt 5. AI-generated analytical section

Generate the central AI-generated content section explaining Chat Everywhere as a case of AI-mediated knowledge-work automation. Keep the analysis grounded in the provided case material.

Prompt 6. Human reflection

Draft a human critical reflection on the AI-generated interpretation, including what the AI analysis contributed and where human review was needed.

Procedural clarification: This prompt was used to generate possible points for critical reflection. The resulting AI-generated draft was not used as the final text of Section 4. The final Human Reflection presented in Section 4 was written by the human author.

Prompt 7. Final article assembly

Assemble the final manuscript with abstract, keywords, article type, methodology, AI-generated content, human reflection, implications, conclusion and references. Keep the title free from the phrase "AI co-created article".

Prompt 8. Editorial check

Review the manuscript for overlap with prior educational practice-paper versions. Strengthen the Robonomics contribution and remove language that suggests a formal impact evaluation.

Prompt 9. Editorial revision in response to editor feedback

Review the manuscript in light of the editor's comments and assist in preparing a revised version. Address each editorial comment directly while preserving the article's central argument and the distinction between the AI-generated analytical section and the human critical reflection. Strengthen the theoretical background where requested, integrate relevant and verified academic references, clarify the human-AI co-creation process, expand the theoretical implications, and elaborate the limitations and future research directions. Do not introduce empirical claims that are not supported by the case material. Treat design-by-incentives as an emergent interpretive principle rather than a validated theoretical mechanism. Check that all in-text citations correspond to the final reference list and format the references according to APA 7th edition, including DOI links where available. Preserve the human author's role as the final decision-maker regarding all revisions and responses to the editor.

B. Human procedure and editorial control

The human author selected the article objective, supplied the case material, checked the final argument, approved the references, controlled the interpretation of the questionnaire responses, removed unsupported claims and accepted responsibility for the final manuscript. The AI system was used to generate structure, wording, conceptual links and the analytical interpretation presented in Section 3, but it was not treated as an author.

The AI system was also used to support the revision of the manuscript following editorial feedback. The editor's comments were provided to the AI system and addressed individually. The AI assisted with identifying possible revisions, integrating additional literature, refining academic wording, expanding selected sections and checking consistency between in-text citations and the reference list. The human author evaluated the proposed changes, selected the revisions to retain and remained responsible for the final responses and revised manuscript.

Supplementary Material 2

Chat Everywhere Case Study Questionnaire

(Robonomics Practice Paper)

About the journal: <https://journal.robonomics.science/index.php/rj/about>

João Pedro Portugal, Ph.D

Documentation: <https://docs.v2.chateverywhere.app/en>

I. Product origin, positioning and evolution

To describe Chat Everywhere accurately in the paper, could you please summarise the platform's origin and evolution, covering the points below (short bullet answers are fine):

Year of creation/launch and key milestones (e.g., first release, v2 launch, major product redesigns, education/business positioning, institutional adoption).

Chat Everywhere launched in late 2023 (first landing page, October 2023) as an education-first AI platform that later expanded into business productivity.

Key milestones:

2023: Launch; awarded Azure startup funding; approved onto Taiwan's government procurement vendor list.

2024: First school purchases; the major **v2 redesign** (private beta Sept, public beta Oct) introduced the portal-based architecture; Prompt Library launched; year-end sales roughly 7 x 2023.

2025: v2 Teacher Portal showcased; first municipal/provincial deal (Chiayi City) and university deployment interest (Tamkang University); Meeting Summarizer paid launch; BenQ partnership outreach.

2026: LINE bot launched.

On positioning: it started education-first, then expanded into white-collar use via the **Meeting Summarizer** (audio/video up to 5 GB → transcript, summary, mind map, and report in minutes; its edge is large-file and wide-format support, speed, and everything in one app). Institutional adoption has been guided by the principle of **enhancing, not replacing, teachers** — automating tedious grading and enabling tailored quizzes, feedback, and personalised practice.

Who initially sponsored/initiated the product (e.g., Explorator Labs, internal development team, educator need, institutional partner, specific stakeholder) and the original target user profile.

Chat Everywhere was self-funded internally by **Explorator Labs** — no external sponsor. Its origin was personal: ChatGPT wasn't accessible in the Macau/Hong Kong region, so the founder built a simple AI "wrapper" so his parents in Macau could use it. Explorator Labs' custom-software client revenue easily covered the modest API costs. The **original target users were, quite literally, the founder's parents**.

The initial problem it was built to solve (the concrete pain point in access to AI tools, classroom use, productivity, multilingual support, prompt reuse, or digital learning workflows).

The first problem was simply **access to AI** — giving the founder's parents AI access in a region where it was unavailable. It wasn't about education initially.

The move into education came from two sources: the founder's mother (a teacher) spotting classroom applications, and a Taiwanese education influencer, **David**, who found the tool via a Facebook group and suggested educator-focused features. As more teachers adopted it and gave feedback, the team built what they asked for. For example, the **Classroom feature** lets students join via a QR or numeric code instead of creating accounts — cutting a 3–5 minute login down to seconds in a 40-minute class.

The first core use case ("day 1" functionality) and what has changed since.

Day one (education): the Classroom feature. Students join a teacher's room via a code, instantly get full paid-tier access (including image generation), and share work in one click — ideal for engaging classroom activities. It remains a core use case.

Since then the focus has shifted toward teacher workflows:

Grading homework and quizzes.

Chat and breakout rooms for AI-assisted group work, with teacher monitoring.

In-class quizzes: AI-graded in minutes, with class-wide analytics showing the weakest points and where to focus teaching.

Pronunciation Practicer: scenario-based sentences with word-by-word audio and text feedback.

Three major evolution milestones (e.g., new portals, file reader, meeting summariser, prompt library, teacher tools, organisation management, model access, image generation, institutional features) and why they happened.

1. **Meeting Summarizer** — opened the business market: drop in a recording (up to 5 GB) and get a transcript, summary, mind map, and report in minutes. Built because user feedback revealed high-value business pain points, and large-file/format support plus speed gave a clear edge.

2. **In-class Quiz** — made personalised learning practical at scale through instant AI grading and class-wide analytics, which was previously impossible for the average student.

3. **White-Label Solution** — enabled institutional deployment: the full suite on the institution's own domain in a private cloud, removing account-management burden. One partner university (~20,000 staff and students) gained instant access.

Current state in 2–3 sentences: what Chat Everywhere does, where it fits in educational/professional workflows, and what it is not intended to do.

Today Chat Everywhere serves two workflows: in **education**, students take AI-graded quizzes, get tailored feedback in minutes, and practise weak concepts while teachers gain class-wide insight; in **business**, the Meeting Summarizer turns recordings into transcripts, summaries, mind maps, and reports. Its purpose is to remove tedious, repetitive work and make AI affordable and accessible, especially for educators. It is explicitly **not intended to replace educators** — only to support them.

One short illustrative example of a typical user or situation that captures this evolution (no sensitive or confidential details).

A middle-school English teacher opens a Classroom; students join in seconds via a QR code, with no accounts to set up. She starts with a short AI-generated quiz on last week's grammar, and within minutes it's graded — her Teacher Portal shows that most of the class missed the same rule. She has each student work on a unique AI-generated practice question targeting that gap, with instant feedback, then submit

it back to her. The class collaborates in breakout chat rooms, and later students drill speaking individually on the Pronunciation Practicer. In one period she has run a diagnostic, personalised practice for every student, and provided tailored speaking practice — work that wasn't possible at that scale before.

2. Core job-to-be-done today

What is the single most important job-to-be-done Chat Everywhere solves best today, and how do you define it operationally (i.e., what “success” looks like in practice)?

The job Chat Everywhere does best is **tailored, AI-graded quizzing**. What sets it apart is the range of accepted answers — including **handwritten work**: a student can photograph a maths solution on paper and the AI identifies which step is wrong, why, and how it connects to the broader concept.

Operationally, success means this becomes the **rhythm of the class**: a quick quiz at the start of each lesson to check prior understanding, and homework delivered as AI quizzes with instant feedback and personalised follow-up questions. Students improve through continuous practice and understanding their mistakes.

3. User groups, portals and workflow map

Which user groups and workflow stages does Chat Everywhere support (e.g., students, teachers, schools/training centres, business teams, individual professionals), and what is the rationale for prioritising these groups and stages?

Chat Everywhere supports **teachers, students, and individual professionals** (via the Meeting Summarizer), plus **schools and universities** at the institutional level. It spans the full teaching cycle — diagnostics, practice, grading, and analytics — and the post-meeting documentation stage for professionals.

These groups were prioritised because education concerns the next generation and teachers face rising workloads, and because professionals lose significant time manually writing up meetings. Both groups **benefit most from AI yet are least familiar with it**, so the product is built to be extremely easy to use with an obvious payoff.

Please indicate how the Worker Portal, Teacher Portal and Organization Portal differ in purpose, user permissions, governance needs, and expected outcomes.

(The “Worker Portal” corresponds to the **General Use** portal.)

General Use portal (“Worker Portal”): a ChatGPT-like general assistant for anyone, with convenience touches such as one-click mind maps. No special role required.

Teacher Portal: all teacher features — the Classroom (whole-class paid access without individual accounts), AI in-class quizzes, and the Pronunciation Practicer. Teacher-level access; outcome is less repetitive workload and personalised practice.

Organization Portal: account and usage management for IT admins — usage tracking, account management, password resets. It also provides org-wide conversation visibility for accounts added to the organisation (data stays in our cloud). This is the institution's governance layer.

Visibility differs by portal: the Teacher Portal sees only conversations inside a classroom a teacher created, whereas the Organization Portal sees all conversations of org-added accounts.

4. Main features and prioritisation logic

Which features are strategically central today (e.g., AI chat, file reader, meeting summariser, code runner, mind map generator, web analyser, image generator, quiz creator, pronunciation tool, prompt library, templates, shared prompts), and why were these prioritised?

The **AI quiz** is the central feature and the hub linking to general chat, the Classroom, and the Pronunciation Practicer. It was prioritised because **tailored grading** — accepting student work (including handwriting) and giving step-by-step, concept-linked feedback — is exactly what ChatGPT and Gemini don't do. The **Pronunciation Practicer** is the other central feature: both an on-ramp to AI and something general chatbots don't offer.

Which features are experimental, secondary, or designed for specific user segments rather than for the whole platform?

Homework Grader (experimental): unlike the standard quiz, it needs no setup — a teacher scans paper homework and it detects each student's name, identifies the questions, grades them, and generates a per-student report. It bridges paper-based classrooms into AI without forcing teachers to change their workflow; it's experimental because a larger roadmap exists but resources are currently limited.

Secondary features (less development effort, not abandoned): sharing student-created HTML, the code runner, and image generation.

5. Human–AI task boundaries

Which tasks are designed to be automated with low supervision, and which tasks should always require human judgement? What is the reasoning behind these boundaries?

Automated, low supervision: tasks with clearly defined answers — grading quizzes, analysing class performance, generating practice questions, and transcribing/summarising meetings.

Always human judgement: deciding where to focus teaching. The AI may miss context the teacher holds — knowing a concept is hard but will resolve once a prerequisite clicks, or knowing each student's quirks — plus all the human interaction.

The boundary exists because the goal is to **enhance, not replace, educators**: automate the repetitive, well-defined work so teachers can spend time on the human connection a machine can't replicate.

Please distinguish between low-risk productivity support (e.g., summaries, organisation, drafts, quizzes) and higher-risk uses (e.g., assessment, feedback, educational guidance, institutional decision-making, compliance-sensitive outputs).

Low-risk productivity support: meeting summaries, question and draft generation, organisation — easily corrected if wrong.

Higher-risk: grading, where the output must be correct, and institutional decision-making, where AI should only **inform** the decision, never make it — humans catch what the model misses.

6. Productivity/efficiency metrics

Which internal metrics do you use to assess productivity gains (e.g., time saved in meeting notes, faster document review, faster quiz generation, fewer repeated prompts, reduced administrative workload), and which 2–3 KPIs do you consider most robust?

We frame business use as "productivity" and education as "effectiveness":

Productivity (business): total meeting duration transcribed, and whether a report was generated.

Education effectiveness: quizzes sent, answers submitted, personalised quizzes generated, and overall usage; engagement is currently proxied by quiz volume.

Most robust KPIs: (1) total minutes transcribed; (2) quizzes sent vs. answered; (3) personalised quizzes generated. Several of these are usage-based proxies rather than direct learning-outcome measures.

7. Learning, teaching and work-quality metrics

Which indicators do you use to assess quality of output or user value (e.g., clarity, usefulness, accuracy, pedagogical fit, task completion, reuse rate, student engagement, teacher workload reduction, meeting-summary usefulness), and how do you detect or prevent quality degradation over time?

We found no model reliably measures output quality — it's highly subjective and varies by market and language, especially in education where wording changes how a student feels. So we rely on our many **heavy users**: after each update we watch for their feedback and fix major issues within hours. Quality degradation is detected mainly through **complaint volume**, and the tight loop lets us **roll back** quickly. This assessment is qualitative and feedback-driven rather than a formal, quantified metric.

8. Common failure modes and misuse

What are the most frequent failure modes or user misuses (e.g., over-reliance, accepting outputs without review, weak prompting, using AI-generated content as final work, inaccurate summaries, inappropriate assessment use, privacy-sensitive uploads), and how were these identified (testing, analytics, support tickets, educator feedback, institutional feedback, etc.)?

We currently identify failure modes and misuse mainly through **user feedback**; we have no separate guard against, say, a student using AI to write an assignment. This is a genuine balance between usefulness and safety — we don't want to block a student who legitimately needs help — so our stance is to be the most helpful assistant possible.

The **AI grader** mitigates answer-harvesting structurally: it's **one-way** — students can't "talk" to it; it only grades. The **chatbot**, by contrast, is an open box, which is where the usefulness/safety balance lives. We haven't fully solved it and prefer to stay open and work with the industry rather than over-restrict.

9. Built-in guardrails, nudges and user education

What guardrails are built into Chat Everywhere to encourage human review and reduce over-reliance (e.g., edit transparency, reminders, warnings, usage guidance, teacher controls, shared prompt curation, organisation-level settings, account permissions), and why were these chosen?

Chat Everywhere is designed around **simplicity** — reducing mental load and not surfacing options on parts that need human review, rather than piling on warnings.

The key guardrail is in the quiz: the AI grades, but students can **dispute a grade**, and we encourage it. Because their score is on the line, students are incentivised to catch the AI's mistakes — learning that AI shouldn't be blindly trusted. Teachers see the same when students report errors. So the "guardrail" is really a set of **incentives** that teach both sides where AI helps and where it doesn't.

How does the platform help users understand the limits of AI-generated outputs, especially in educational or institutional contexts?

We help users understand AI's limits **through experience, not rules** — it's unrealistic to enumerate every do and don't. Students and teachers are exposed to AI output with their incentives tied to it: a student's chosen output determines their assessment, so they must verify what they include; a teacher must check that AI-generated questions and answers are correct, or face manual rework. This **design-by-incentives** approach makes them learn AI's limits for themselves.

10. Responsible AI use, academic integrity and disclosure

What is Chat Everywhere's position on transparency of AI use, responsible disclosure and academic/workplace integrity? How can the platform support responsible use without creating "compliance theatre"?

This is again **design-by-incentives**. The core mechanism is **conversation visibility** at two levels:

Teacher Portal: a teacher sees a student's conversations within a classroom they created.

Organization: if a student's personal account is added to the organisation, its IT admins and teachers can see all that student's conversations. All conversations stay within our cloud and are never shared externally. Knowing it's visible keeps students from misusing or over-trusting the AI — accountability comes from the **real structure of the environment, not disclosure checkboxes**, which is how it avoids "compliance theatre." Planned: an automated system that flags concerning interactions (e.g., overly intimate or persuasive) to teachers or parents.

Are there recommended disclosure templates, institutional policies, teacher-facing guidelines, or user-facing reminders that help clarify when and how AI use should be declared?

We don't currently offer formal disclosure templates, institutional policies, or guidelines. Accountability has been handled **structurally** — through visibility and incentives — rather than documented disclosure. This is an area for future development.

11. Privacy, data protection and institutional governance

How do you approach privacy and data protection for educational and professional users (e.g., data retention, processing controls, uploaded files, meeting data, student data, enterprise/institutional options, deletion requests, regional requirements), and what do you recommend to schools, universities, training centres or business teams?

We take privacy seriously: we never share or resell user data, and even internal debugging **strips personally identifiable information first**.

This is why some institutions want local deployment — but local models perform far worse than cloud models, undermining the AI's usefulness, so we don't currently recommend fully local deployment. Instead, our **white-label solution** runs in an **enterprise-grade cloud environment**, giving similar protection with cloud speed and intelligence. We recommend a cloud solution done properly: clear data hosting, database protection, and audited authentication.

What changes when the platform is deployed institutionally rather than individually in terms of governance, permissions, monitoring, integrations, procurement, security review and data-processing agreements?

Deployed **institutionally**, the institution holds the highest privilege — access to all its data — and can run its own security review, set its own data-processing agreement, choose hosting, and define its security level; essentially it **tailor-makes the environment**. **Individuals** use what we provide, which is enterprise-grade but without that elevated control and configurability.

12. Individual vs institutional adoption

What changes when adoption is institutional (schools, universities, training centres, companies) rather than individual (requirements, integrations, governance, onboarding, teacher/admin roles), and what onboarding practices have worked best?

Our onboarding is deliberately **self-driven, not top-down**. We start with a handful of users who've tried the public version, run training and Q&A focused on their most relevant feature, then hold sessions through the deployment — **without forcing usage**, since turning it into a KPI distorts incentives. We seed a small group of **enthusiasts** who teach their colleagues, and let adoption spread, supported by FAQs and documentation. The best practice we've found is **enthusiast-led, voluntary diffusion** rather than mandated rollout.

Which adoption barriers are most common (e.g., AI literacy, trust, privacy concerns, cost, language, teacher readiness, internal policy gaps), and how does Chat Everywhere address them?

The two most common barriers are **AI literacy** and **privacy concerns**:

AI literacy: addressed experientially — people try the tool, see how much time it saves, and adopt because they felt the benefit.

Privacy: public platforms and shared accounts expose private chat records. We address this with the **white-label solution** (everyone gets an isolated account in the institution's cloud) and by advising against shared accounts. Account isolation is strong — no user can ever see another's data.

13. Aggregated impact evidence and conditions

Can you share 2–3 non-confidential, aggregated examples of typical impact (productivity, teaching support, learning support, organisational efficiency and/or quality), and under what conditions those results are most likely (user type, task, discipline, language, institution size, onboarding quality, stage of use, etc.)?

(All figures are aggregated platform analytics with their measurement window; several are conservative undercounts because instrumentation post-dates launch.)

1. **Organisational efficiency (Meeting Summarizer):** over **3.2 million minutes (~53,000 hours)** of audio/video transcribed across **~99,000 meetings for ~19,700 users** (Aug 2025–Jun 2026). Most impactful for professionals with frequent or long meetings, where the alternative is manual transcription.

2. **Learning support (AI quizzing):** nearly **21,000 quiz submissions across ~990 teacher-created quizzes** (since Jul 2025), each with instant tailored feedback. Most effective in active classrooms where a teacher builds the quiz rhythm into well-defined subjects.

3. **Access and scale:** over **110,000 registered users** and **6.6 million+ chat messages** (Oct 2023–Jun 2026), plus large supporting usage (126,000+ images, ~39,000 pronunciation analyses). Strongest where onboarding is enthusiast-led and the product stays simple.

Caveats: transcription minutes are instrumented only from Aug 2025 (an undercount); quiz figures are submissions, not distinct students; active-user counts include guests and are upper bounds.

14. Roadmap, differentiation and limits

What differentiates Chat Everywhere from general-purpose AI chatbots and from other education/productivity AI platforms?

Vs. general chatbots: though it started as one, Chat Everywhere is now an **education- and productivity-focused platform**, adding the Pronunciation Practicer, Meeting Summarizer, AI and personalised quizzes, the Classroom, and the Organization Portal — covering most classroom AI use cases.

Vs. other edtech/productivity tools: rather than juggling multiple separate tools with different interfaces, **everything is aggregated in one platform**, and the white-label option puts the whole suite under the institution's own domain.

Which future developments are strategically important, and which limitations should be acknowledged honestly in a scientific case-study paper?

Roadmap: the next step is a new product with a much **simpler, more intuitive UI** that keeps Chat Everywhere's power but removes the barrier of entry — many teachers, especially senior ones, struggle with chatbot-style interfaces, so the aim is a tool so intuitive users barely need to type. (Other directions: a school-level "talent bank" and an automated safety-flagging system.)

Limitations to acknowledge honestly:

The bigger open question — how AI should be incorporated into education ethically and practically — is industry-wide; we want to help define best practices.

AI still makes mistakes; the field needs a shared framework.

Accuracy is mitigated, not eliminated; the platform relies on human review.

Impact evidence is engagement/usage-based, not a controlled study of learning outcomes.

The evidence is single-vendor and self-reported.

Validation is concentrated in Taiwan/Macau and Chinese-language education, so generalisability is unproven.

THANK YOU!