July-September 26
Aryavart Journal of Multidisciplinary Research
(NATIONAL PEER REVIEWED E-RESEARCH JOURNAL)
E-ISSN 3107-6009
Vol - 2 Issue - 3

Working with Artificial Intelligence, Not Only Using It

(A Review of Trust, Shared Work and Human Control in human–AI Teams)

Shubham Rajak
Assistant Professor, School of Computer Science
Aryavart University, Sehore (M.P.)

ABSTRACT

Many offices, colleges and research labs now use artificial intelligence every day. People ask a system to write a draft, check a plan, or help with a decision. This is often called human–AI collaboration. The phrase sounds simple, but the work is not. A person still has to judge the answer, correct mistakes, and take responsibility if the output is wrong. This paper reviews recent studies on how people and AI systems share work. It looks at trust, control, team roles and the habit of saying that a "human is in the loop." The main point is modest. Good collaboration does not come only from a stronger model. It comes from a clear split of tasks, honest signs of uncertainty, and a human who can really check the result. Studies also show that people may trust smooth language too quickly, and that a label such as "human supervised" does not always raise confidence. The paper ends with simple questions that later research can test, and with advice for teaching and workplace design.

KEYWORDS

Human–AI collaboration, trust, human control, AI teammate, overreliance, shared work, human-in-the-loop, team performance, knowledge work, artificial intelligence.

1. INTRODUCTION

A few years ago, most software waited for a clear command. The user clicked, typed, or pressed a key, and the program followed a fixed path. Today's AI tools behave in a different way. They write paragraphs, suggest next steps, and sometimes act through other programs. Teachers, doctors, managers and students now sit with a machine that talks back. The question is no longer only "Can the tool do this task?" The harder question is "How should a person and a machine share the task?"

That question matters because the tools are uneven. On some jobs they are fast and useful. On other jobs they sound sure and still get the facts wrong. A person who accepts every sentence will copy errors. A person who rejects every sentence will waste the help. The useful skill sits in the middle. The person must know when to accept, when to edit, and when to start again.

Workplaces often use friendly words for this mix. They speak of copilots, assistants and digital teammates. Words can hide the real problem. If the system cannot keep track of the same file, cannot say "I am not sure," and cannot show who changed what, it is not a teammate. It is a

talkative tool. People then guess how much to trust it. Guessing is a poor method when the output may reach a student, a patient, a client or a public record.

This paper is a review, not a new laboratory test. It brings together recent findings on trust, team process and human control, and it writes them in plain language. The aim is to help teachers, students and working professionals see Human–AI collaboration as a research topic with clear parts: who does which step, how trust is formed, and how responsibility is kept after the work is done.

2. RESEARCH SCOPE

2.1 Review of Literature

Older writing on computers and people focused on ease of use. If the screen was clear and the steps were few, the design was called good. Later work on group software asked how people coordinate with one another. Human–AI collaboration sits on both lines of work and adds a new difficulty. The other "member" of the team can invent text, change plan in the middle of a task, and speak in a confident tone even when it has no stable picture of the user's goal.

A large part of the recent literature is about trust. Trust here does not mean liking the tool. It means a person's readiness to depend on the output. Too little trust and the tool is ignored. Too much trust and mistakes pass through. Several studies find that people lean on answers that read well, even when the content is weak. Wrong or made-up details, often called hallucinations, make this worse. They do not only spoil one answer. They also lower the user's later willingness to work with the system.

Another part of the literature looks at real team work, not only ratings on a form. In large tests where people made public-facing content with either a human partner or an AI partner, the mixed teams often produced more text and better written copy. Human pairs were stronger on some visual work. Talk with the AI was more about the task and less about social chat. People also handed over more of the writing and made fewer small edits. This helps speed. It can also make many outputs look alike.

A third line of work tests a popular claim: if a human watches the AI, people will trust the result more. The claim is common in policy notes and product pages. Direct tests are less kind. In one set of studies, the same report was shown with different source labels, such as human, AI, or human-supervised AI. The source label did little to change later action. In some cases, the "human supervised" label was the least trusted. A name on the page is not the same as real checking.

A fourth line looks at how teams think together. People work better when they know who knows what. The same idea now appears in human–AI writing tasks. When the person and the system bring different useful knowledge, the team does better and members feel more satisfied. A clear and calm team mood helps. Leaders who manage both human staff and AI agents still tend to prefer people, and they lose trust in agents quickly when skill or honesty looks weak.

Aryavart Journal of Multidisciplinary Research
(NATIONAL PEER REVIEWED E-RESEARCH JOURNAL)
E-ISSN 3107-6009
Vol - 2
Issue - 3

2.2 Research Gap

The field has grown fast, but some holes are easy to see. Many papers still treat the AI as a box that gives one answer. Newer systems plan, call other tools, and work across many steps. Old trust questions do not fully cover that behaviour. Second, a lot of scoring looks only at the final text or the final mark. It does not record who decided, when control moved, or whether a person could have caught the error. Third, "human in the loop" is used as a comfort phrase. Few papers define what the human must actually do. Fourth, offices buy tools faster than they change work rules. Technical papers rarely study that delay. This review tries to place those gaps next to each other so later work can pick one and test it.

2.3 Hypotheses and Research Questions

The review uses four working claims. Later studies can support or reject them.

H1. After a system reaches a fair level of accuracy, the quality of joint work depends more on well-matched trust than on small extra gains in accuracy.

H2. A human-supervision label, with no sign of real review, can lower trust instead of raising it.

H3. Human–AI teams do better on some well-structured writing and drafting tasks, and worse on tasks that need human taste, talk, or careful looking. The gain is uneven.

H4. A shared picture of roles and knowledge explains much of the benefit of mixing human skill with machine skill. Without that picture, extra output may only add noise.

Three questions follow. First, which simple signals help a user see the difference between fluent text and sound content? Second, how should the right to act be shared when a system can do more than give advice? Third, what should a fair test of "team quality" include, beyond the last answer on the screen?

2.4 Objectives

1. To gather recent evidence on Human–AI collaboration as a problem of trust and shared work, not only as a problem of model scores.
2. To name the main paths through which joint work goes well or goes badly: trust, handing over tasks, talk style, shared roles and team mix.
3. To show why popular fixes, such as long explanations or a human label, are not enough on their own.
4. To set a short list of tests and design steps that students and working teams can use.

3. METHODOLOGY

This paper is a review article. It does not report a new survey of local staff or a new lab experiment. The method is a careful reading of recent work, mostly from 2023 to 2026, in human–computer interaction, information systems and nearby fields.

Aryavart Journal of Multidisciplinary Research
(NATIONAL PEER REVIEWED E-RESEARCH JOURNAL)
E-ISSN 3107-6009

Sources include conference papers, journal articles and a small number of widely discussed preprints. Search terms included human–AI collaboration, AI teammate, trust, and human-in-the-loop, shared mental models and mixed teams. A paper was kept if it said something clear about people and AI working on a task together. Papers that only reported a model score, with no human role, were left out.

For each useful study the review asked four plain questions. What process does the author claim is at work? How was the claim tested? What was measured? What would a designer or teacher do with the result? The answers were then grouped by theme.

This method has limits. It is not a full count of every paper in a database. It gives more space to recent and visible studies. It cannot produce one number for "how much AI helps," because writers do not measure help in the same way. Those limits fit a review whose job is to clear the ground, not to close the debate.

4. DISCUSSION

4.1 A Tool Is Not the Same as a Teammate

It helps to keep a simple difference in mind. A tool is picked up, used, and put down. A teammate shares a goal, may start an action, and should notice when the plan is going wrong. Many products use teammate language while they still act like tools. They wait for a prompt. They explain after the fact. They do not keep a stable picture of the job.

The difference is practical. If two people edit one document, they can see the same page. If an AI system forgets the file, drops the last instruction, or treats each request as a new world, the human cannot build a shared plan. Smooth talk does not fix that. Shared work needs a shared place to work, a memory of what was already done, and a rule for who may change the result.

4.2 Trust Needs a Middle Path

Trust is the most common word in this literature, and it is also the most mixed. People need enough trust to try the tool. They also need enough doubt to check it. The useful state is not blind faith. It is a match between trust and real skill.

Wrong details hurt that match. When a system invents a source, a number, or a step, the user may still feel that the text "sounds right." Later, when the error is found, trust can fall hard. Some studies also find that people forgive human partners more easily than machine partners when the record is mixed. If that is true, hiding doubt in order to look strong is a poor design choice. A short note such as "this part is uncertain" may do more good than a long, confident paragraph.

4.3 A Human Label Is Not the Same as Human Checking

Many guides say that a human should stay in the loop. The advice is sensible as a moral rule. As a design rule it is too thin. People can see a badge that says "reviewed by a person" and still not know what was reviewed. In tests, that badge has not always helped. Sometimes it has hurt. Readers may think the system was weak and needed rescue. Or they may think the human only signed a page.

Useful checking leaves a mark. A person who accepts one claim, rejects another, and writes a remaining doubt gives the next reader a way to judge the work. Logs, source links and short uncertainty notes are not extra decoration. They are how the public claim of human control becomes something a later person can see.

4.4 Speed Can Rise While Variety Falls

When people work with AI on drafting tasks, they often finish more pieces and spend less time on small wording changes. That is a real gain. It is not the whole story. If many users ask similar prompts, the outputs can crowd around one style and one set of ideas. A team that wants range, not only speed, has to add a later step: compare drafts, keep odd but useful lines, and refuse the first fluent answer.

The same studies also show that the gain is not even across all parts of a job. Writing may improve while image work, taste, or careful human talk does not. A school or office that tells everyone to "use AI for all work" will lift some tasks and weaken others. A better habit is to map the task first. Where the tool is strong, hand over the draft and check it. Where the tool is weak, the person should do the first work and use the system only to question or search.

aryavart
Journal of Multidisciplinary Research

4.5 Teams Need a Clear Split of Knowledge

People already know, from ordinary group work, that confusion about roles wastes time. The same holds when one member is a machine. If nobody says what the human owns and what the system owns, both sides may repeat the same step or skip a needed check. Studies of joint writing support this. When the person and the system bring different useful knowledge, and when that split is clear, the team does better.

Leaders add another layer. A manager who treats an AI agent like a junior staff member, but gives it no clear duty and no clear limit, will not get a stable team. Trust in the agent falls fast after a weak or dishonest-looking act. Role design is therefore part of the computer problem. Permission, logs and review steps must sit in the software, not only in a memo.

Aryavart Journal of Multidisciplinary Research
(NATIONAL PEER REVIEWED E-RESEARCH JOURNAL)
E-ISSN 3107-6009

4.6 One Default Voice Will Not Fit Every User

Some recent experiments change the "personality" of the AI partner and pair it with different kinds of people. The mix matters. A style that helps one user can slow another. This does not mean the system should pretend to be a friend. It means the way the system talks—how much it pushes, how many options it gives, how much it disagrees—changes how the human edits.

A single average score can hide that fact. A tool that looks fine "for most users" may be a poor match for a careful checker or for a fast planner. Allowing a person to choose a stricter or lighter style is a small design step with a large daily effect.

4.7 What Teachers and Designers Can Do Now

A few steps do not need a new model. Show doubt early, not after a long explanation. Keep the person and the system on the same page or file. Make hand-over visible, so a third person can see who did what. At the start of a project, write two short lists: what the human must judge, and what the system may draft. For class work, students can measure how often they accept a wrong fluent answer, or how an uncertainty note changes their checking. Those exercises need no private computer cluster, and they teach the real skill.

Tests of new agents should also change. A score that only asks "Did the final answer look right?" will keep rewarding systems that do well in a short demo. Joint work lasts longer than a demo. A fairer test asks whether the system can recover after a mistake, whether it can wait, and whether a person can still see the trail of choices.

5. CONCLUSION

Human–AI collaboration is a strong research topic because it sits between machines, people and daily duty. The recent record supports a plain conclusion. Help is real on some tasks and overstated on others. Trust is needed, but it is easy to set at the wrong level. A human label helps only when the human actually checks. Teams improve when roles are clear and when each side brings something the other side lacks. A system that cannot keep state, admit doubt, or accept a limit is not a teammate, however smooth its sentences may be.

Later papers can test the four claims in this review. Does better matching of trust beat a small rise in accuracy? Does a visible review trail fix the weak effect of a supervision badge? Can a simple task map predict which school or office jobs should be handed over? Does a shared role sheet help more than a larger model on the same job? Each test is small enough for a student project and useful enough for a workplace.

The last line can stay short. Do not ask only how strong the model is. Ask who may act, who must check, and how both sides will know that the plan has failed. If those points are settled in the work process, collaboration is real. If they are not, the word "teammate" is only a polite name for a tool.

REFERENCES

1. Dell'Acqua, F., McFowland, E., Mollick, E., and others. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. Harvard Business School Working Paper.
2. Dell'Acqua, F., and others. (2026). Collaborating with AI agents: A field experiment on teamwork, productivity, and performance. arXiv preprint.
3. Endsley, M. R. (2017). From here to autonomy: Lessons learned from human–automation research. Human Factors, 59(1), 5–27.
4. European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act).
5. Gondocs, D., and others. (2025). Uncovering the dynamics of human–AI hybrid performance. International Journal of Human-Computer Studies.
6. Ju, H., & Aral, S. (2026). Personality pairing improves human–AI collaboration. arXiv preprint.
7. Kaur Nagpal, A. (2026). Human in the loop, or perceived oversight? The psychological inference that drives AI credibility. Psychology & Marketing.
8. Klein, G., Woods, D. D., Bradshaw, J. M., Hoffman, R. R., &Feltovich, P. J. (2004). Ten challenges for making automation a "team player" in joint human-agent activity. IEEE Intelligent Systems, 19(6), 91–95.
9. Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. Human Factors, 46(1), 50–80.
10. Li, and others. (2026). The role of shared mental models in driving knowledge complementarity: Enhancing human–AI team effectiveness. International Journal of Human–Computer Interaction, 42(16).
11. National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0).
12. Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. Human Factors, 39(2), 230–253.
13. Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe and trustworthy. International Journal of Human–Computer Interaction, 36(6), 495–504.
14. Song, D. (2026, August). On the promise and perils of agentic AI [Interview]. Berkeley News.

Aryavart Journal of Multidisciplinary Research
(NATIONAL PEER REVIEWED E-RESEARCH JOURNAL)
E-ISSN 3107-6009

15. Technology in Society. (2026). The trust crisis in artificial intelligence: AI hallucinations and human–AI collaboration, 86, Article 103286.
16. Technology in Society. (2026). Towards an integrative model of organizational human–AI collaboration: A semi-systematic review, 84, Article 103064.
17. Vaccaro, M., Almaatouq, A., & Malone, T. (2024). When combinations of humans and AI are useful: A systematic review and meta-analysis. Nature Human Behaviour.
18. CHI Conference on Human Factors in Computing Systems. (2026). "Should I rely on you or the AI?" Leaders' trust and perceptions in mixed human–AI teams. Proceedings of CHI 2026.
19. Review article. (2026). AI as teammate or tool? A review of human–AI interaction in decision support. arXiv preprint.
20. Computers in Human Behavior: Artificial Humans. (2026). Trusting in the competence of humans and artificially intelligent agents varying in generosity, 7, Article 100265.

