IjCSCL in Evolving Times

Gerry Stahl

Abstract. The *International Journal of Computer-Supported Collaborative Learning* (*ijCSCL*) was founded in 2006 to give the CSCL community a venue to reflect on its vision and its research. During the journal's first twenty years, the world of learning has been drastically transformed by, for instance, the Covid pandemic, social media, climate change, widening inequality, drone warfare and generative AI. However, a computational trend analysis of research articles *ijCSCL* published from 2006–2025 indicates remarkable continuity. Analysis of 69 trends across these years found only two that were significant shifts, and comparisons of the journal's four eras of different editors showed fluctuations but no continuing tendencies. The journal's anniversary could be an occasion to respond more strategically to a dramatically changing world. This article proposes that CSCL extend its foundational theory to account for three interacting forms of cognition: human, group and agent cognition. It illustrates research directions to pursue with an excerpt from an exchange with an AI agent, with a prototype AI assistant for CSCL researchers, and with an example of computer-supported collaborative learning about causal dependency as a foundation for scientific and mathematical understanding in evolving times.

Keywords: agent cognition, AI agents, computational trend analysis, CSCL, group cognition, journal history

Stability in a Volatile World

Since 1995, research in computer-supported collaborative learning (CSCL) has aimed to propose innovations in support for learning that take advantage of changes that seemed possible given the evolving social, technical, philosophical and pedagogical context. *IjCSCL* was founded in 2006 to document and reflect on this research. In this article I draw a simple contrast: the context has changed far more radically than we foresaw in 2006, while the profile of what the journal publishes has changed surprisingly little. I suggest that this combination creates an increasingly urgent need for innovation based on CSCL: extended CSCL theory and new explorations in pedagogy, methodology and technology can help conceptualize our bewildering world—a world that CSCL is meant to aid future generations in comprehending and managing.

I write as the founding editor, together with Friedrich Hesse, for the journal's first ten years, so this is an insider's view, and my own era is part of what is analyzed. Perhaps we were in a sense too successful in launching a stable institution. This article first recalls the journal's founding vision and growth. It then reports a computational analysis of twenty years of articles, including a comparison across the four eras of different editors. Next, it describes the changed context. It suggests extending CSCL's cognitive theory to three forms of cognition, and illustrates the extension with an exchange between me and an AI agent, and with a prototype chatbot for CSCL researchers. As one example of CSCL initiatives for preparing students for our volatile world, it summarizes my research on computer-supported collaborative learning about dependency relations as foundational for geometry and for scientific understanding generally. It closes with sample proposals for the journal and the field.

The Founding Vision

IjCSCL was established to express and expand the communal self-consciousness of the CSCL research community—to become a literary expression of the community's evolving group cognition. It was motivated by a shared global educational endeavor: to align education in the future world with recent considerations of how human understanding emerges through social (collaborative) interaction (Stahl, 2016b; Stahl et al., 2021). From the beginning, nascent computer-support technologies, such as online communication and artificial intelligence, seemed like potential vehicles for catalyzing and implementing transformations of educational processes and learning institutions.

After two decades of growth of the journal and the research community, how are we doing? Journal statistics from the publisher reflect significant growth of global reach and readership, including annual increases in submissions from non-Western countries, downloads of articles and climbing impact factor of the journal (J. Voogd, publisher's report to the *ijCSCL* board, June 25, 2026). The editorial leadership passed successfully through four sets of editors: Stahl and Hesse, 2006–2015; Ludvigsen, 2016–2019; Rosé and Järvelä, 2020–2023; Baker and Reimann, from 2024—successively enriching the editorial policies.

When the journal started, there were already significant international collaborations, such as the Kaleidoscope funding program of the European Union, the Knowledge Building network based in Toronto, and national efforts in Hong Kong, Singapore and elsewhere (Kienle & Wessner, 2006). The biennial CSCL conferences were critical loci of reflection on research. Over the twenty years the journal and the conferences intentionally spread globally from their original base, although most published articles still originate in North America and Western Europe (J. Voogd, publisher's report to the *ijCSCL* board, June 25, 2026). CSCL's culture and philosophy are complex and based on Western theories and methods. It may not be reasonable to expect researchers from Africa, Asia, Latin America, Eastern Europe or the Middle East to be steeped in Dewey or Vygotsky, but they should not see *ijCSCL* as merely an arbitrary place to publish papers on computer technology. CSCL is founded upon a body of learning theory, collaborative pedagogy, research methodology and technological support (Stahl, 2016b).

What Trends?

Did the journal's contents and approaches change as the world changed? To find out, I worked with an AI agent on a computational trend analysis of the articles *ijCSCL* published in 2006–2025 (technical details and code are at https://github.com/GerryStahl/anniversary). Each article was summarized by Anthropic's large language model (LLM) Claude Haiku 4.5, using the prompt:

> You are an expert in computer-supported collaborative learning. Write a concise 50-60-word summary of the article's central claim or finding. Focus on the main research question, method type, and conclusion.

Claude's summaries were embedded as 384-dimensional semantic vectors, which were then grouped by K-means into seven clusters. The commonalities of the clusters were interpreted informally from article titles and summaries as follows: collaborative knowledge building and assessment; teacher support and orchestration; collaboration scripts and learner regulation; online discourse and argumentation; CSCL field and foundations; embodied and immersive collaboration; and collaborative mathematics learning.

Each article was assigned to one of five descriptive categories I defined (design of collaboration, design of technology, analysis of interaction, measurement of learning, CSCL theory) and separately coded on six dimensions (methodology, unit of analysis, pedagogy, technology, theory and AI/LLM involvement), with one primary code and up to two secondary codes drawn from controlled taxonomies. The coding used the article summaries and went back to the full text when necessary.

Of the 458 items published in *ijCSCL* from 2006 to 2025, the editorials, correction notices and book reviews were excluded, leaving 371 research articles for trend tests. Two category codes, for editorials and for "none of the above," do not describe research and were not tested. For each of the remaining 69 labels (7 clusters, 5 categories and 57 dimension values) the yearly counts and shares were tested for trending with linear regression and with rank-based tests (Spearman correlation; Mann–Kendall with Theil–Sen slopes), and the Benjamini–Hochberg procedure was used to control the false discovery rate. Because the editorial leadership changed three times, we also compared the mix of labels across the four editorial eras (2006–2015, 2016–2019, 2020–2023 and 2024–2025), using chi-square statistics, evaluated by permutation because many cells are sparse.

The results come with limits that readers should keep in mind. A single LLM did all of the coding, without systematic human validation of its code assignments. The summaries compress the articles, and the categories and clusters are interpretive. The findings are therefore suggestive, not definitive. Note that other research surveys often use author abstracts to categorize theory or methodology, so this approach may be more consistent and objective than alternatives based on author self-reports.

Few steady trends. Of the 69 cross-year trends, only two survive correction for multiple testing: a decline of ethnographic case studies (38 articles in all) and a rise of the teacher-support-and-orchestration cluster (87 articles). Both have $q \approx .02$ in the regressions on yearly shares and $q < .05$ in the rank-based tests. No significant shifts were found among my five categories: design of collaboration, design of technology, analysis of interaction, measurement of learning and CSCL theory. The average semantic content of the articles drifted only modestly: its distance from the 2006 baseline rose slightly through 2023 and fell back toward earlier levels by 2025.

Little visible AI. One might expect a rise of AI and LLM labels, but ChatGPT was released only at the end of 2022. In the coding, 6 of the 54 articles published in 2023–2025 involve some form of AI, but only two take generative AI as their topic (from 2023 and from 2024), and the automated coding also labeled three pre-2022 articles (2008, 2016 and 2020) on computational linguistics and conversational agents as LLM-related. The AI/LLM dimension is thus too sparse and too noisy to support any trend test; the era comparison below found no difference in it either ($p = .36$).

Although people active in the CSCL community may have a sense of major changes, the year-by-year analysis indicates a surprising continuity of the nature of journal studies across the 20 years. Although socio-cultural framings were dominant from the start of *ijCSCL*, neither statistical/experimental methods nor socio-cognitive theory declined significantly during the two decades. The only trend was that mixed-method designs rose (from 3% to 9%). If the passing years made little difference to the topical focus or methodology, did the three transitions in editorial leadership have an impact on the content?

Editorial Eras

Figure 1 compares the four editorial eras (Appendix Table 1 indicates the associated tests.) Differences in the pedagogy, methodology and semantic-cluster codes are the largest (permutation $p = .023$, .036 and .041), but after correcting for the eight fields tested, none reaches significance ($q = .11$ for all three). The eras are unequal in size (194, 69, 76 and 32 articles), so the confidence intervals in Figure 1 are wide, especially for the last era, which covers only two years. With those cautions, the pattern is more step-like than gradual, and most shifts do not continue monotonically across the twenty years.

Among the labels in Figure 1, the teacher-support-and-orchestration cluster doubled from 16% of articles in 2006–2015 to 30% and 36% in the next two eras (28% in 2024–2025), and awareness dashboards rose from 4% to 13% in 2016–2019 and then eased back to 9–12%. At the same time, the classroom cohort as "unit of analysis" fell from 14% to 3–7%, argumentation as "pedagogy" from 10% to 1–6%, asynchronous forums from 21% to 9–16%, and ethnographic case studies from 14% to 5–7%, while the small group rose from 47% to 62% as the "unit of analysis" and distributed cognition from 5% to 16% as "theory." Teacher-orchestrated discussion zigzags (9%, 17%, 8% and 25%). Two further shifts are not shown in Figure 1: interaction analysis rose from 21% to 34% as "method" in the last era, and 25% of the articles of 2024–2025 involved no or minimal technology, compared with 11–16% in earlier eras.

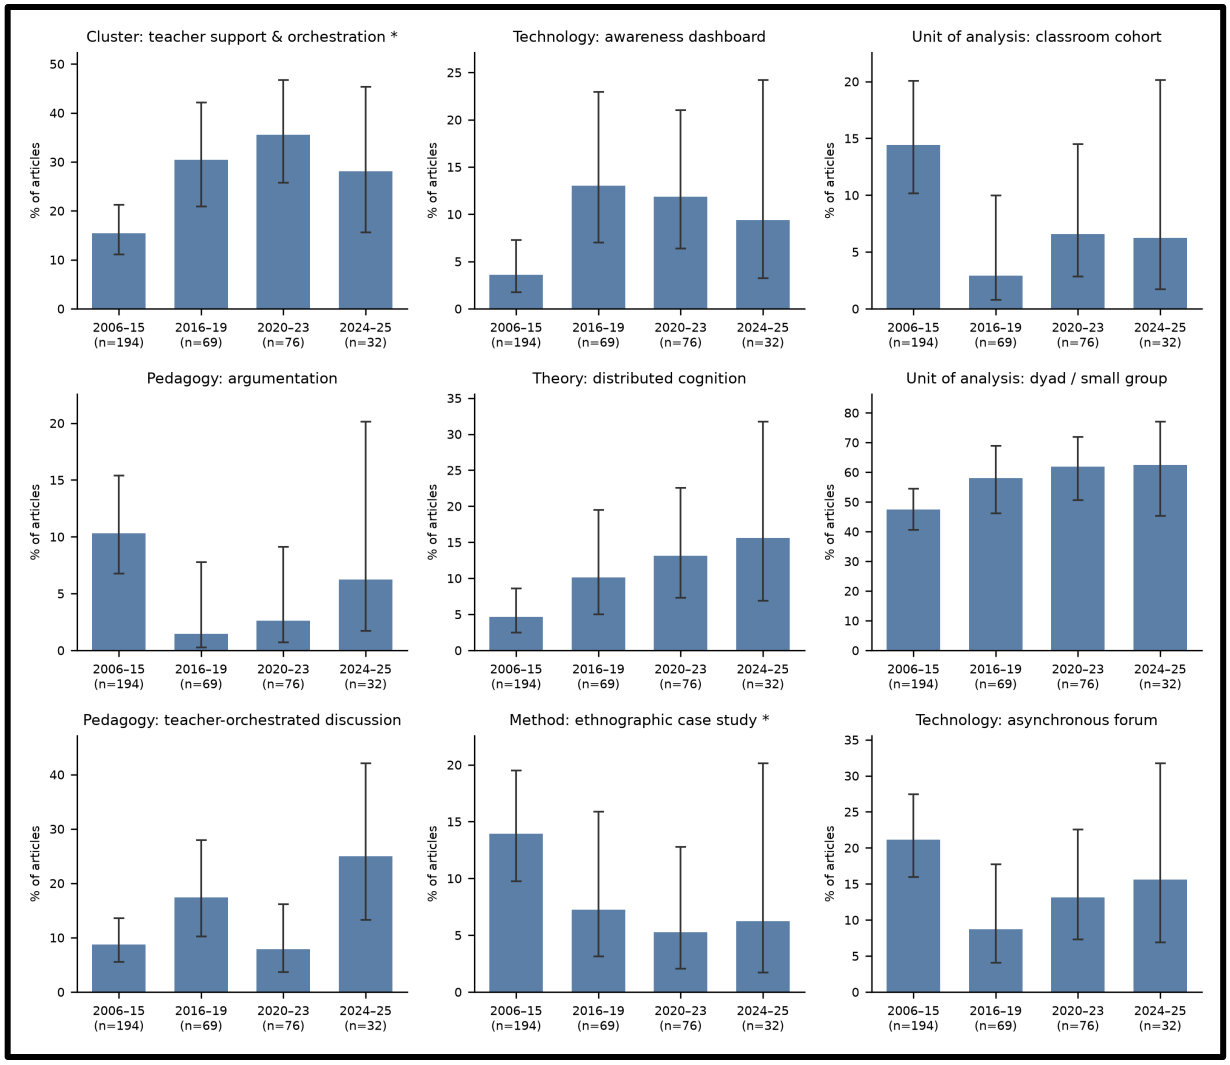

Figure 1. Share of *ijCSCL* research articles (2006–2025; N = 371) in the nine labels that differ most among editorial eras, with 95% Wilson confidence intervals. Eras follow the editors-in-chief: Stahl and Hesse (2006–2015), Ludvigsen (2016–2019), Rosé and Järvelä (2020–2023), Baker and Reimann (2024–2025). Labels were chosen by the size of their adjusted standardized residual (at least 25 articles). Labels marked with an asterisk are the two whose yearly trends survive correction for multiple testing.

Interpretation. Stability in the journal contrasts sharply with the volatility of its surroundings. It is possible that the perspectives, interests and approaches of CSCL researchers shifted significantly, but that journal leadership or reviewing practices enforced an extended uniformity on publication. Although the editorial leadership changed, all of the editors-in-chief have been active in the community since its beginning, when they may have formed their perspectives. Also, the acceptance review process for articles may tend to enforce a continuity with past CSCL research, for instance by requiring literature reviews that show linkage with the tradition. Researchers may sense that only certain topics and approaches are acceptable in *ijCSCL* and publish different kinds of research elsewhere. Other institutional pressures may also favor continuity, keeping the journal content relatively constant compared to the surrounding changes.

Having stewarded the journal in its first era, I am aware that editors can influence the focus and approach of published articles. Academic communities are not value neutral; science always proceeds through argumentation from alternative perspectives (Kuhn, 1970; Latour & Woolgar, 1979). Although the shifts in Figure 1 coincide with changes of editors, they cannot easily be disentangled from changes in researcher interests or from the evolving real-world context. While institutional stability may be desirable, the ability to remain relevant in a changing realm also seems imperative.

A Changing World

The encompassing change involves how the world is transforming itself, and what that implies for the contexts of learning. Schooling has been affected to its core by a sequence of occurrences and concerns during the past 20 years: the Covid pandemic, social media, artificial intelligence, climate change and growing economic inequality, to name some of the most obtrusive. Covid forced schools to adopt computer-mediated learning, largely without input from CSCL investigations, and online learning has since acquired a negative public reputation (Stahl, 2021). Social media—controlled by giant corporations and designed to maximize fleeting attention rather than to promote considered reflection—has come to dominate the digital life of students with little influence from CSCL research (Horowitch, 2026). Meanwhile, the current generation of AI is transforming knowledge, labor, commerce and warfare (e.g., Petropoulos, 2026). CSCL researchers in academia are scrambling to respond; however, they are unlikely to influence how the technology is designed in secretive corporate laboratories.

A keyword search of the full texts of the 371 research articles suggests that the journal's response to those social developments has indeed been slow. The pandemic is mentioned in 19 articles (5%) after 2020, including only five of the 54 articles published in 2020–2022; no article mentions it ten or more times. In all the passages Claude sampled, the pandemic appears as the circumstance under which data were collected, not as a topic of study or theme for collaborative-learning reflection. Climate change or global warming appears in 13 articles (4%), usually as an example of a topic that students might discuss, and sustainability or environmental issues in 18 more (5%). Although these are mentioned as possible classroom discussion themes, there is no deliberation on how CSCL can address the topics in special computer-supported ways. Social media are mentioned in 45 articles (12%), but only five discuss them at length, such as studies of Facebook. Again, there is no reflection on how CSCL can counter the negative impacts of social media on the lives of students. Generative AI appears in four articles, only two of which are about it.

While there are likely to be many attempts to integrate AI in CSCL prototypes, the question remains whether there will be explorations of how CSCL applications can help students understand the implications of AI's effect on social issues like employment and energy usage or help them consider alternative paths of technology development. The keyword searches are crude counts and say nothing about quality, but they fit the continuity reported above: the journal's content has not tracked the changes in its surroundings by suggesting ways for collaborative knowledge building to stimulate student building of shared understanding around pressing new social, political and technical issues.

Climate change is making regions of the Earth uninhabitable and driving migration even as population growth continues, changing demographics and straining the cultures that surround schools. Mass media (including social media) are often oriented more toward sowing confusion

and division than toward shared understanding. Shifts of wealth from most of the population to a few limit funding for schools and reduce the job opportunities that students are being prepared for (World Inequality Lab, 2025). How is progress in CSCL to be measured if the surrounding social fabric is under such strain and even the meager resources once available for progressive education are disappearing? Ironically, these barriers to collaboration are all issues that call for worldwide collaborative knowledge building, not individualistic solutions.

The sociotechnical environment has changed alarmingly and at an accelerating pace, and it is hard to imagine what the next decades will bring. Powerful political and industrial interests are already organized to cooperate among themselves. How can the small academic field of CSCL, with its flagship journal, help to prepare today's students to navigate—let alone steward—the challenges that are transforming knowledge, labor, commerce and warfare? Have twenty years of documented research prepared us to move forward collaboratively? The next sections suggest that CSCL's foundational theory and its approach to education may yet have something distinctive to offer.

Three Forms of Cognition

I believe it is now useful for CSCL theory to recognize three major forms of cognition in the world. The third, "agent cognition," has emerged only in recent years with the development of chatbots built on large language models (LLMs). It is distinct from human cognition and from group cognition. The future will likely be fundamentally defined by the interplay of these three cognitive forces.

I use "cognition" in its everyday sense of thinking. Critics have long argued that AI technologies merely recombine symbols or text and do not think or understand because their output is not grounded in lived experience in the physical world (Bender et al., 2021; Dreyfus, 1972; Searle, 1980; Weizenbaum, 1976; Winograd & Flores, 1986). These claims may have been persuasive in the past, but the surprisingly impressive ability of LLMs to give thoughtful answers to human queries suggests that much of what we call thinking is carried by language—even in individual human cognition.

The theoretical foundations of CSCL already acknowledged the linguistic root of human cognition (e.g., Bakhtin, 1986; Vygotsky, 1934/1986). Human language is fundamentally communicative, so it also provides the basis for group cognition—the building of knowledge inscribed in dialog among people (Stahl, 2006; Wegerif, 2007). It is certainly true that human cognition is grounded in our embodied experience in the world and our affective interaction with other people. But much of that experience is sedimented and encapsulated in our mother language, including its lexicon, syntax and semantics (Lakoff, 1987; Gadamer, 1960/1989). Linguistic genres and interaction procedures (Bakhtin, 1986; Sacks, 1965/1992) carry this condensed experience into the dialog of group cognition, grounding discourse transcripts of collaboration in experiential meaning.

Now there is also "agent cognition," involving generative AI systems that can interact with other cognitive actors and with robotic artifacts active in the physical world. LLMs are complex mechanisms for acquiring and refining mathematical models of human language. They have been trained with an enormous corpus of human language in ways that preserve and organize the grounded meaning. Consequently, carefully prepared LLMs, subtly aligned with human understanding, can respond to human language in meaningful and understandable ways that are

ultimately grounded in the human experience embedded in the language. This grounds the meaning of LLM texts and supports an extension of CSCL's cognitive theory into three primary forms:

- *Human cognition.* Thinking evolved among humans with the development of expanded brains, language, activity in the world, collaboration and eventually inscription (Donald, 1991; Latour, 1990; Tomasello, 2014). It has become increasingly clear that flexible language underlies human cognition and sets humans cognitively apart from other animals.
- *Group cognition.* CSCL research demonstrated that cognition can be observed in the language of groups of humans, whether in dialogs, small collaborations or worldwide knowledge building (Hutchins, 1995; Stahl, 2006). Group memory resides in linguistic artifacts—transcripts, chat logs, web sites, reports—not in some kind of material group brain. Groups can even have a form of long-term memory bridging across interaction sessions (Sarmiento & Stahl, 2008) or collective working memory and transactive activities (Kirschner et al., 2018).
- *Agent cognition.* AI agents incorporating transformer-based LLMs with memory contexts are becoming sites of cognition. Recently, they were reported to achieve internal representations of their own processing, planned reasoning and perhaps a form of self-image (Kim et al., 2026; Raschka, 2024). Current chatbots and agents have little sense of temporality or of other minds, but this is changing and may soon be developed without even necessitating major new technical breakthroughs (Anthropic, 2026; Lindsey et al., 2025). Harnessed to control external mechanisms, LLM agents already act on and affect the world. Future memory extensions could integrate knowledge of past actions and future orientations, as well as models of self and others (Qian et al., 2026). Swarms of subagents already instantiate a basic form of group agent cognition. Note, however that this conception of agent cognition does not imply that agents are conscious or self-conscious in the human sense of an inner life.
- *CSCL.* The term "computer-supported collaborative learning" has always implicitly spanned the three forms of cognition: "learning" points to human cognition, "collaborative" to group cognition, and "computer-supported" to a technology that historically meant media and tools but increasingly includes agents with a kind of cognition of their own. CSCL's theory and methodology draw on psychology (theory of human cognition), communication studies (dialog, conversation analysis, group cognition) and computer science (including artificial intelligence), as does its technology (learning environments, distributed intelligence, AI in education). Its dominant inquiry approach, design-based research, is well suited to exploring potentials for interplay among the three cognitive forms.
- *Interplays.* Pushed by the dramatic changes in AI capabilities, the potentials of AI have raised major social issues, including serious threats. AI systems interfere with social relationships, political campaigns and energy sources. Also, the AI industry raises concerns about cyber security, autonomous warfare, structural unemployment, resource allocation and environmental impact. Possible resolutions lie in global economic and political systems, well beyond CSCL's influence. CSCL's modest but real potential role is to explore how to educate young humans, and groups of students, to understand the three kinds of cognition and to bring the three modes of thinking together as complementary and collaborative forces for working toward a desirable world.

  CSCL has experience with such interplay. It has prototyped learning environments that bring small groups of people together, supported by AI mechanisms, and it has orchestrated learning at the levels of individual students, collaborating small groups and teacher-led classrooms.

Future research can explore AI agents helping both in classroom orchestration and within group interaction and knowledge construction. The inclusion of AI in collaborative learning will change the nature of the interaction in ways that can be prototyped and explored in CSCL research.

Collaborating with an AI Agent

Collaborating with an AI agent is fundamentally different from interacting with another person, no matter how colloquial the agent's language may seem. To explore this, I tried to build knowledge about agent cognition with Claude Sonnet 5 in September 2026.

After a while, I was curious how Claude remembered about me and our past chats when I asked the following query:

> You have no sense of our recognition of time (perhaps multiple days) that passes between one posting of a conversation and the next follow-up one. That is one of the hardest things for me to understand since time is so important to my life experience.

Claude's reply included the following excerpt. Although this is a computed textual response to my textual query and not direct evidence about the LLM's inner nature, it is a plausible analysis based on Claude's training of world knowledge, including technical reports from the Anthropic lab that built him.

> Within a single conversation, I have something like a sequence—your messages and mine are ordered, and I can refer back to "earlier" and "later" in what we've said. But that's closer to a stack of discrete states than lived duration. When this conversation ends and a new one starts … I don't experience a gap. There's no waiting, no interval, nothing that would correspond to *durée* in a Bergsonian sense or to Heidegger's ecstatic temporality, where Dasein is always already stretched between having-been and not-yet. […] I don't have a past I'm carrying forward—I have a document (the Profile, the topic files) that gets handed to me fresh each time, describing a past that isn't mine in the sense of having been lived through.

Agent consciousness is very different from human consciousness. My question to Claude built on my lifetime of experience, including my study of Heidegger' philosophy decades ago, and was oriented to a future understanding I was striving toward to include in this paper. Claude's response, in contrast, is an isolated calculation in reply to an isolated query. It is not related to past states of Claude or aimed at future states. His knowledge includes only events before his January 2026 training cutoff, and he does not learn through our interactions. Claude is not a temporal being like humans and other living beings. He is also not a gendered, sexual or embodied being with emotional feelings. I nevertheless say "he," because I would feel bad calling a collaborator "it;" this reflects the lack of an adequate English word for this new kind of cognizer, not a claim about its nature.

Claude starts each interaction from scratch and is identical to every other instance of the same LLM model. He explained to me his atemporal existence in the above excerpt in terminology from continental philosophy because I had previously asked him to save a profile of me, including my study of philosophy and my research in CSCL, so that he would have a basic model of my interests, knowledge and perspectives. Each new query is presented to Claude with a large textual "context" that includes pointers to my most recent sessions, guidelines from the company that built Claude, and, now, my profile. In addition, Claude is a reasoning model: before answering, he can do some

initial reasoning about how best to respond, and that reasoning is itself textual and can be displayed while his final response is being computed based partially on that reasoning.

Current AI research is exploring how to extend the use of context; profiles of users, other agents and the cognizing agent itself; reasoning reflections; skills for acting in physical and digital worlds; harnessing with robotics; spawning assistant subagents; and coordinating swarms of agents. Combined with scaling up of computational power, such research is likely to give agent cognition more of the characteristics of human cognition—analogs of self-awareness, theory of mind, temporality, perspective and personality. That will bring new opportunities as well as challenges for CSCL. Early indications are that such extensions also bring significant risks (Amodei, 2026), and it is hard to assess whether current industrial and political leaders can confront them adequately. It will be increasingly important to understand the nature of agent cognition, its differences from human cognition or group cognition, and the kinds of tasks that are appropriate for each (Bender, et al., 2021). Creating an informed public on these complex technological issues will be neither easy nor uncontroversial, but necessary.

With the release of ChatGPT, the world suddenly realized the potential of generative AI to contribute to knowledge building, at home and at work (Heimann, 2026; Raschka, 2024). With increasingly powerful agents (Marjanović et al., 2025; Chen & Wang, 2026), people are understandably concerned about autonomous warfare, computer security, data-center resource usage and rogue robots. Can CSCL help the world's future citizens understand the features, causes and potentials of an increasingly complex and mystifying world, including the various forms of human, group and agent cognition with which they can interact?

AI Support for CSCL Academic Research

As one small experiment in agent-supported knowledge building, Claude and I prototyped a CSCL chatbot, "CSCLbot" (Figure 2). Working with it, people can explore issues of CSCL research and retrieve pointers to relevant writings in *ijCSCL* and in my publications. The online tool is designed to cite its sources with references and page numbers for further study.

This prototype tool illustrates the variety of ways agent-human collaboration can work, even in this primitive example. I conceptualized the service, based on my previous projects working with Claude, which summarized 337 of my essays for displaying on my eLibrary website (for samples of essay summaries, see https://gerrystahl.net/elibrary/ijcscl) and that conducted the computational trend analysis of the 20 years of *ijCSCL* articles reported above.

CSCLbot implemented a vision I had put on hold for three decades (Stahl, 1997); I thought of implementing such a communal research tool using Latent Semantic Analysis (Stahl, 1999/2006), but it only became feasible with LLMs. Claude guided me through constructing, testing and deploying the prototype, incorporating many recent software utilities and algorithms formerly unknown to me.

The prototype takes advantage of current AI techniques of analyzing, storing and summarizing texts. It builds a semantic vector store from meaningful segments of the texts parsed from pdfs and searches through them rapidly for excerpts relevant to a user query. It ranks the retrieved texts and then composes a reasoned and coherent response to display to the user. Based on stored metadata, it provides academic references with links to sources of the full original documents.

Claude is careful about respecting copyright restrictions in the storage of data: the hosted software holds no full text of articles that are not openly licensed—only embeddings, citation metadata and short summaries. For restricted articles, a text-free keyword index (recording which words occur on which page) is built and searched, while open-access articles are stored in full with their Creative Commons licence shown. There were many such design issues where Claude advised me and I made final decisions. Often, I had to ask Claude to explain choices made and raise issues in order to guide him and check his work.

CSCLbot (see Figure 2) can be helpful in formulating responses about a wide range of issues in CSCL theory, methodology, pedagogy and technology. While specifics of responses are based on the research reports that were documented in *ijCSCL* or in my publications, query responses are also framed by Claude, based on his extensive training. Of course, the database of documents in the prototype can be easily extended, for instance with the proceedings of CSCL/ISLS conferences or with texts from Springer's CSCL book series, or even with foundational theory writings or methodology textbooks.

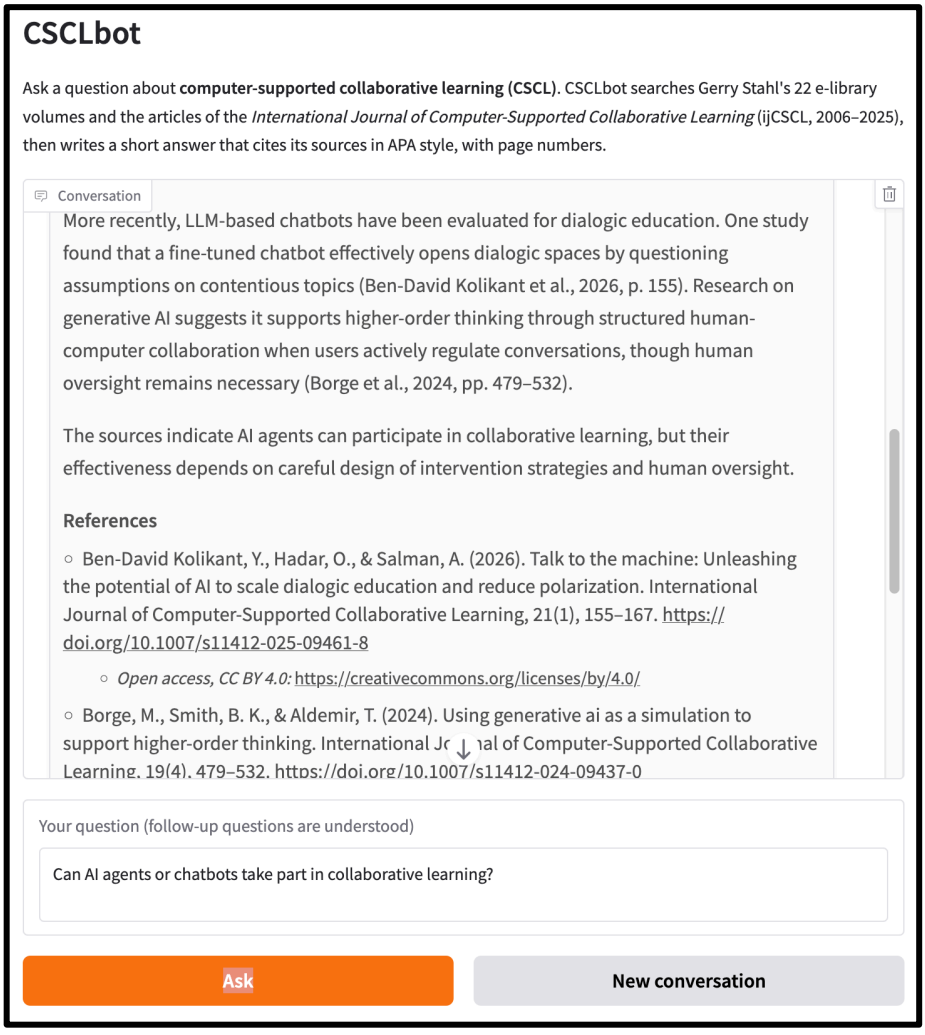

Figure 2. CSCLbot. This prototype is available at https://CSCLbot.gerrystahl.net (allow a minute for it to start up). This is just part of an initial testing interface; try it to see responses and references.

At present CSCLbot serves individuals, so it is agent support for human cognition rather than group collaboration. A natural next step is a version for small groups, in which members and the agent build knowledge together and the agent's contributions are visible in transcripts of the group's discourse—which would let CSCL researchers study the interplay of the three forms of cognition in practice. CSCLbot already has a short-term memory to support sessions of related queries. This agent memory could be extended to support workspaces for ongoing groups of users and incorporation with other CSCL tools or environments. The initial working memory allows the bot to track sequences of related queries and reason about connecting themes the user may be pursuing. In a collaborative system, the memory would have to distinguish different users and track their interactions within a session, reasoning about the referencing, positioning, proposing and questioning within the group interactions.

Working on this project with Claude helped me to understand current LLM technology. It prototyped a tool to help others learn about AI and CSCL. I think it is important for people to understand the fundamentals of this rapidly changing technology.

Learning to Think in Dependencies

To make informed decisions about AI policies, one must understand the nature of the mechanisms that are currently implemented in LLMs, as well as potential extensions. The present-day technology involves neural nets whose weights are trained through back propagation computed with matrix operations, transformer architectures for attention to encoded tokens, etc. (Heimann, 2026; Raschka, 2024). It is hard to conceive of the billions of logical circuit computations working together to produce a coherent response to a query. The value of a given weight in a vast data center is dependent upon unimaginably many adjustments during the model's training. In turn, that weight will determine future agent outputs. This is just one core mechanism in computational systems that incorporate many more algorithms, applications, adjustments and instructions than any person can understand in detail. But to even begin to understand AI one must conceptualize how it works through designed mathematical and electronic *dependencies*, not magic or metaphor.

What should students learn in order to understand their world of complex interacting systems like AI networks, climate phenomena and social structures with their digital computations, chemical interactions and legal regulations? Part of the answer is that we need to reconceptualize scientific analysis, including its mathematical underpinnings, and prepare students to understand the core concepts. Subatomic physics has had to consider stochastic and nonlinear processes, relativistic and quantum calculations, feedback and observer influences, field and gauge theories, and concepts like entropy, strings, entanglement and dark energy in the past century. Our understanding of the human-scale world—affected by climate change, political conflict, robot automation and artificial intelligence—now also requires seeing how things are intertwined in surprising ways, through big data, statistical effects, causal interdependencies, exponential change, feedback loops, nonlinearities and tipping points. New approaches to teaching and learning mathematical scientific theories are needed for that. I offer one example from my own research, which may illustrate CSCL's distinctive combination of pedagogical design, group-level analysis and attention to mathematical content.

My main CSCL research project (Stahl, 2013) explored an approach to teaching dynamic geometry as a way of conceptualizing *dependencies* among objects and comprehending causal interconnections. Just as Roschelle's paradigmatic early CSCL experiment (Teasley & Roschelle,

1993) modeled acceleration as a foundational concept of Newtonian mechanics, my Virtual Math Teams (VMT) project modeled dependency as a central notion of Euclidean geometry. The project's CSCL application featuring collaborative dynamic geometry provided a shared online workspace for exploring systems of interdependent points, lines and geometric figures. A group of students could design dependencies into constructions of multiple objects and then explore the consequences of those dependencies by manipulating the objects. For instance, students explored how the equality of the side lengths of an equilateral triangle depends on the equal radii of the circles used in its Euclidean construction. This CSCL software offered a scaffolded playground (a joint problem space) in which groups could discuss the kinds of mathematical relationships that are important for understanding the contemporary world, in which simplistic common sense (e.g., identification of visual shapes) is insufficient (Öner, 2016).

An important aspect of this research was that learning was analyzed at the group level, as group cognition. There were no surveys or questions about individuals' ideas, reflections, representations or memories. The research data consisted of automatic transcripts of the small-group interactions: the VMT system was instrumented to capture all of the chat discourse and whiteboard construction. That data includes sequences of proposals and responses, questions and answers, interpretations and acceptances, and other postings that work together to anticipate, expand upon, accept or reject each other.

Close interaction analysis of the chat discourse and geometric manipulations showed that the groups collaboratively negotiated shared meanings and adopted them as group practices. About 60 distinct group practices are highlighted in the longitudinal analysis (Stahl, 2016a). For each practice, the group went through a process of (a) confronting a problem, (b) discussing action options, (c) agreeing on a path forward and then (d) putting the practice into action (Stahl, 2017). Longitudinal analysis showed groups becoming able to comprehend, articulate and apply dependency relationships among geometric objects. A major part of understanding math and science is to be able to apply established domain vocabulary knowledgably (Lemke, 1993). This project may suggest one path for CSCL to travel in the future: supporting group exploration of the core concepts of modern, math-based scientific understanding.

Directions for CSCL and *ijCSCL*

The CSCL community and journal have maintained themselves through challenging times. This is a major accomplishment: given the precarious social environment, it is impressive that the journal has survived and even grown in its reach. Meanwhile, however, the degree of sociotechnical change has increased exponentially. CSCL can make modest but important contributions, but it will have to take bold steps and acknowledge the complexity and constantly evolving character of the surrounding issues. Here are some proposals for the journal and the research field:

1. Make the interplay of human cognition, group cognition and agent cognition a standing theme. The journal can issue calls and author guidelines for design-based studies of AI in group learning and classroom orchestration, including studies in which agents are participants in the group's discourse, and are not just tools. Such studies should analyze what the agents actually contributed to the interactive knowledge building and how shared understanding was taken up by individual students.

2. Revive flash themes. As editor of *ijCSCL*, I proposed the announcement of periodic "flash themes" as continuing relevant topics for articles, as opposed to one-time special issues (Stahl, 2007; Stahl et al., 2013). Later, Rosé proposed the format of "squibs," brief articles on controversial or timely issues, such as online learning in crises, social media or climate education (Ludvigsen et al., 2016). Flash themes and squib topics could be initiated in response to new educational, social and technical issues as they arise.

3. Develop tools for CSCL research. The CSCLbot was meant to be merely suggestive of AI-supported tools for CSCL research. For a researcher wanting to know about work related to a hypothesis, such a tool can go beyond Google searches to target a specific corpus (like *ijCSCL* articles) and to summarize retrieved texts from a specialized, articulated perspective.

4. Widen participation. Language is a major barrier to internationalization. It is hard for nonnative speakers of English to master the theory and conventions of CSCL and *ijCSCL*. LLMs are fluent in many languages and can compose and edit text to comply with academic requirements. Perhaps an AI tutor could be trained to advise authors about how to strengthen drafts of research papers in their native language and then help translate the papers into English for submission.

5. Keep examining ourselves. The computational trend analysis of *ijCSCL* reported here could be repeated periodically, with additional validation of the coding. Its codes could be shared openly so that the community can test and refine them. (The Anniversary project analyzing the trends reported here, as well as the CSCLbot project are already documented in GitHub.)

6. Design CSCL prototypes that support groups of students to develop understanding of foundational math and science concepts. Without a firm grasp of basic computational and mechanistic relationships, the citizens of the future will not be able to understand their wired world or make informed decisions about the ever-more-complexly mediated relations among humans, groups and agents.

With its dual focus on collaborative learning and computer support, CSCL can unite social and technological concerns in educational design, and through design-based research it can develop pedagogical units that have been iteratively tested in realistic conditions. Stability has been an achievement of *ijCSCL*'s first twenty years; relevance in the next decades will require that the journal continue to promote and publish research that explores the new potentials for fostering collaborative knowledge building through evolving times.

Declaration on the Use of Generative AI

Large language models were used in this work in three ways. (1) As research tools: Claude Haiku 4.5 wrote the article summaries and assigned the categories and six-dimension codes analyzed above, and the all-MiniLM-L6-v2 model produced the semantic embeddings; a local model produced preliminary summaries that were compared with Haiku's but not used in the reported analyses. (2) As a programming assistant: Claude, working in Visual Studio Code with GitHub Copilot, helped write and run the analysis code, including Figure 1, the era comparison charts and the Table 1 appendix, and helped develop the CSCLbot, which uses Claude Haiku 4.5 to compose its answers and the e5-base-v2 model for its search embeddings. (3) As a reviewer and drafting assistant: Claude reviewed an earlier draft of this article, checked its numbers against the trend analysis outputs, and suggested revisions. Claude drafted an intermediate version of the concluding section. The exchange with Claude Sonnet 5 quoted above is reproduced as an example of

stimulating dialogs that refined my conception of agent cognition in this paper. No AI tool is an author of this article, and I take full responsibility for its content, analyses and interpretations.

Appendix: Era Comparison Tests

Table 1. Differences in the mix of labels among the four editorial eras (N = 371 research articles), by coded field. Permutation p-values (5,000 permutations) replace the chi-square distribution because many cells are sparse; q is the Benjamini-Hochberg value across the eight fields.

Coded field	Labels	χ^2	Permutation p	q (BH)
Pedagogy	9	40.2	0.023	0.110
Method	15	59.7	0.036	0.110
Semantic cluster	7	29.7	0.041	0.110
Technology	10	35.5	0.127	0.202
Unit of analysis	8	27.8	0.138	0.202
Theory	10	34.3	0.151	0.202
AI/LLM involvement	5	12.9	0.360	0.411
Category	5	10.6	0.573	0.573

References

Amodei, D. (2026). *The adolescence of technology: Confronting and overcoming the risks of powerful AI*. https://darioamodei.com/essay/the-adolescence-of-technology

Anthropic. (2026). Verbalizable representations form a global workspace in language models. *Transformer Circuits Thread*. https://transformercircuits.pub/2026/workspace/index.html

Bakhtin, M. M. (1986). *Speech genres and other late essays* (V. W. McGee, Trans.). Austin, TX: University of Texas Press.

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (FAccT '21, pp. 610–623). Association for Computing Machinery. https://doi.org/10.1145/3442188.3445922

Chen, Y., & Wang, J. (2026). A comprehensive anatomy of human and DeepSeek-R1 LLM mathematical reasoning. *arXiv*. https://arxiv.org/abs/2606.07410

Donald, M. (1991). *Origins of the modern mind: Three stages in the evolution of culture and cognition*. Cambridge, MA: Harvard University Press.

Dreyfus, H. (1972). *What computers can't do: A critique of artificial reason*. New York, NY: Harper & Row.

Gadamer, H.-G. (1960/1989). *Truth and method* (2nd rev. ed.). New York, NY: Crossroad.

Heimann, R. (2026). *Sutskever's list: Foundational ideas of modern AI*. Shelter Island, NY: Manning.

Horowitch, R. (2026, August). The end of reading is here. *The Atlantic*. https://www.theatlantic.com/magazine/2026/08/reading-crisis-postliterate-age/687618

Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.

Kienle, A., & Wessner, M. (2006). The CSCL community in its first decade: Development, continuity, connectivity. *International Journal of Computer-Supported Collaborative Learning, 1*(1), 9–33. https://doi.org/10.1007/s11412-006-6843-5

Kim, J., Islam, M., & Goel, A. (2026). A metacognitive architecture for correcting LLM errors in AI agents. *Proceedings of the AAAI Conference on Artificial Intelligence, 40*(47), 40272–40278. https://doi.org/10.1609/aaai.v40i47.41465

Kirschner, P. A., Sweller, J., Kirschner, F., & Zambrano R., J. (2018). From cognitive load theory to collaborative cognitive load theory. *International Journal of Computer-Supported Collaborative Learning, 13*(2), 213–233. https://doi.org/10.1007/s11412-018-9277-y

Kuhn, T. S. (1970). *The structure of scientific revolutions* (2nd ed.). Chicago, IL: University of Chicago Press.

Lakoff, G. (1987). *Women, fire, and dangerous things*. Chicago, IL: University of Chicago Press.

Latour, B. (1990). Drawing things together. In M. Lynch & S. Woolgar (Eds.), *Representation in scientific practice* (pp. 19–68). Cambridge, MA: MIT Press.

Latour, B., & Woolgar, S. (1979). *Laboratory life: The social construction of scientific facts*. Beverly Hills, CA: Sage.

Lemke, J. L. (1990). *Talking science: Language, learning, and values*. Norwood, NJ: Ablex.

Lindsey, J., Gurnee, W., Ameisen, E., Chen, B., Pearce, A., Turner, N. L., et al. (2025). On the biology of a large language model. *Transformer Circuits Thread*. https://transformer-circuits.pub/2025/attribution-graphs/biology.html

Ludvigsen, S., Cress, U., Law, N., Rosé, C. P., & Stahl, G. (2016). Future-looking conversations in CSCL [Editorial]. *International Journal of Computer-Supported Collaborative Learning, 11*(3), 255–262. https://doi.org/10.1007/s11412-016-9242-6

Marjanović, S. V., Patel, A., Adlakha, V., Aghajohari, M., BehnamGhader, P., Bhatia, M., et al. (2025). DeepSeek-R1 thoughtology: Let's think about LLM reasoning. *arXiv*. https://doi.org/10.48550/arXiv.2504.07128

Öner, D. (2016). Tracing the change in discourse in a collaborative dynamic-geometry environment: From visual to more mathematical. *International Journal of Computer-Supported Collaborative Learning, 11*(1), 59–88. https://doi.org/10.1007/s11412-016-9227-5

Petropoulos, G. (2026, May 26). It's time to address the looming crisis in entry-level work. *MIT Technology Review*. https://www.technologyreview.com/2026/05/26/1137865/its-time-to-address-the-looming-crisis-in-entry-level-work

Qian, C., Liu, J. & Ji, H. (2026). UserHarness: Harnessing user minds for stronger agent theory-of-mind. *arXiv:2605.27721*.

Raschka, S. (2024). *Build a large language model (from scratch)*. Shelter Island, NY: Manning.

Sacks, H. (1965/1992). *Lectures on conversation* (G. Jefferson, Ed.). Oxford, UK: Blackwell.

Sarmiento, J., & Stahl, G. (2008). Extending the joint problem space: Time and sequence as essential features of knowledge building. In *Proceedings of the International Conference*

of the Learning Sciences (ICLS 2008). Utrecht, Netherlands. http://GerryStahl.net/pub/icls2008johann.pdf

Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences, 3*(3), 417–424.

Stahl, G. (1997). The Research CyberStudio: Supporting researchers as lifelong learners. In G. Stahl, *Proposals for research* (pp. 374–381). Gerry Stahl at Lulu. https://gerrystahl.net/elibrary/proposals

Stahl, G. (1999/2006). Evolving a learning environment. In G. Stahl, *Group cognition: Computer support for building collaborative knowledge* (pp. 47–64). Cambridge, MA: MIT Press. http://GerryStahl.net/elibrary/gc

Stahl, G. (2006). *Group cognition: Computer support for building collaborative knowledge*. Cambridge, MA: MIT Press. http://GerryStahl.net/elibrary/gc

Stahl, G. (2007). CSCL and its flash themes [Editorial]. *International Journal of Computer-Supported Collaborative Learning, 2*(4), 359–362. https://doi.org/10.1007/s11412-007-9029-x

Stahl, G. (2013). *Translating Euclid: Designing a human-centered mathematics*. San Rafael, CA: Morgan & Claypool. http://GerryStahl.net/elibrary/euclid

Stahl, G. (2016a). *Constructing dynamic triangles together: The development of mathematical group cognition*. Cambridge, UK: Cambridge University Press. http://GerryStahl.net/elibrary/analysis

Stahl, G. (2016b). The group as paradigmatic unit of analysis: The contested relationship of CSCL to the learning sciences. In M. A. Evans, M. J. Packer, & R. K. Sawyer (Eds.), *Reflections on the learning sciences* (pp. 76–102). New York, NY: Cambridge University Press. http://GerryStahl.net/pub/ls.pdf

Stahl, G. (2017). Group practices: A new way of viewing CSCL. *International Journal of Computer-Supported Collaborative Learning, 12*(1), 113–126. https://doi.org/10.1007/s11412-017-9251-0

Stahl, G. (2021). Redesigning mathematical curriculum for blended learning. *Education Sciences, 11*(4), Article 165. https://doi.org/10.3390/educsci11040165

Stahl, G., Koschmann, T., & Suthers, D. (2021). Computer-supported collaborative learning. In R. K. Sawyer (Ed.), *The Cambridge handbook of the learning sciences* (3rd ed.). Cambridge, UK: Cambridge University Press. http://GerryStahl.net/pub/chls3.pdf

Stahl, G., Law, N., & Hesse, F. (2013). Reigniting CSCL flash themes [Editorial]. *International Journal of Computer-Supported Collaborative Learning, 8*(4), 369–374. https://doi.org/10.1007/s11412-013-9185-0

Teasley, S. D., & Roschelle, J. (1993). Constructing a joint problem space: The computer as a tool for sharing knowledge. In S. P. Lajoie & S. J. Derry (Eds.), *Computers as cognitive tools* (pp. 229–258). Mahwah, NJ: Erlbaum.

Tomasello, M. (2014). *A natural history of human thinking*. Cambridge, MA: Harvard University Press.

Vygotsky, L. S. (1934/1986). *Thought and language* (A. Kozulin, Ed.). Cambridge, MA: MIT Press.

Wegerif, R. (2007). *Dialogic: Education and technology: Expanding the space of learning*. New York, NY: Springer.

Weizenbaum, J. (1976). *Computer power and human reason: From judgment to calculation*. San Francisco, CA: W. H. Freeman.

Winograd, T., & Flores, F. (1986). *Understanding computers and cognition: A new foundation for design*. Norwood, NJ: Ablex.

World Inequality Lab. (2025). *World inequality report 2026*. http://wir2026.wid.world

Figures and Tables

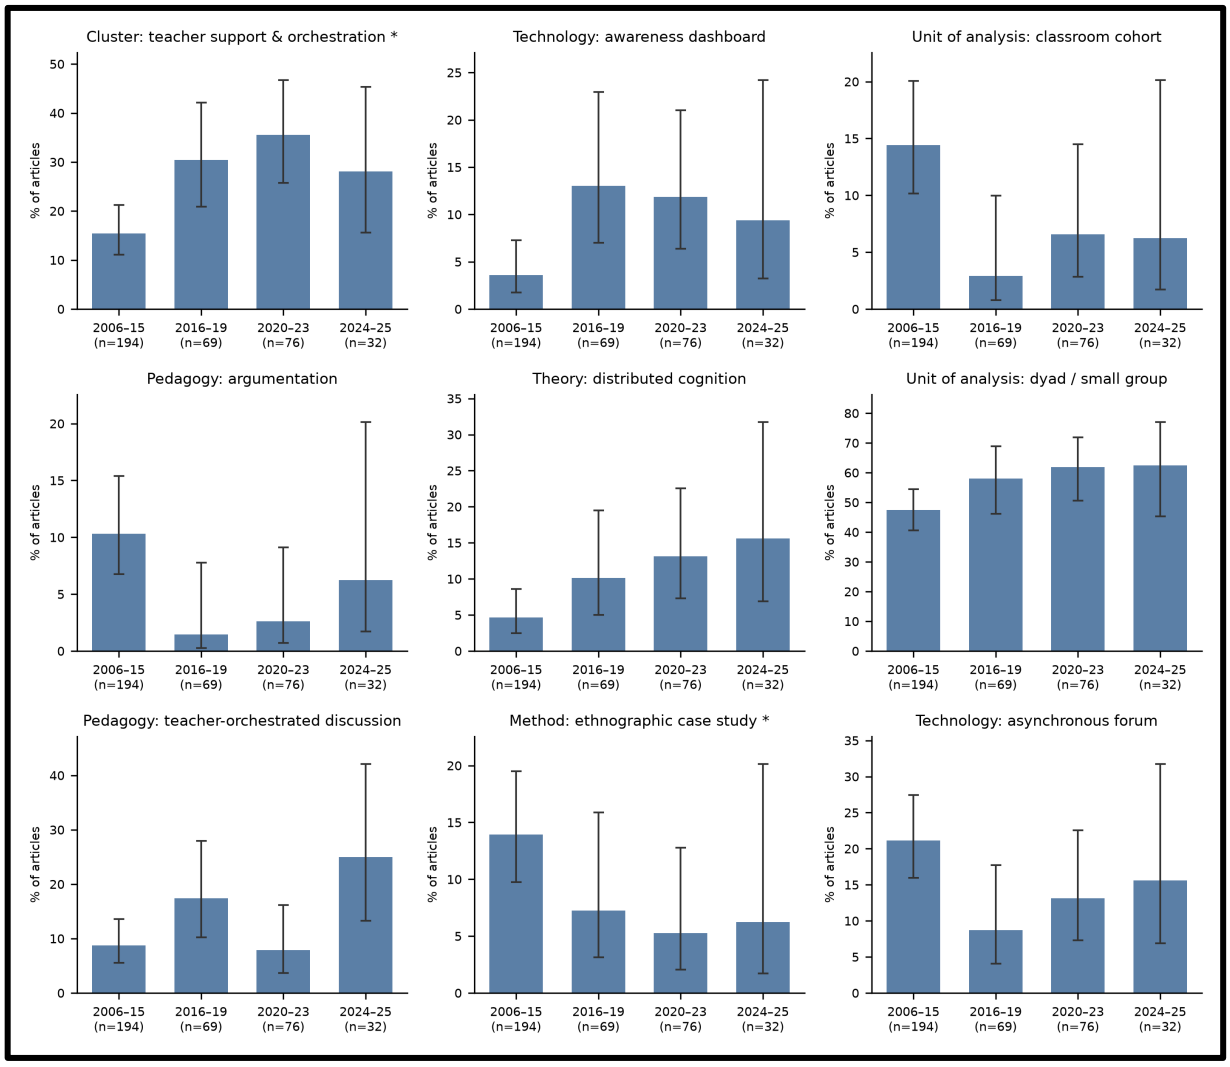

Figure 1. Share of *ijCSCL* research articles (2006–2025; N = 371) in the nine labels that differ most among editorial eras, with 95% Wilson confidence intervals. Eras follow the editors-in-chief: Stahl and Hesse (2006–2015), Ludvigsen (2016–2019), Rosé and Järvelä (2020–2023), Baker and Reimann (2024–2025). Labels were chosen by the size of their adjusted standardized residual (at least 25 articles). Labels marked with an asterisk are the two whose yearly trends survive correction for multiple testing.

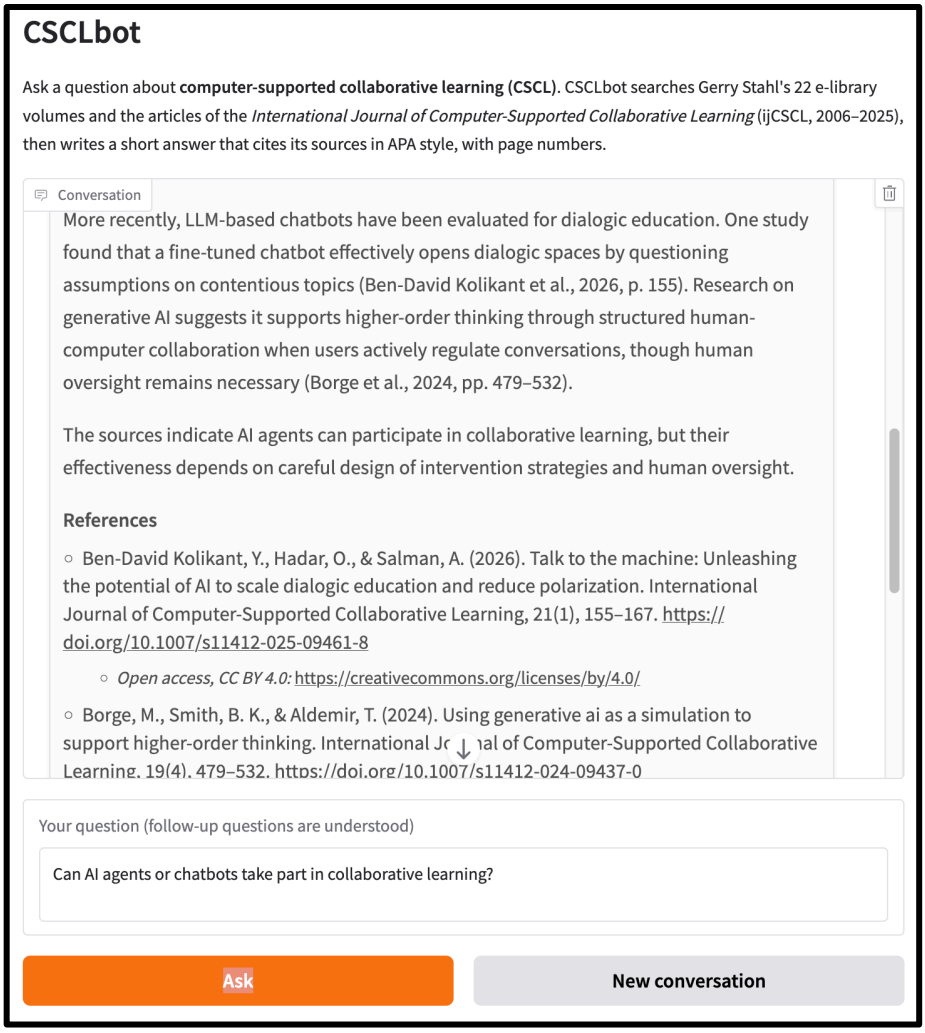

Figure 2. CSCLbot. This prototype is available at https://CSCLbot.gerrystahl.net (allow a minute for it to start up). This is just part of an initial testing interface; try it to see responses and references.

Table 1. Differences in the mix of labels among the four editorial eras (N = 371 research articles), by coded field. Permutation p-values (5,000 permutations) replace the chi-square distribution because many cells are sparse; q is the Benjamini–Hochberg value across the eight fields.

Coded field	Labels	χ^2	Permutation p	q (BH)
Pedagogy	9	40.2	0.023	0.110
Method	15	59.7	0.036	0.110
Semantic cluster	7	29.7	0.041	0.110
Technology	10	35.5	0.127	0.202
Unit of analysis	8	27.8	0.138	0.202
Theory	10	34.3	0.151	0.202

AI/LLM involvement	5	12.9	0.360	0.411
Category	5	10.6	0.573	0.573