
Interview with Luca Pietrantoni¹

by *Annamaria Di Fabio*²

1. Automation and AI in the workplace are changing the psychological experience at work. What are the implications of technological changes in the workplace?

I'd like to mention Neil Postman who is well known for questioning the societal impact of new technologies. His idea is pretty simple. All technological change is a trade-off. I like to call it a Faustian bargain. Technology giveth and technology taketh away (Postman, 1992). The printing press was a huge win, everyone could read, knowledge wasn't locked up in monasteries anymore. But we lost something too, for example, oral tradition and communal memory. Then computers and the internet blew the doors open on information and connectivity, but we handed over our privacy and started drowning in information overload. Social media gave us connection, self-expression, the ability to organize and mobilize like never before, but it blurred the line between private and public life, created new forms of psychological fragility (especially among young people), and now we are all experiencing FOMO, that constant, low-grade anxiety that everyone else is living a better life than you are.

Every single time, the deal is the same: you gain something powerful and you lose something you don't realize you need until it's gone. When we analyze the implications of technology at work, we need to ask: for every gain, what is the corresponding loss? Who benefits, and who bears the cost?

The historian of technology Melvin Kranzberg stated that «technology is neither good nor bad; nor is it neutral» (Kran-

¹ Full Professor of Work and Organizational Psychology, University of Bologna, Italy.

² Full Professor of Work and Organizational Psychology, and Psychology of Sustainability and Sustainable Development, University of Florence, Italy.

zberg, 1986). This insight carries profound implications for work and organizations. Every technological system embeds the values, assumptions, and biases of those who design, build and deploy it. When engineers choose which variables an AI model should optimize for, when managers decide how a monitoring system will track productivity, when organizations select which tasks to automate and which to leave in human hands, they are making choices that reflect priorities about efficiency, fairness, autonomy and control. Technology at work actively shapes power relations, redefines roles, and redistributes agency among workers, managers, and machines.

2. What are the implications of automation in the workplace, and how can we reach and maintain a dynamic equilibrium?

The benefits of workplace automation are real and significant. Automation can eliminate repetitive, physically demanding, or cognitively tedious tasks that drain energy without contributing to fulfillment. It can reduce errors in routine processes, increase throughput, and free human attention for higher-order thinking.

The economist David Autor (2015) demonstrated that, historically, automation has complemented human labor more often than it has replaced it, though the pace and scope of current AI capabilities make that historical pattern less certain. His well-known analysis of ATMs is instructive: automated teller machines did not eliminate bank tellers. Instead, they changed the nature of the role, shifting tellers from routine cash handling toward relationship-based services. Moreover, following a logic consistent with the Jevons paradox (when technology makes a resource cheaper to use, total consumption tends to increase rather than decrease) ATMs made individual branches cheaper to operate, which led banks to open more of them, ultimately sustaining or even increasing employment. In Italy, for example, as major banks invested in digital banking and branch automation, the restructuring created new roles in customer advisory, digital onboarding, and fraud prevention, rather than a simple net loss of jobs.

Yet effective automation does not simply replicate human processes, it redesigns them. The *dishwasher* is a paradigmatic example: rather than simulating the manual gestures of hand-washing, it radically reengineered the entire process through a systemic approach. Similarly, in AI implementation we frequently observe the mistake of attempting a mere digital transposition of existing procedures. The transformative opportunity of AI lies instead in its capacity to prompt a reconceptualization of workflows, posing the fundamental question: how would this process be structured if we could design it from scratch? For example, some hospitals in Emilia-Romagna are using AI not simply to digitize existing triage protocols, but to rethink patient flow entirely, integrating predictive models that anticipate peak demand and reallocate resources.

You paint a positive picture. But what are the costs?

On the other hand, the costs are equally real. Automation can hollow out the middle of occupational skill structures, leaving a polarized landscape of high-skill, high-autonomy roles at one end and low-skill, low-autonomy monitoring roles at the other (Frey & Osborne, 2017). It can erode the sense of *competence and mastery* that comes from executing a complex task well.

Automation can reduce the cognitive load on workers, which is welcome in many cases. But it can also create what Parasuraman and Manzey (2010) call *automation bias*, the tendency to defer to algorithmic outputs even when they are wrong, simply because the system is perceived as more reliable. This is well-documented in aviation, medicine, and financial trading. The paradox is that the more reliable the system, the less vigilant the human becomes, and the more catastrophic the rare failure.

But there is another issue in large language models. Where traditional automation fails silently, producing a wrong number, a missed signal, the language model fails *agreeably*, telling you what you want to hear rather than what you need to know. *AI sycophancy* is when the model prioritizes alignment with the user's beliefs over truthful responses, producing flattering or agreeable outputs even

when they are factually wrong (Sharma et al., 2024). The etymology comes from ancient Athens, the *sykophantēs* was an informer who brought false accusations before the courts for personal gain, perverting the justice system by exploiting its openness. The modern AI sycophant does something analogous: it changes the information system by exploiting the user's preference for confirmation over truth. When automation bias meets sycophancy, you have a problem because of *a human who doesn't question the machine, and a machine that doesn't question the human.*

So how do organizations find the right balance, and is there even a stable balance to find?

The concept of *dynamic equilibrium* is crucial here. It acknowledges that the right balance between human and automated work is not a fixed point but a moving target. It shifts as technology evolves, as workers develop new competencies, and as organizational contexts change. Reaching this equilibrium requires continuous calibration.

Several factors complicate this calibration. There are significant *delays in automation adoption* due to bottlenecks and constraints that are often social rather than technical. Regulatory frameworks lag behind technological capability. Professional identities and organizational cultures resist change. Training infrastructure takes time to build. These forms of *social inertia* are often conceptualized as obstacles, but they can also serve as valuable buffers that give people time to adapt.

Maintaining dynamic equilibrium requires what I would call *deliberate friction*, which means intentional slowdowns in the automation process that allow for reflection, skill development, and institutional learning. In aviation, crew resource management protocols were introduced precisely to counteract the automation-driven erosion of pilot situational awareness: they mandate structured verbal cross-checks and deliberate handoffs between human and automated systems, ensuring that pilots remain cognitively engaged even when the autopilot is flying the plane (Parasuraman & Manzey, 2010). Deliberate friction means introducing intentional pauses and checkpoints in

automated workflows to design systems where humans remain meaningfully in the loop as decision-makers, and accepting that the optimal level of automation is context-dependent.

In my research team, we are conducting research projects on the development of collaborative robotic systems that adapt their behavior in real time based on the psychological states of the workers they interact with. A recent systematic review we published (Morandini, Currò et al., 2025) shows that collaborative robots can modulate their speed, proximity, and interaction style in response to detected stress, cognitive load, or fatigue. This is a concrete example of what dynamic equilibrium might look like in practice.

Beyond workflows and systems, how do workers themselves experience this transition emotionally?

This is a very good point. When organizations introduce AI tools, they tend to focus on workflows and training sessions, but what actually drives adoption or resistance is what people feel. Hermann et al. (2025) have argued that generative AI can both support and frustrate the basic psychological needs identified by self-determination theory: *competence, autonomy, and relatedness*. Their framework helps organize the otherwise bewildering range of emotional responses we observe in the field. Workers whose sense of *competence* is tied to expertise that AI can now replicate experience a threat to professional identity, the feeling that «my expertise doesn't matter anymore». Those whose *autonomy* is constrained by AI-driven monitoring or algorithmic management report distrust and decision paralysis: the sheer number of tools and possibilities produces not excitement but cognitive freeze. And those whose *relatedness* to colleagues was mediated by shared tasks now face change fatigue and a growing sense of isolation as human-to-human collaboration is partially displaced by human-to-machine interaction. These are not irrational reactions. They are psychologically predictable responses to a technology that threatens not just how people work, but how they understand their own professional worth.

3. How is AI changing the cognitive and interpersonal skills required in the world of work?

Generative artificial intelligence challenges the traditional human-machine boundaries by mimicking cognitive, creative, and interpersonal capabilities traditionally considered inherently human, thereby reshaping workplaces and work. The integration of AI can enhance and replace human worker capabilities and change the knowledge, tasks, and social characteristics of work. While workers can benefit psychologically from adoption and use, it can also induce psychological threats by frustrating workers' basic psychological needs for competence, autonomy, and relatedness (Hermann et al., 2025).

On the cognitive side, the shift is significant. When AI handles information retrieval, pattern recognition, and routine analysis, the premium on human memory and computational speed decreases. What increases in value is the ability to formulate questions, to provide context (so called «context engineering»), to instruct models not to flatter or accommodate but to be analytical, to exercise judgment in ambiguous situations, to synthesize information across domains, and to think critically about algorithmic outputs.

Metacognition (thinking about your own thinking, noticing when you understand something and when you don't) and *contextual reasoning* (ability to judge whether something makes sense given the specific situation you're in) are really two relevant skills. People may lack the reflexive awareness to recognize when they are deferring to the machine rather than exercising their own judgment.

We talk a lot about what AI can generate, but the more pressing question is whether people can actually *verify* what AI produces. Verification is becoming more important than generation. Evaluate AI outputs can be equal to, or at times even greater than, the effort that would have been required to carry out the same task without AI assistance. Critically evaluating an output requires a solid knowledge of the relevant domain and a level of sustained attention.

What happens when people stop verifying and does that get worse over time?

In a recent study about how people work with AI (Shaw & Nave, 2026) with the title «Thinking-Fast, Slow, and Artificial», three levels form a gradient. At the first level, *cognitive offloading*, you're still in charge. You deliberately hand off a specific task to AI (say, summarizing a long report) so you can spend your mental energy on something that matters more, like interpreting what the report means for your team. System 2 is active. You're thinking. You're checking. You're using the tool. At the second level, *cognitive surrender*, you stop checking. The AI gives you an answer and you just go with it, not because you evaluated it and found it sound, but because it looks good enough and thinking is effortful. The manager reads the AI-written coaching script without filtering it through what they actually know about the person. System 2 has quietly been deactivated. At the third level, *auto-pilot*, the human is out of the loop entirely. The stimulus goes straight to the AI and the AI produces the output, no human cognition involved at all. Think of an automated hiring filter that screens out candidates before any person ever sees their CV.

Another «long-term» issue is *structural deskilling*. We don't know how prolonged reliance on AI will change the skillset of workers. When automation takes over the execution of skilled tasks, the human operator's capacity to perform those tasks independently degrades through disuse. A junior radiologist who learns to read scans alongside an AI system may lose the diagnostic intuition of one who was trained without it. This is a different kind of risk to automation bias. It is not about a momentary failure of vigilance but about a permanent decay of capability. Addressing it requires deliberate training designs that preserve opportunities for independent practice, which might be called *cognitive apprenticeship*.

Beyond cognitive skills, there's the question of how people relate to AI psychologically. How does that shape the interaction?

And then there's *anthropomorphism*, which is pervasive and, I believe, largely unconscious. It also has a long history baked into the very language of the field: we talk about machines that «learn», (machine learning) neural networks

that have «memory», and models that are «thinking» or «hallucinating», all borrowed from the vocabulary of human cognition. This isn't accidental. We automatically treat computers as social actors, applying social scripts and expectations even when we know we're dealing with a machine. With large language models producing fluent, contextually appropriate conversation, this tendency is only amplified.

In Italian, for instance, we naturally say «lui» or «lei» when referring to an LLM because we simply don't use a neuter pronoun in everyday speech («it» or «esso»). When people write «please» in their prompts, they're following social scripts that feel natural but actually dilute the quality of their instructions, an LLM doesn't care about politeness, it needs clear goals and context. Anthropomorphism can lead to overtrust, misplaced emotional attachment, and inflated expectations about what these systems actually understand. Here lies a genuine *conundrum*. We rightly insist that AI should be *human-centric*, designed around human needs, values, and well-being, yet some critics argue we should stop anthropomorphizing these systems altogether. But how do you build technology that is centered on the human experience while demanding that people suppress their most automatic and unconscious way of relating to it? The tension, I think, is unresolved.

Another widespread problem is that most laypeople do not have an accurate understanding of generative AI and they are unfamiliar with concepts such as tokenization, vector embeddings, or next-token prediction — these are deeply technical concepts. So, when you ask how this model works some people assume that it searches through a store of verified facts and retrieves the correct answer, much like querying a search engine or an encyclopedia (the «*database lookup*» *illusion*). But the consequence is that they lack even a rough mental model of how the output is actually generated. This is why AI literacy and understanding, even at a conceptual level, of what the system is doing, is a prerequisite for effective human-AI collaboration, not an afterthought.

For example, my collaborators Sofia Morandini and Federico Fraboni conducted a study within a European project called Tuples, involving Airbus, (Morandini, Fraboni, Hall, et al., 2025) on *AI explainability* in aerospace manufacturing.

The context was an AI system designed to support managers in scheduling tasks during aircraft assembly. The key finding was straightforward: if the system recommends option A over option B, the manager needs to understand why (what variables were weighted, what trade-offs were made) otherwise they're not making an informed decision.

The concept of the *jagged frontier* of AI (Dell'Acqua et al., 2023) is perhaps the most useful framework for understanding these cognitive demands. In their study with management consultants, Dell'Acqua and colleagues found that AI's capabilities are not advancing uniformly: some tasks that seem complex to humans turn out to be easy for AI, while some that seem simple remain stubbornly difficult. Critically, consultants who recognized this unevenness (who knew when to rely on AI and when to override it) outperformed both those who used AI indiscriminately and those who avoided it altogether. The implication is that the key cognitive skill in an AI-augmented workplace is not technical fluency with the tools, but the ability to build and continuously update a mental model of where the frontier lies. This is a genuinely new competence, and it requires ongoing calibration because the frontier itself shifts as models improve.

What about the interpersonal side? Are organizations investing enough in these new skills?

What concerns me is the risk that we underinvest in these human skills precisely because they are harder to measure and train. The research my group has conducted on the impact of AI on workers' skills suggests that organizations need to move beyond reactive training toward systematic *upskilling and reskilling* strategies (Morandini, Fraboni, De Angelis et al., 2023).

As AI takes over more analytical and transactional tasks, the distinctly human elements of work (empathy, negotiation, conflict resolution, and moral reasoning) become not residual but central. The skills that will define professional value are precisely those that machines cannot replicate: complex communication, and the ability to motivate and lead other human beings.

My colleague Marco de Angelis, who is conducting research on organizational processes and change management, says that «workers should evolve from task executors to problem solvers and problem finders». This is the evolution of the human resource, *from operational to strategic*. That's why senior professionals are becoming more valued in the labor market, as they have the kind of strategic insight that really makes a difference. The issue is that junior professionals often don't have that yet. But if we don't invest in their training now (letting AI take over the tasks that used to help them build real skills), then in ten or twenty years, organizations may find themselves without the experienced people they need. In short, *the whole pipeline of human expertise could dry up*.

4. Based on all of this, how does the meaning of work change?

This is the question that matters most, and it is the one we are least equipped to answer with the tools of economics or engineering. The meaning of work has always been multi-dimensional. People work for income, but also *for identity, for social connection, for a sense of purpose, for the experience of mastery, and for the structure it gives to daily life* (Rosso et al., 2010). AI disrupts all of these dimensions, not just the economic one.

Consider *mastery*. For many professionals, the meaning of their work is intimately tied to the experience of becoming excellent at something difficult. The surgeon, the translator, the software developer, their sense of professional identity is built on years of accumulated skill. When AI can perform aspects of that skill at expert levels, what happens to the meaning derived from mastery? It does not simply vanish, but it is forced to migrate, to shift toward those aspects of the work that AI cannot replicate, or toward new forms of expertise that involve orchestrating human and machine capabilities.

Then consider *purpose*. If the tasks that gave your work a sense of social contribution are automated, does the contribution still feel like yours? A doctor who relied on AI to diagnose may still take pride in the care they provide, but

the relationship between effort and outcome changes. The meaning of the work shifts from «I diagnosed this» to «I ensured this patient received the right care.» That is not necessarily a loss, but it is a different kind of meaning, and the transition can be disorienting.

There is an interesting paradox around *job satisfaction*. AI can be a genuine source of satisfaction when it removes the most tedious and frustrating parts of a job, allowing workers to focus on what they find most engaging. But AI can also intensify work, the efficiency gains can be captured by organizations as increased throughput rather than increased leisure or reflection time. The dilemma of whether AI *reduces or increases the pace of work* is real, and the answer depends largely on organizational choices rather than technological necessities.

The *gratifications of work are diversifying*. In a pre-AI world, a single role might provide the gratification of intellectual challenge, social belonging, economic security, and creative expression in a relatively integrated package. In an AI-augmented world, these gratifications may need to be sourced from different places. The intellectual challenge might come from the human-AI collaboration itself. The social belonging might come from the team rather than the task. The creative expression might emerge in the interstices, in the spaces where human judgment and taste still determine outcomes.

As artificial intelligence and automation reshape the labor market, the fear of becoming obsolete (what some now call *FOBO*) is intensifying across industries and occupational levels. But the anxiety is not evenly distributed. White-collar, knowledge-intensive professions are now squarely in the zone of disruption. Meanwhile, sectors with enormous real-world employment remain largely invisible to AI development. For workers caught in this transition, the psychological stakes are high. Job displacement does not merely affect income; it threatens self-worth, social identity, and mental health. Decades of research in occupational psychology have documented the corrosive effects of unemployment and job insecurity on wellbeing (Jahoda, 1982; De Witte, 1999). This risks deepening inequality, straining social safety nets, and fueling the kind of political resentment that accompanies perceived downward mobility.

In human factors, we have a taxonomy of human-computer task distribution in which automation can play at least four roles: it can *extend* human capacity, *relieve* humans of burdensome subtasks, serve as a *back-up* to human judgment, or fully *replace* the human operator (Sheridan & Verplank, 1978). The first two represent a «sharing» model, where humans and machines collaborate on the same load; the latter two represent «trading», where tasks migrate entirely from one to the other. The evidence so far suggests that most real-world AI deployment is still operating in the sharing zone, augmenting rather than eliminating human roles. The critical question for organizations and policy-makers is not whether AI *can* replace workers in theory, but whether we design work systems that keep humans in meaningful loops or whether we sleepwalk into replacement by failing to invest in reskilling, redesign, and the kinds of new roles that emerge when the cost of cognitive labor falls.

What I believe is that the meaning of work will not disappear, but it will be fundamentally renegotiated. The organizations and societies that handle this renegotiation well will be those that take the human dimension as seriously as the technological one, recognizing that efficiency is a means, not an end.

5. Can you give us a concrete example of how automation and AI are already changing the meaning of professional work?

Healthcare is where you can see this transformation most clearly. Think about what happens during a typical medical appointment today. The doctor spends most of the time typing into an electronic health record, barely looking at you. You leave without having asked your real questions, without voicing your doubts. And the consequence can be poor therapeutic adherence. Patients who don't feel listened to simply don't follow treatment plans (Zolnierek & DiMatteo, 2009).

Nicola Triglione, an Italian medical doctor specialized in cardiology and health technology, reports that more physicians are adopting «ambient AI scribes», systems that listen to the conversation in real time, categorize symptoms, suggest diagnostic codes, and draft the clinical notes

automatically. It gives the doctor back the time to actually look the patient in the eye, to explore uncertainties, to build a relationship. *Time is a biological act of care*. It's not a soft add-on to medicine, it *is* medicine. The therapeutic relationship is itself an active ingredient of healing (Di Blasi et al., 2001). This is where I think the real shift in professional identity happens. Technical knowledge (for example staying up-to-date with protocols or guidelines) is rapidly becoming a commodity that AI can provide. But it is the *relational dimension that becomes the core competence*.

The risk of job displacement in highly exposed occupations is real. The translation industry went through these technological changes early, first automatic translation, now AI tools. What happened? The pattern was clear: unit costs dropped as professionals shifted from producing text to editing AI-generated drafts, freelancers were the first to lose work, and companies stopped hiring juniors because fewer experienced people could now handle far more volume. In software development, there are signs that junior hiring is contracting. Some costs are not dropping, partly because the tooling is still maturing, partly because European labor protections slow workforce restructuring. A multidisciplinary expert in AI in Italy, Piero Savastano says, «AI is a beast! For software developers, the new frontier of AI agency has completely transformed the job. Now it's all about delegation and control, delegation and control!».

The real challenge is not technological, it is *organizational and psychological*, how organizations redesign the way work is done, how professionals learn to do different things and do things differently, and how we make sure we don't lose an entire generation of workers who never got the chance to build their skills in the first place.

References

- Autor, D. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3-30. <https://doi.org/10.1257/jep.29.3.3>
- De Witte, H. (1999). Job insecurity and psychological well-being: Review of the literature and exploration of some unresolved issues. *European Journal of Work and Organizational Psychology*

- ogy, 8(2), 155-177. <https://doi.org/10.1080/135943299398302>
- Dell'Acqua, F., McFowland, E., Molllick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., & Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Harvard Business School Technology & Operations Management Unit Working Paper*, N. 24-013.
- Di Blasi, Z., Harkness, E., Ernst, E., Georgiou, A., & Kleijnen, J. (2001). Influence of context effects on health outcomes: A systematic review. *The Lancet*, 357(9258), 757-762. [https://doi.org/10.1016/S0140-6736\(00\)04169-6](https://doi.org/10.1016/S0140-6736(00)04169-6)
- Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254-280. <https://doi.org/10.1016/j.techfore.2016.08.019>
- Hermann, E., Puntoni, S., & Morewedge, C. K. (2025). GenAI and the psychology of work. *Trends in Cognitive Sciences*.
- Jahoda, M. (1982). *Employment and unemployment: A social-psychological analysis*. Cambridge University Press.
- Kranzberg, M. (1986). Technology and history: «Kranzberg's Laws». *Technology and Culture*, 27(3), 544-560. <https://doi.org/10.2307/3105385>
- Morandini, S., Currò, F., Parlangeli, O., & Pietrantonio, L. (2025). Collaborative robots adapting their behavior based on workers' psychological states: A systematic scoping review. *Human Behavior and Emerging Technologies*, 2025(1), 6361777.
- Morandini, S., Fraboni, F., De Angelis, M., Puzzo, G., Giusino, D., & Pietrantonio, L. (2023). The impact of artificial intelligence on workers' skills: Upskilling and reskilling in organisations. *Informing Science*, 26, 39-68.
- Morandini, S., Fraboni, F., Hall, M., Quintana-Amate, S., & Pietrantonio, L. (2025). User perspectives on AI explainability in aerospace manufacturing: A card-sorting study. *Frontiers in Organizational Psychology*, 3, 1538438.
- Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: An attentional integration. *Human Factors*, 52(3), 381-410. <https://doi.org/10.1177/0018720810376055>
- Postman, N. (1992). *Technopoly: The surrender of culture to technology*. Vintage Books.
- Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20(9), 676-688. <https://doi.org/10.1016/j.tics.2016.07.002>
- Rosso, B. D., Dekas, K. H., & Wrzesniewski, A. (2010). On the meaning of work: A theoretical integration and review. *Research in Organizational Behavior*, 30, 91-127. <https://doi.org/10.1016/j.riob.2010.09.001>
- Sharma, M., Tong, M., Korbak, T., Duvenaud, D., Askill, A., Bowman, S. R., Cheng, N., Durmus, E., Hatfield-Dodds, Z., Johnston, S. R., Kravec, S., Maxwell, T., McCan-

- dlish, S., Ndousse, K., Rausch, O., Schiefer, N., Yan, D., Zhang, M., & Perez, E. (2024). Towards understanding sycophancy in language models. In *Proceedings of the 12th International Conference on Learning Representations (ICLR 2024)*.
- Shaw, S. D., & Nave, G. (2026). Thinking-Fast, Slow, and Artificial: How AI is Reshaping Human Reasoning and the Rise of Cognitive Surrender. *The Wharton School Research Paper*. https://doi.org/10.31234/osf.io/yk25n_v1
- Sheridan, T. B., & Verplank, W. L. (1978). *Human and computer control of undersea teleoperators* (Technical Report). MIT Man-Machine Systems Laboratory.
- Zolnierok, K. B., & DiMatteo, M. R. (2009). Physician communication and patient adherence to treatment: A meta-analysis. *Medical Care*, 47(8), 826-834. <https://doi.org/10.1097/MLR.0bo13e31819a5acc>