frontiers
in Astronomy and
Space Sciences

TYPE Original Research
PUBLISHED 02 June 2026
DOI 10.3389/fspas.2026.1829012

Check for updates

OPEN ACCESS

EDITED BY
Didier Fraix-Burnet,
UMR5274 Institut de Planétologie et d'Astrophysique de Grenoble (IPAG), France

REVIEWED BY
Yanxia Zhang,
Chinese Academy of Sciences (CAS), China
Massimiliano Bonamente,
University of Alabama in Huntsville, United States

*CORRESPONDENCE
Alessandro Cabras,
alessandro.cabras@inaf.it

RECEIVED 12 March 2026
REVISED 20 April 2026
ACCEPTED 20 April 2026
PUBLISHED 02 June 2026

CITATION
Cabras A, Marra M, Capasso G, Petrella A, Balestra A, Lo Cicero U, Riggi S, Smareglia R, Arcidiacono C and Costanzo S (2026) Adoption of generative AI at the Italian institute for astrophysics: a survey on usage trends, ethics, literacy, and Shadow IT. *Front. Astron. Space Sci.* 13:1829012. doi: 10.3389/fspas.2026.1829012

COPYRIGHT
© 2026 Cabras, Marra, Capasso, Petrella, Balestra, Lo Cicero, Riggi, Smareglia, Arcidiacono and Costanzo. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Adoption of generative AI at the Italian institute for astrophysics: a survey on usage trends, ethics, literacy, and Shadow IT

Alessandro Cabras[1]*, Monica Marra[2], Giulio Capasso[3], Amedeo Petrella[4], Andrea Balestra[4], Ugo Lo Cicero[5], Simone Riggi[6], Riccardo Smareglia[7], Carmelo Arcidiacono[8] and Saverio Costanzo[6]

[1]INAF (National Institute for Astrophysics), Osservatorio Astronomico di Cagliari, Selargius, Italy, [2]INAF (National Institute for Astrophysics), Osservatorio di Astrofisica e Scienza dello Spazio, Bologna, Italy, [3]INAF (National Institute for Astrophysics), Osservatorio Astronomico di Capodimonte, Napoli, Italy, [4]INAF (National Institute for Astrophysics), Osservatorio Astronomico di Padova, Padova, Italy, [5]INAF (National Institute for Astrophysics), Osservatorio Astronomico di Palermo, Palermo, Italy, [6]INAF (National Institute for Astrophysics), Osservatorio Astrofisico di Catania, Catania, Italy, [7]INAF (National Institute for Astrophysics), Osservatorio Astronomico di Trieste, Trieste, Italy, [8]INAF (National Institute for Astrophysics), Osservatorio Astrofisico di Arcetri, Firenze, Italy

Introduction: Generative Artificial Intelligence (GenAI) is rapidly transforming scientific workflows, yet empirical data on its adoption within national research infrastructures remain scarce.
Methods: This study presents the results of a comprehensive survey conducted at the Italian National Institute for Astrophysics (INAF) in spring 2025, collecting 371 responses (~20% of the workforce). Data were analyzed quantitatively, including user segmentation through K-Means clustering, while qualitative feedback highlighted user attitudes and ethical concerns.
Results: Our analysis reveals widespread adoption, with 50% of staff using GenAI weekly, primarily for coding and linguistic refinement. We identify a critical shift in usage patterns: AI is valued more for speed (efficiency) than for quality improvement, acting as an accelerator rather than a quality enhancer. The clustering analysis identifies three user personas: Enthusiasts, Pragmatists, and Skeptics, highlighting a hidden digital divide where 18% of researchers invest personal funds to access premium tools ("Shadow IT"), creating equity concerns. Furthermore, we find a correlation between AI literacy and user satisfaction, suggesting that skepticism may stem from lack of training.
Discussion: The study concludes that while the community is predominantly eager to adopt these tools, the current absence of formal guidelines poses risks regarding data privacy and scientific integrity. We provide evidence-based recommendations for institutional governance, prioritizing specific technical training and centralized provision of suitable tools over too restrictive policies.

KEYWORDS
AI governance, AI literacy, astrophysics community, generative artificial intelligence, large language models, research infrastructure, Shadow IT, AI concerns

1 Introduction

The last 4 years have witnessed exponential growth in the capabilities and adoption of Large Language Models (LLMs). With the release of tools such as ChatGPT in late 2022 (Ouyang et al., 2022; Brown et al., 2020), the global public gained access to artificial intelligence systems capable of generating text, code, and complex analyses with unprecedented fluency. In the scientific realm, and particularly within STEM disciplines, this technological revolution promises to radically transform workflows, offering new possibilities for ideation, literature review, code writing, and manuscript drafting. Recent studies estimate that a significant portion of STEM professionals' working hours could be automated or assisted by generative AI in the near future (Noy and Zhang, 2023).

However, the impact on productivity is not uniform. Recent literature describes a "Jagged Technological Frontier", where AI excels at specific tasks while failing at others that appear similar, creating a complex landscape where effectiveness depends heavily on user competence and the nature of the task (Dell'Acqua et al., 2024). While some experimental studies highlight an increase in productivity, particularly for profiles with lower initial skills (Noy and Zhang, 2023), others warn against the risks of uncritical adoption.

The astronomical community has responded rapidly to these innovations. Beyond the use of generalist tools, we are witnessing the development of domain-specific models, such as AstroLLaMA (Nguyen et al., 2023) or semantic frameworks like Pathfinder for knowledge discovery (Iyer et al., 2024). Recent analyses of the literature suggest that the use of tools like ChatGPT in drafting astronomical papers is already widespread and growing, often without explicit disclosure (Astarita et al., 2024). This raises urgent questions regarding transparency, and scientific integrity, especially considering the propensity of these models for "hallucinations" (Birhane et al., 2023) and the risk of generating "illusions of understanding" that could pollute the scientific record (Messeri and Crockett, 2024).

Despite the growing enthusiasm, a significant gap emerges in our practical understanding of the phenomenon. Although several institutions and journals have begun publishing general guidelines on the ethical and practical use of LLMs, solid empirical data on how the astronomical community is effectively integrating these tools into their daily work is still lacking. Existing literature has so far focused primarily on theoretical analyses or qualitative studies based on small groups of participants (Fouesneau et al., 2024; Sharma et al., 2025). For instance, Fouesneau et al. have explored the potential of LLMs through the experiences of restricted working groups of astronomers, highlighting promising use cases but without providing a statistically representative overview. The focus of Sharma et al. (2025) is on Indian research policy in astrophysics; the survey—three questions to twenty-nine respondents—is simply touched on, revealing anyway that the majority of the participants (58.6%) presently do not use AI/ML in mainstream astronomy research, but most of them (48.3%) plan to do so in the near future.

On the quantitative front, while we can build on disciplinary area results emerging from global surveys of scientific communities (Mohammadi et al., 2026; Oxford University Press, 2024; Van Noorden and Perkel, 2023), studies focused on the specificities of national research infrastructures in astrophysics are lacking. A recent exception is the large-scale study conducted within the Max Planck Society and the Fraunhofer Society in Germany, which highlighted how adoption varies significantly based on role and familiarity with the tool (Chugunova et al., 2026). Furthermore, new ethical concerns are emerging regarding the use of AI not only in the production but also in the evaluation of science, as highlighted by recent studies on the use of LLMs in peer review (Liang et al., 2024) and, more broadly, by Trotta (2025); these issues require internal institutional reflection.

This work aims to bridge these gaps by shifting the analysis from a qualitative or generic context to one that is also quantitative and specific to the Italian research infrastructure. We present here the results of a survey conducted among the staff of the National Institute for Astrophysics (INAF), collecting nearly 400 responses from researchers, technologists, administrative and technical staff at various career stages, as well as pre-tenured research collaborators. Unlike previous studies based on small groups of "early adopters", this paper offers an extensive and detailed "snapshot" of the state of the art in Italy. Through the analysis of this large dataset, we explore real adoption rates, prevalent use cases, community sentiment towards AI, and the perception of ethical and practical risks, thus also providing an empirical basis for future institutional policy decisions.

2 Materials and methods

To assess the penetration and impact of Generative AI tools within the Italian National Institute for Astrophysics (INAF), we designed and distributed an anonymous online survey. While previous exploratory studies in astronomy have provided valuable initial insights, they primarily relied on qualitative analysis of restricted working groups or limited survey samples—such as the 13-person focus group and 27 survey respondents analyzed in Fouesneau et al. (2024). Our goal was to complement these findings by obtaining a statistically significant representation of the national community, comparable in scope to recent large-scale studies conducted in other European research organizations (Chugunova et al., 2026).

2.1 Survey design and administration

The survey was conducted over a period of about 3 weeks, from 17 April to 9 May 2025. The questionnaire was distributed via institutional mailing lists to ensure it reached personnel across all INAF structures and across various professional areas. The inclusion of technical and, especially, administrative staff appears to be rare in previous surveys of this kind and is aimed at giving a complete picture of INAF as a multiform community.

The survey comprised 37 questions organized into five logical sections. To ensure data relevance, the questionnaire employed branching logic: specific questions regarding the impact of AI on tasks such as software development or academic writing were displayed only to respondents who explicitly indicated using AI for those purposes.

The sections were structured as follows:

1. Demographics and Professional Profile: Characterization of respondents by role, age group, contract type, and research area (RSN).
2. Adoption and Literacy: Investigation of frequency of use and specific tools employed. This section included a self-assessment of AI literacy using a qualitative scale to measure confidence in managing hallucinations and crafting effective prompts.
3. Task-Specific Applications: A deep dive into workflows for Scientific Research, Writing, Evaluation, and Coding. Respondents rated the perceived impact on work quality and efficiency using a 5-point Likert scale.
4. Concerns and Governance: Respondents' sentiment identification including limiting factors and opinions on future institutional policies.
5. Qualitative Feedback: Open-ended questions collecting real-world examples of, for example, AI failures and success stories.

2.2 Quantitative data analysis

Data processing was performed using Python, specifically relying on the Pandas (McKinney, 2010) and SciPy (Virtanen et al., 2020) libraries. A preliminary data cleaning phase was necessary to harmonize professional roles and normalize tool names. Specifically, unstructured free-text inputs identifying "Other" professional roles were manually mapped to standard institutional categories. Missing data within the dataset was primarily due to the survey's branching logic—where entire sections (e.g., specific questions on software development) were automatically skipped if not applicable to the respondent's profile—and the optional nature of open-ended qualitative fields. For core mandatory questions, the completion rate was near-complete, and any remaining missing values were handled using pairwise deletion to maximize the use of available responses.

For the user profiling, we employed Unsupervised Machine Learning via K-Means clustering (MacQueen, 1967) implemented in Scikit-learn (Pedregosa et al., 2011). The input variables (usage frequency, satisfaction index, literacy index, and personal financial investment) were first standardized using zero-mean and unit-variance scaling (StandardScaler). The optimal number of clusters ($K = 3$) was determined by evaluating two standard validation metrics: the Elbow Method (which visually identifies the point where adding more clusters no longer significantly reduces the within-cluster variance) and the Silhouette score (a metric ranging from −1 to +1 that measures how similar an object is to its own cluster compared to other clusters, indicating strong cohesion and separation). Both metrics indicated a robust segmentation at three clusters. The stability of the solution was verified across multiple random initializations (seeds).

2.3 Qualitative data analysis

To complement the quantitative metrics, the survey included open-ended questions designed to capture nuanced user experiences, ethical, legal, and epistemological concerns, as well as real-world examples of AI applications. The free-text responses, originally submitted in Italian, were subjected to an inductive thematic analysis.

The responses were reviewed to identify and code recurring macro-themes (e.g., "Shadow IT", "Linguistic Equity", "Fear of Deskilling"). To ensure reliability, the thematic coding and the selection of representative quotes were discussed and agreed upon by the research team, resolving any discrepancies through consensus. The selected quotes presented in the Results section were subsequently translated into English. The translations are verbatim, with only minor grammatical adjustments applied to ensure readability in English while strictly preserving the original meaning and tone of the respondents.

2.4 The INAF context and participant demographics

INAF operates as a highly distributed network comprising headquarters and 17 research structures across Italy. With a total of 371 valid responses, our survey achieved a penetration rate of approximately 20% of the institute's active workforce, consistent with similar European studies (Chugunova et al., 2026).

To validate the robustness of our data, we compared the demographic breakdown of the respondents against the official INAF population statistics. As illustrated in Figure 1, while permanent staff (Researchers, Technologists, and Administrative/Technical personnel) are well-represented, the comparison highlights an under-representation of early-career scientists (i.e., PhD students, postdoctoral researchers, research fellows, and grant holders; 3.7% in our sample, $N = 13$, versus 27.3% of the total INAF population). Although this deviation from the population baseline represents a limitation of our sampling, the massive response rate from the permanent staff provides a solid and reliable foundation for evaluating the institute's core operational workflows. Furthermore, the responses gathered from the early-career group, albeit smaller in number, remain crucial for identifying emerging usage trends among younger generations, who are typically early adopters of AI tools.

Furthermore, we verified the disciplinary coverage by mapping respondents to the five National Scientific Groupings (RSN). Figure 2 confirms that the survey represents a holistic cross-section of the institute. Notably, the higher engagement from RSN 5 (Advanced Technologies) reflects the specific relevance of AI tools in software-intensive domains, yet all major research areas (RSN 1–4) retain statistical significance.

3 Results

The survey reveals a widespread integration of Generative AI tools within the institute. In this section, we analyze adoption rates, tool preferences, and use cases, while also highlighting the barriers and perceived risks that currently limit deeper integration.

3.1 Adoption rates and frequency of use

The penetration of AI tools among INAF personnel is significant. 50% ($N = 186$) of respondents reported using GenAI tools at least

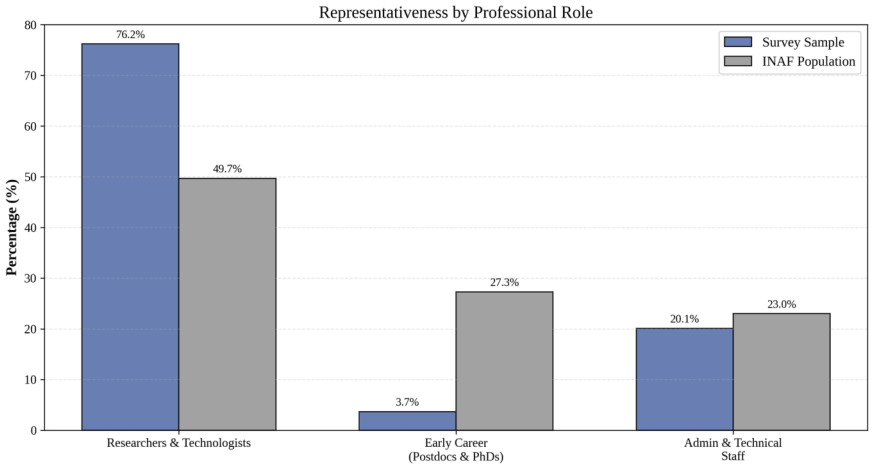

FIGURE 1
Sample Representativeness by Role. Comparison between the survey respondents and the total INAF population. The sample successfully captures a large portion of researchers, technologists, and administrative/technical personnel, while highlighting an under-representation of early-career researchers compared to the institutional baseline. Permanent staff is categorized into Researchers (primarily focused on scientific production), Technologists (focused on advanced instrumentation, software, and infrastructure), and Administrative/Technical personnel.

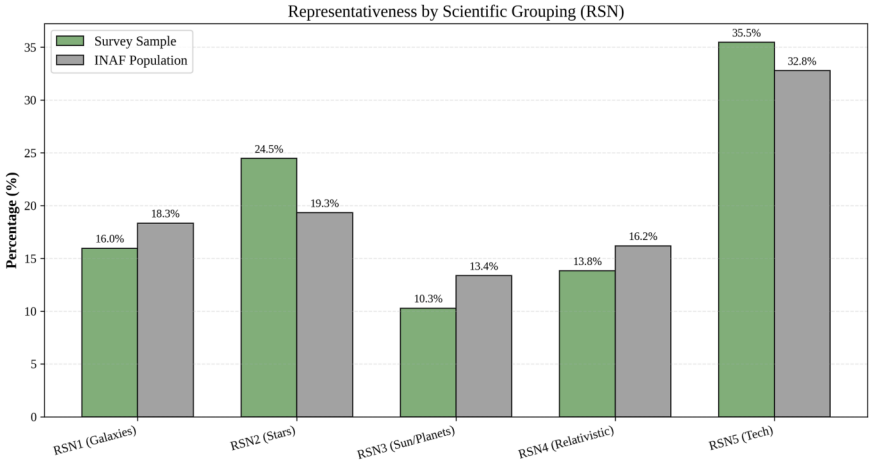

FIGURE 2
Scientific Distribution (RSN). The breakdown of respondents by research area mirrors the institute's broad scope. The slight peak in RSN5 (Advanced technologies) correlates with the high usage of AI for coding tasks identified in the results.

once a week, with a 17% ($N = 63$) reporting daily use. Only a minority (13.8%, $N = 51$) declared they never use such tools.

However, adoption is not uniform across the institute. As shown in Figure 3, there is a clear correlation between professional role and usage intensity. Early Career researchers display the highest density of "Daily" and "Weekly" usage, whereas the frequency tends to decrease slightly among permanent Staff and administrative roles.

3.2 The toolkit and the "Shadow IT" phenomenon

In terms of specific platforms, the landscape is dominated by OpenAI's solutions. However, Figure 4 reveals a critical detail regarding access and equity. While the majority of users rely on free versions (grey bars), a significant fraction of respondents—particularly for ChatGPT and GitHub Copilot—pay for subscriptions using personal funds (red segments). This points to a "Shadow IT" phenomenon where researchers subsidize their own work tools to access advanced features, lacking in institutional licenses.

3.3 Primary use cases: quality vs. speed

Respondents identified a versatile usage pattern. While coding remains a primary driver (45.8%, $N = 170$), linguistic support and writing assistance are nearly as prevalent (38.8%, $N = 144$).

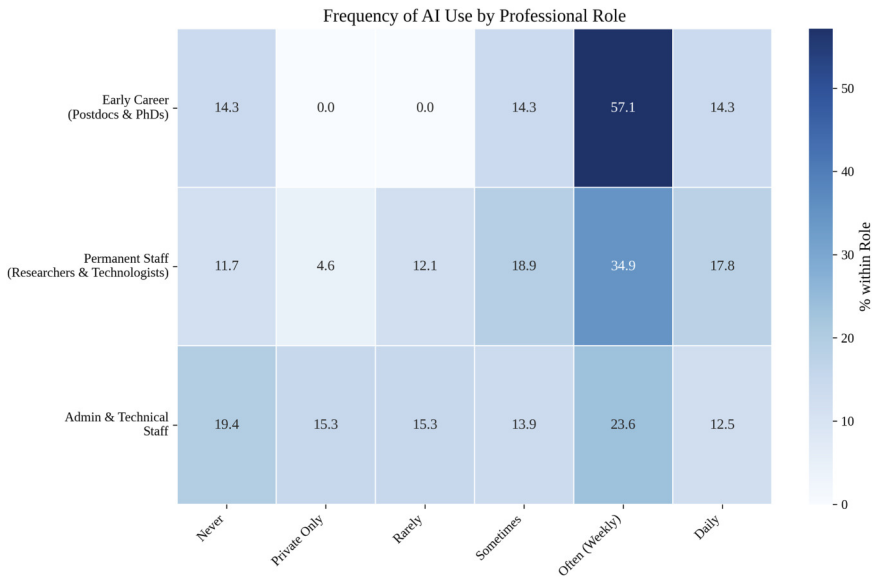

FIGURE 3
Frequency of AI Use by Professional Role. The heatmap displays the percentage of respondents within each role for every frequency category. Darker shades indicate higher concentrations. Notably, Early Career (N = 14) show an exceptional adoption rate (57,1 weekly usage, N = 7), while administrative staff tend to limit usage to private contexts or non-adoption.

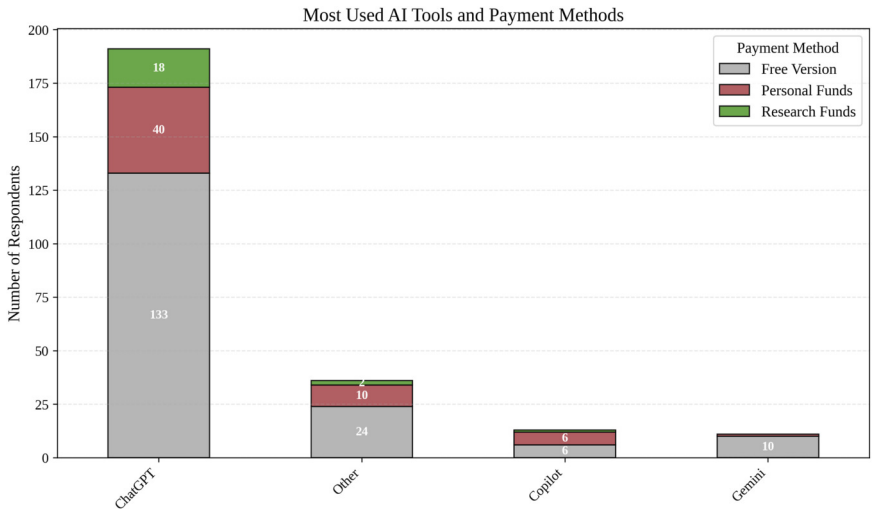

FIGURE 4
Most used AI tools and payment methods. Stacked bars show funding source by tool. ChatGPT dominates usage, while personal funding (red), especially for coding tools such as Copilot, indicates a gap in institutional provision. The "Other" category (about 2%) includes niche models such as *Anthropic Claude* (*N* = 4) and *Perplexity AI* (*N* = 2).

To understand the *nature* of this help, we analyzed the perceived impact on different tasks (Figure 5). Two key patterns emerge:

- The Efficiency Gap in Writing: For writing tasks, the perceived benefit is clearly skewed towards *Speed/Time Saving* (orange box, mean ≈ 3.2) rather than pure *Quality Improvement* (blue box, median ≈ 3.5). This confirms that AI is primarily valued as an accelerator of routine drafting rather than a qualitative enhancer of scientific content.
- Task Maturity: *Coding* displays the most compact and high-ranking distribution, indicating a mature and consistent utility. Conversely, *Research* shows the highest variance (wide whiskers), reflecting a polarized experience where the utility heavily depends on the specific query complexity.

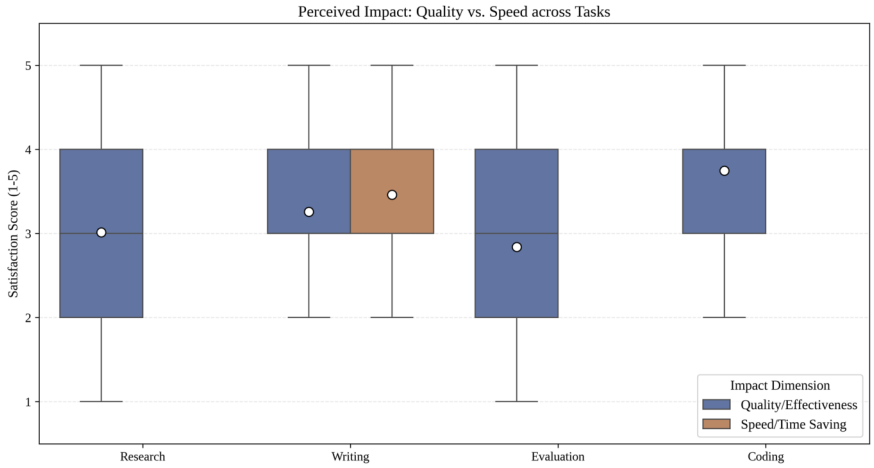

FIGURE 5
Perceived Impact: Quality vs. Speed. Distribution of satisfaction scores (1-5 scale). In each boxplot, the central horizontal line represents the median, while the white circle indicates the mean score. The colored box encloses the interquartile range (IQR, from the 25th to the 75th percentile), and the whiskers extend to represent the rest of the data distribution. Note: The 'Speed/Time Saving' dimension was specifically surveyed only for the academic writing task.

3.4 Barriers and the literacy gap

Despite high adoption, reliability remains the top concern (64.3%). A crucial finding is the link between user expertise and the perceived utility of the tools. To quantify this relationship, we calculated the Pearson correlation between a composite *Literacy Index* (derived from the mean of six self-assessment items) and an aggregate *Satisfaction Index* across all reported activities. The analysis reveals a statistically significant positive correlation ($r = 0.35$, $p < 0.001$).

This relationship is visualized in detail in Figure 6, which breaks down satisfaction levels by expertise group. The data suggests that literacy is strongly associated with perceived quality, particularly in complex domains. In *Software Development*, experts more frequently report maximum satisfaction (Score 5) compared to novices (31% vs. 14%). A similar trend is observed in *Research Problem Solving* (10.5% vs. 3.9%). This pattern suggests that advanced prompting skills and awareness of model limitations are key factors associated with achieving professional-grade results. However, being a cross-sectional observation, this correlation should be interpreted as an association rather than a strict causal mechanism.

3.5 User personas: the hidden digital divide

To move beyond standard demographic segmentation, we employed Unsupervised Machine Learning (K-Means clustering) as an exploratory tool to detect behavioral patterns based on usage intensity, satisfaction, AI literacy, and willingness to pay. The analysis ($N = 307$ active users) identified three distinct user segments (Figure 7), which we heuristically define as Enthusiasts, Pragmatists, and Skeptics.

- The Enthusiasts (18.6%, $N = 57$): This cluster represents the "power users." Visually, they cover the largest area of the radar chart. Their defining feature is the *Personal Investment* axis (normalized score = 1.00): they are the only group paying out-of-pocket to access premium models. This financial commitment correlates with the highest *Literacy Score* (0.58) and peak *Usage Intensity* (0.79), suggesting that advanced features (e.g., GPT-4) are necessary to fully satisfy expert requirements.
- The Pragmatists (49.2%, $N = 151$): This is the most critical finding of our study. As shown by the blue profile in Figure 7, this group is quantitatively almost identical to the Enthusiasts in terms of *Frequency* (0.75 vs. 0.79) and *Satisfaction* (0.70 vs. 0.71). They use AI daily and derive comparable value from it. However, they diverge sharply on the Investment axis (0.00 vs. 1.00): they rely strictly on free tools. This identifies a massive productive segment operating in a sub-optimal environment—highly active and skilled (*Literacy* = 0.49), but unsupported by institutional provision and unwilling to use personal funds.
- The Skeptics (32.2%, $N = 99$): The red profile occupies the center of the chart. Characterized by scores consistently below the 0.4 threshold, they report low *Frequency* (0.32) and the lowest *Satisfaction* (0.35). The compact shape of their cluster, combined with the lowest *Literacy Score* (0.24), indicates a self-reinforcing pattern: without sufficient theoretical knowledge to guide the AI effectively, the resulting output is often perceived as poor, which in turn sustains their initial skepticism.

The statistical overlap between *Enthusiasts* and *Pragmatists* on the "Utility" axes proves that cultural resistance is not the primary barrier for the majority of users (approx. 68% of the active workforce, $N = 208$). Since Pragmatists are already highly active and satisfied using free tools, they are not strictly "held back" in their daily workflows. Providing enterprise-grade licenses would bring this massive cohort of independent users into a safe, privacy-compliant environment, mitigating the risks of unmanaged AI adoption. Conversely, for the *Skeptics*, the data suggests that targeted training programs aimed at improving AI literacy could

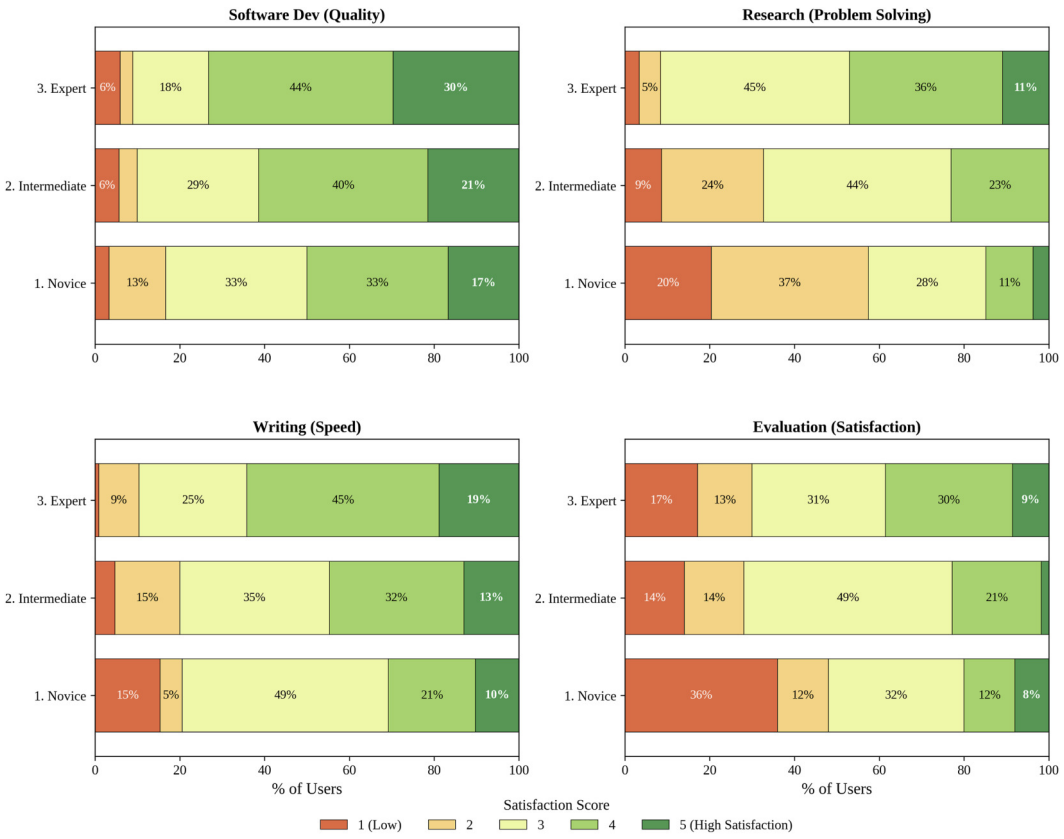

FIGURE 6
Impact of AI Literacy on Satisfaction across Tasks. The charts display the distribution of satisfaction scores (1–5) for Novice, Intermediate, and Expert users across four domains. While the statistical correlation between overall literacy and satisfaction is moderate ($r = 0.35$), the breakdown reveals a strong "Expert Premium" in technical tasks: experts are twice as likely to report maximum satisfaction in Coding and Research compared to novices.

address this pattern of low competence and low satisfaction, progressively increasing their engagement and trust in these technologies.

3.6 Ethical stance and transparency

A central aspect of the survey concerned ethics. Ethical concerns are clearly represented within our sample. Respondents had the opportunity to express them through two multiple-answer questions:

1. "*What aspects raise your concern about the use of generative AI at INAF?*", with multiple ethical options allowed and depicted in Figure 8
2. "*Which factors limit your use of AI tools?*" — including the option "*I am concerned by the ethical aspects (e.g., risk of plagiarism, unethical use of the generated content etc.)*".

The first question was answered by 91.64% of the total survey respondents (N = 340 out of 371), with an average of ~4 answers per respondent. At least one reason for concern—of any kind—was expressed by as many as 94.11% of the researchers who answered the questionnaire, followed by non-tenured research collaborators (research assistants, PhD students and post-doc students: 92.30%), technologists (~92%), administrative staff (~81%) and technicians (~77%)[1]. This scenario reveals a clear need for information, training and, possibly, community debate. However, there appears to be no correlation between the presence of ethical concerns and self-reported levels of AI literacy.

The answering options to this question were intentionally included rather broadly, so that concerns could be expressed both with regard to proper ethical issues, and to the legal sphere. The outcome reveals that the carbon footprint of AI, the possible substitution of the human workforce by AI and the likely progressive deskilling of human authors due to AI use—which can be considered as the "core" ethical answers—altogether account for ~25% of all answers. 61.23% of the total respondents ($N = 208$) to the survey, across areas, chose at least one of the core ethical answers, with non-tenured researchers ahead (69.23% within the area) followed by researchers (~66%) and technologists (~59%).

Deskilling emerges as the second most representative answer to the whole question 1 (13.45%, $N \approx 183$ answers), in line

1 These percentages of areas have been calculated excluding the five respondents who gave only free answers. After managing these cases for as much as possible (three out of five cases), the percentages become: researchers: ~95%; technologists: ~93%; non-tenured research collaborators: ~92%; administrative staff: ~81%; technicians: ~77%.

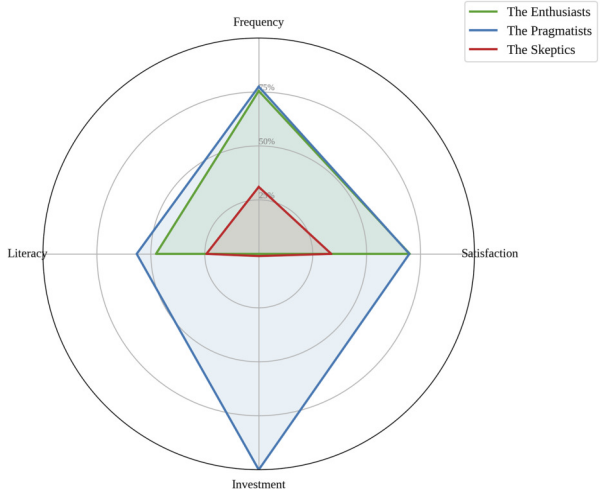

FIGURE 7
User Personas Radar Chart. Three distinct user clusters identified via K-Means. The axes represent normalized scores. Notably, "Enthusiasts" (green) and "Pragmatists" (blue) overlap on usage and satisfaction, identifying a large productive workforce that diverges solely on the "Personal Funds" axis, highlighting the demand for institutional licenses.

with the results of recent literature (Chugunova et al., 2026; Trotta, 2025; Oxford University Press, 2024)[2], and revealing concern for a possible change of paradigm in the performance of intellectual tasks, in the context of a new human-machine balance.

Some significant open answers reveal the concern that, ultimately, an uncritical use of AI may challenge the very nature of the scientific endeavour. As highlighted by several respondents:

> "When all papers are written by the AI, there will be no more science, but the 'absolute truth' of AI. Today, for the single researcher, it looks like a miracle: it's like having dozens of PhDs in support. But it's an illusion: as soon as the tool is within everyone's reach, it will become indispensable for remaining competitive, only accelerating our hamster wheel. The result? Not discoveries, but more and more complex combinations of what we already know. And with this, science will die."

> "[I am concerned by] The trivialization of the creative process of the human mind: which is not a simple collection and stitching together of content drawn from the internet and linked by analogy, but the result of patient meditation and connection of one's own rational and emotional experiences."

> "I think that in the long run [the AI tools] could actually worsen the quality of research, understood as the processing of thoughts, ideas, hypotheses."

> "The methodology at the basis [of AI] is based upon a statistical approach which is fundamentally different from the logical process in place during a scientific activity. I think [AI] can be useful, but only for managing large quantities of data. I think today it is used inappropriately, without adequate knowledge."

> "We must not look for solutions that speed up work but solutions that improve and enhance the acquisition of work knowledge by humans not by an algorithm."

> "The AI can be an accelerator, but we must prevent it from becoming the pilot. If we just follow it, maybe we go faster, but only where it allows us to go."

If considered in broader terms, the ethical scenario appears to be more fragmented, probably as it captures different sets of values and beliefs more easily. While a respondent states:

> *"I do not believe ethical issues should be considered, as ethics is a para-religious framework."*

another respondent thinks:

> *"I would like it if, before spending public money on proprietary product licenses that give an advantage to a few, we first considered ethical and moral issues. Those who want to use AI do so privately and not academically. It's a Trojan horse that, once it's let in, we will not be able to drive out."*

For another respondent:

> *"In my opinion, we should start to think about AI also in philosophical terms, not only technically and ethically."*

To these three 'core ethical' options, we deem it important to add what we recognize as a specific subset in the ethical area, which calls into question research, or professional, integrity: "*lack of attention in using AI answers due to their inaccuracy and presence of hallucinations, biases,* etc.," and "*misrepresentation of AI-generated content as human-created content*". The former is the most concerning issue for all respondents to question 1 (17.50% of answers, $N \approx 238$), showing how important it is for the community to be able to rely on robust professional contributions from peers; the latter highlights the importance of networks of mutual trust. This subset adds another ~30% ($N \approx 408$ answers) to the total, thus making it clear that ethical concerns do impact the community significantly, although in different nuances.

Two legal concerns about intellectual property rights ("*the generated content might contain copyrighted materials, without…warning*" and "*the uncertainty about the attribution of intellectual property to the generated content*") account for an

2 In (Mohammadi et al., 2026) this concern ranks third (p.9) and the whole set of concerns as well as its internal ranking is largely comparable to our findings. In the Oxford U.P. Report (Oxford University Press, 2024), "the top theme that most concerns researchers is that AI may lead to a loss of critical thinking" (p. 12).

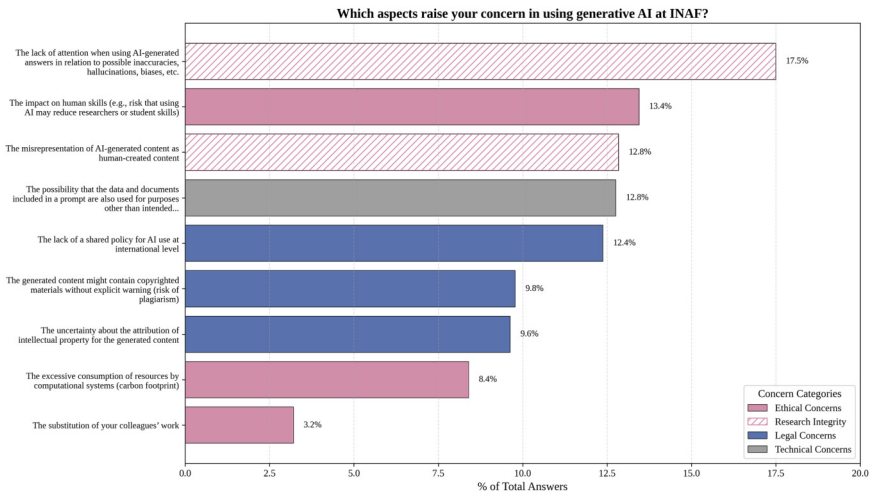

FIGURE 8
Analysis of the answers to question 1, "*Which aspects raise your concern in using generative AI at INAF?*" (%). Colour pink indicates ethical concerns; colour blue is for legal concerns; grey indicates a technical concern. Diagonal pink lines on a white background is for worries concerning research integrity.

additional ~19% ($N \approx 258$ answers). The further choice "*lack of a shared policy at international level*" (12.38%, $N \approx 168$ answers) suggests concern about the possibly debatable use of AI on a global scale, as a shared normative framework on AI has not been agreed on yet—apart from the pioneering actions of the European Union both for general scope norms (AI Act, 2024) and in form of non-binding guidelines for research (2025)[3] (thus, some international publishers' policies on AI seem to presently act as *de facto* worldwide references for astrophysics).

3.7 Ethical concerns as a possible limitation to the use of AI

Ethical concerns emerge also as possible limitations to the respondents' use of AI (question 2). This occurs to a lesser extent compared to the more general question 1. The total percentage of ethically concerned personnel who might be limited in their use of AI is ~18% of respondents to this specific question (N = 64 out of 352). Those who take this option are found mostly in the age tier 46–55 y (~35%, $N = 22$), followed by the youngest researchers (≤ 35 y.o.) and by those older than 55 y (22.58% each, $N = 14$ each). Researchers in the 36–45 age tier stop at ~19% ($N = 12$), although one of them writes overtly "*I'm ideologically against*".

The percentages on the actual members of the categories show that technologists are proportionately the most concerned category (38.46%), followed by the other groups as shown in Figure 9.

Continuing to use the percentage on the total number of respondents in each category, we find that the most concerned research sub-areas are "*Sun and solar system*" (~31%) and "*Stars, stellar populations and interstellar medium*" (~30%), whereas the advanced technologies area exhibits a limited concern (~13%).

As in question 1, we found a more limited correlation between self-reported limitations to use due to ethical concerns and the self-perception of inadequate AI literacy. ~61% ($N = 39$) of these ethically concerned respondents feel "*proficient*" or "*rather proficient*" in AI usage, indicating that their ethical concerns are unlikely to stem from a perceived lack of technical competence.

Interestingly, the correlation between reported ethical constraints and actual non-use of AI tools, while not negligible, is not as high as expected. In fact, "*I've never used it*" and "*I use it only in private life*" altogether represent ~33% ($N = 21$) of these respondents. Instead, as many as ~31% ($N = 20$) use it "*often*" and ~8% ($N = 5$) choose "*always*". A gray area of infrequent use lies in between.

Of special interest to us are the reasons for this behaviour. Noting that multiple answers were allowed, we find that the great majority of answers (36) are "*[AI] makes some activities faster*". Only 16 answers are "*[AI] helps improving the quality of my job*", with 6 more answers divided between "*it allows me to do things that I would not be able to do otherwise*" and "*I'd be disadvantaged compared to those who use it*". Once again, the most substantial advantage of AI is perceived to be quickness more than quality, which makes it strategic to address the situation. It could be argued that, in a competitive, internationally intertwined domain such as astrophysics, the possibility to boost one's productivity may overcome the awareness of using a somehow still debated, and/or still inadequately mastered, tool.

Curiosity seems to be another driver for some respondents' experimentation with AI. As highlighted by two different respondents:

> "I think I've never used it, but I'm curious about it."

3 This critical issue is pointed out by (Chugunova et al., 2026), the general-scope EU AI Act (https://aiactinfo.eu/, published July 2024) is being enforced gradually. The abridged "Living guidelines on the responsible use of generative AI in research" provide clear advice to researchers and research institutions, https://research-and-innovation.ec.europa.eu/document/2b6cf7e5-36ac-41cb-aab5-0d32050143dc_en

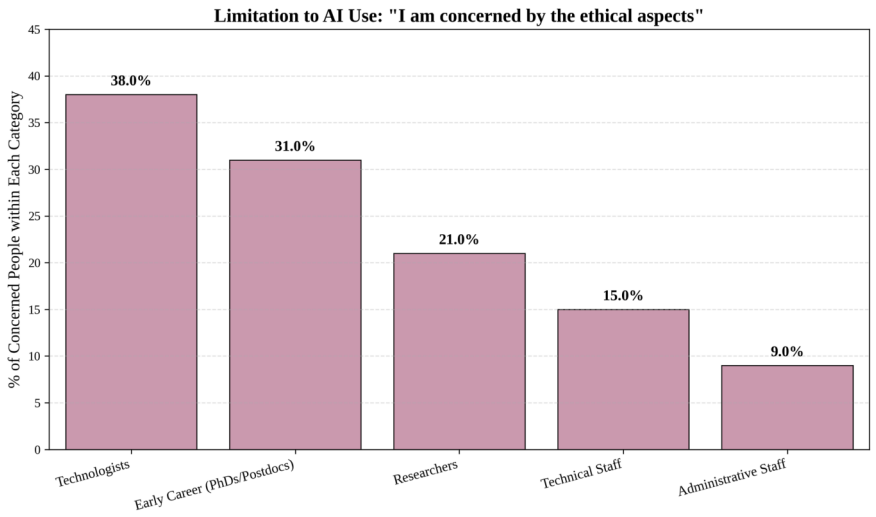

FIGURE 9
Analysis of the answers to question 2, *"What factors limit or reduce your use of generative AI tools?"* – Option *"I am concerned by the ethical aspects"* (% of concerned people in each category). Please note: these percentages are higher in question 1.

> "I'm satisfied by the free tools, but I'd be curious to test a paid alternative."

When asked whether AI usage should be explicitly declared in scientific publications, responses revealed a complex landscape. As shown in Figure 10, the general survey population displays a polarized stance: only 18% ($N = 67$) advocate for unconditional transparency ("Yes, Always"), while a significant portion aligns with the pragmatic view of declaring usage only if required by policy. This reliance on external mandates is particularly critical given that recent large-scale studies (He and Bu, 2026) have demonstrated a stark transparency gap: despite the widespread adoption of AI-use policies by academic journals, actual disclosure remains virtually nonexistent (~0.1%), suggesting that declarative policies alone fail to curb undisclosed AI assistance.

Interestingly, the subset of respondents who previously expressed ethical limitations shows a much higher openness to transparency: within this specific group, unconditional transparency ("always") reaches ~38% ($N = 24$), and disclosure in case the AI contributed significantly accounts for ~34% ($N = 22$).

Furthermore, researchers and technologists show a slightly higher propensity for conditional declaration based on contribution significance (~34%), contradicting the hypothesis of a simple linear correlation between age and ethical openness. The data suggests that across all roles, the community is split between "normative compliance" and "scientific transparency".

3.8 Institutional roadmap and future outlook

Finally, we asked respondents to identify priority actions for the institute. The community is oriented towards the future: 57% ($N = 211$) of respondents expect to use AI more frequently in the next 2 years.

As illustrated in Figure 11, the requests for institutional intervention are multifaceted but strongly prioritize practical enablement over mere regulation. The dominant need is educational: 73.0% ($N = 270$) of respondents identified "Specific Training" as the most urgent intervention, outpacing the demand for general information seminars by a factor of two. This indicates that the majority of INAF personnel are eager to move beyond the experimental phase and are seeking support, including structured technical upskilling (e.g., prompt engineering, RAG workflows).

Finally, the community calls for clear guidelines to navigate the ethical and legal ambiguities of AI usage, particularly for administrative tasks and scientific evaluation. Overall, the data provides a clear message: the majority of the staff is ready to fully integrate GenAI, provided they are equipped with the right training, secure tools, and unambiguous policies.

3.9 In-depth analysis of use cases

While the previous section outlined the quantitative landscape of adoption, the qualitative feedback from the survey offers a nuanced view of how these tools are actually reshaping the daily workflows at INAF. To illustrate the variety of sentiments and applications within the community, Table 1 presents a selection of representative quotes collected from the open-ended responses. In the following subsections, we analyze the four use cases—writing, coding, scientific support/brainstorming, and administration—highlighting the specific benefits and risks reported by the staff.

3.9.1 Academic writing

Writing is the primary vehicle for scientific dissemination, yet for non-native English speakers, it often imposes a significant cognitive load. While Fouesneau et al. (2024) identified writing assistance as a secondary use case (72%), our survey reveals

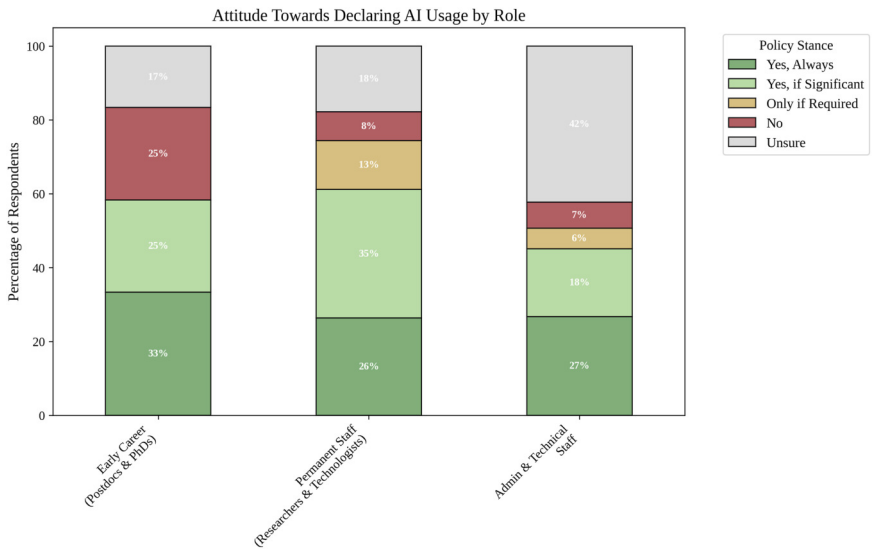

FIGURE 10
Attitude Towards Declaring AI Usage by Role. Stacked bar chart showing the ethical stance across different professional categories. The "Early Career" group (PhDs and Postdocs) shows a highly diversified approach to transparency. Notably, compliance with external editorial or institutional policies remains a strong driver across all demographics, rather than unconditional proactive disclosure.

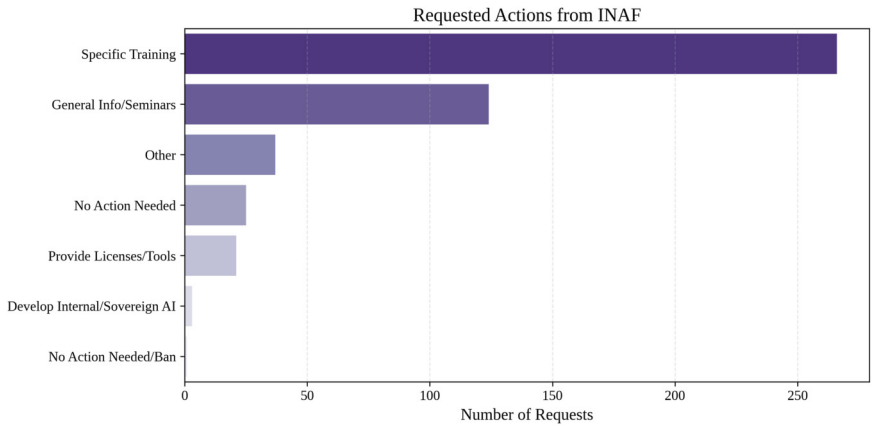

FIGURE 11
Requested Actions from INAF. Respondents could select multiple options. The overwhelming demand for "Specific Training" (requested by ~270 users, or 73%) highlights a critical skills gap that cannot be addressed by simple policy guidelines alone. Alongside education, the need for enterprise licenses, privacy-compliant tools, and clear policy frameworks reflects a mature demand for a secure AI ecosystem within the institute.

that for the INAF community, linguistic support is paramount. Approximately 38.8% ($N = 144$) of respondents utilize AI tools explicitly for writing tasks.

Crucially, our quantitative data (see Figure 5) clarifies the nature of this support: respondents rated *"Time Saving"* significantly higher than *"Quality Improvement"*. This suggests that the AI is not replacing the author's voice but rather acting as a *high-speed editor*.

Our data shows that only a small minority (~9%, $N = 33$) relies on GenAI to generate initial drafts from scratch. Instead, the dominant applications are linguistic verification (21.6%, $N = 80$) and stylistic refinement (17.2%). This trend is corroborated by the widespread adoption of DeepL alongside standard LLMs like ChatGPT. As noted by one respondent, these tools allow researchers to "correct and improve English starting from a text I wrote," effectively acting as a real-time proofreader. This usage pattern implies that AI acts as a "great equalizer" (Fouesneau et al., 2024), lowering the barrier for Italian researchers to produce publication-quality manuscripts without outsourcing the core scientific narrative.

However, trust remains a critical issue. Qualitative feedback highlights a persistent fear of flattening the scientific style or introducing subtle errors. As one participant noted, *"I only use it when forced …I refuse to train algorithms to do my work"* reflecting a tension between the utility of the tool and the preservation of intellectual authorship.

TABLE 1 Selected qualitative feedback from INAF staff representing key themes identified in the open-ended responses.

Theme	Representative quote (original user feedback)
Linguistic equity	"It is very useful for reviewing texts written in English, suggesting synonyms, and checking grammar. It speeds up the writing process significantly."
Coding efficiency	"I asked for a Python script to plot a specific dataset …it wrote 90% of the code correctly in seconds, leaving me only the refinement work. It acts like a junior programmer."
Shadow IT/Cost	"I personally pay for the subscription because the free models are not capable of complex reasoning. Institutional access to advanced models (e.g., GPT-4) is necessary."
Skepticism/Risk	"I asked for a bibliography on planetary habitability. It produced a list of plausible-sounding but completely non-existent papers. It invented the citations."
Administrative tasks	"I asked it to draft a formal bureaucratic email regarding a procurement procedure. It set it up flawlessly, even including the correct regulatory references."
Brainstorming	"I use it to clarify my ideas. Even if the answer is wrong, the dialogue forces me to better structure my own reasoning. It's useful for finding connections I had not thought of."

3.9.2 Software development

In modern astrophysics, coding is as essential as observing. Consistent with international trends, software development represents a major driver for AI adoption at INAF, with 45.8% ($N = 170$) of respondents indicating it as a primary activity. Tools like GitHub Copilot and ChatGPT are employed not just to write code, but to debug, optimize, and translate between languages (e.g., converting legacy IDL scripts to Python).

The perceived value of this application is exceptionally high. As shown in Figure 4, coding is the domain where researchers are most willing to invest personal funds to access premium tools ("Shadow IT"), underlining that for software tasks, the efficiency gain outweighs the cost.

The qualitative feedback offers a nuanced view of this acceleration. Respondents report that AI is particularly effective for generating boilerplate code (standard structures) or specific kernels (e.g., *"writing an RK45 integrator with variable step size"*). However, the human in the loop remains indispensable. Experienced developers noted that while AI enormously speeds up code writing, it does not necessarily improve the quality of the software architecture. In some cases, it can introduce subtle logical bugs ("hallucinations in code") that require expert supervision to detect. This reinforces the finding that AI acts as a junior pair programmer rather than a replacement for senior engineering skills.

3.9.3 Support to scientific activities

Beyond the execution of technical tasks like coding or editing, a significant portion of the community (~26%, $N = 96$) employs GenAI for a strictly cognitive purpose: *ideation and brainstorming*.

In this context, the AI is not asked to produce a final output, but to act as a "sounding board" or a conversational partner. Respondents described using these tools to "clarify confused ideas, explore alternative approaches to a problem," or "find connections between different topics." Crucially, this use case appears less prone to the "hallucination anxiety" described earlier. When used for brainstorming, the user does not expect factual precision but rather lateral thinking. As one researcher aptly summarized in the qualitative feedback: *"I use it to have a dialogue. Even if it says something wrong, it forces me to clarify my own reasoning to correct it. It's like talking to a rubber duck that talks back."*

This suggests that for a quarter of the respondents, GenAI has already transcended the role of a tool and entered the domain of a (synthetic) collaborator, supporting the creative phase of the scientific method itself.

3.9.4 Administrative and institutional support

A unique finding of our institute-wide survey, often overlooked in purely academic studies, is the application of GenAI to administrative and bureaucratic tasks. While the Administrative staff (PTA) shows a polarized adoption pattern (Figure 3)—with many non-users but a highly active core—those who do adopt AI report significant benefits in navigating institutional complexity.

Respondents described using AI to draft formal bureaucratic emails, summarize long technical reports, or verify regulatory references (e.g., procurement codes). One respondent explicitly mentioned success in asking the AI to *"formulate a formal email …which was set up flawlessly with proper regulatory references."* This administrative optimization represents a significant efficiency gain for a public research institute, potentially freeing up time for core research or management activities.

However, this domain is also where the "hallucination" risk is most treacherous. Unlike a coding bug which causes a crash, an invented law or regulation (a specific type of hallucination reported in our open-ended answers) could have serious legal repercussions. This vulnerability is exacerbated by the fact that LLMs struggle to keep pace with the rapid evolution of legislation. Unlike scientific principles or programming syntax, regulations change frequently, and models limited by training cut-offs often fail to retrieve the most current administrative provisions. This underscores the urgent need for institutional guidelines specifically tailored to the drafting of official administrative documents.

4 Discussion

The results of our survey paint a picture of a community at a pivotal moment. Far from being resistant to change, INAF personnel are rapidly adopting Generative AI, yet they are doing so in the absence of formal guidelines and with a healthy dose of scientific skepticism. In this section, we interpret these findings, highlighting the unique role of AI in non-Anglophone research environments,

the equity risks posed by "Shadow IT", and the urgent need for institutional governance.

4.1 Bridging the language gap: AI as an efficiency multiplier

While previous studies focused heavily on coding as the primary application of GenAI, in our sample, language refinement is the most used function. This is in line with interdisciplinary research at the global level (Mohammadi et al., 2026; Oxford University Press, 2024; Van Noorden and Perkel, 2023). In an international landscape where English is the lingua franca, non-native speakers often face a "linguistic tax" (Amano et al., 2023; Hosseini et al., 2023). The widespread adoption of tools like DeepL and ChatGPT acts as a powerful equalizer. However, our impact analysis (Figure 5) adds a crucial nuance: the perceived benefit is skewed towards speed rather than quality. AI is not necessarily making the science "better", but it is drastically reducing the time required to produce publication-quality manuscripts, effectively boosting the institute's productivity.

4.2 The equity dilemma: Addressing "Shadow IT"

A concerning trend emerges from the intersection of our "User Personas" analysis (Figure 7) and payment habits (Figure 4). We identified a cluster of "Enthusiasts" (~18%) who are not only heavy users but are also willing to pay personally for premium access (e.g., GPT-4, GitHub Copilot). This creates a potential internal digital divide (Vipra and Korinek, 2023): researchers with personal disposable income or specific grants have access to superior reasoning engines and coding assistants, while the majority ("The Pragmatists", 49%) relies on free, less capable models. For a public research body, this creates an equity issue. Institutional licensing is no longer just a matter of convenience, but a necessary step to ensure that all researchers, regardless of career stage or personal funds, have equal access to state-of-the-art computational support.

4.3 The "Skeptical Adopter" and the literacy barrier

Our survey uncovers an interesting paradox: adoption is high, but trust is low (64% fear hallucinations). While this skepticism is a positive signal of a "human-in-the-loop" mindset, our correlation analysis (Figure 6) suggests a deeper structural issue. We found a significant positive correlation between AI literacy and satisfaction. This implies that much of the reported frustration or skepticism might stem not from the tools' inherent limitations, but from a lack of advanced prompting skills ("The Knowledge Gap") (Dell'Acqua et al., 2024). Therefore, the "Skeptical Adopter" is often simply an untrained adopter. The path to unlocking value lies in moving beyond generic skepticism toward informed competence, as validated by the overwhelming demand for specific training.

4.4 Navigating the policy vacuum and ethical complexities

Perhaps the most critical finding for governance is the gap between usage and safety. The risk of data leakage is tangible, with only 18% of users confident in identifying safe contexts.

This vulnerability is structurally exacerbated by the reliance on free-tier accounts. Unlike Enterprise or API-based solutions, the Terms of Service for consumer-grade models typically grant providers the right to utilize user inputs for model training. With nearly half of the workforce relying on free tools ("The Pragmatists"), there is a concrete risk that unpublished scientific results, proprietary code, or sensitive administrative data are being inadvertently ingested into public models. This turns the institute's intellectual property into a "silent leak," effectively feeding public datasets with internal knowledge.

Regarding transparency, our data reveals a more nuanced reality than a simple "generational divide". While we hypothesized that younger scientists would champion open disclosure, the survey shows that both Early Career researchers and senior Staff are navigating a gray area, often deferring to external mandates ("If required by journal policy") (Stokel-Walker, 2023; Van Noorden and Perkel, 2023). This suggests that the "policy vacuum" makes it more difficult to choose ethical options: without clear institutional guidelines, researchers of all ages are defaulting to minimal compliance rather than proactive transparency. Any future policy must therefore be prescriptive, clarifying exactly when and how AI contribution should be acknowledged to remove this uncertainty.

4.5 Limitations of the study

We acknowledge that online surveys inherently suffer from self-selection bias; it is likely that respondents were already interested in AI topics. Therefore, the adoption rates might over-represent the "tech-savvy" portion of the institute. Furthermore, as shown in the demographic comparison (Figure 1), early-career scientists (3.7% of our sample vs. 27.3% of the population) are under-represented. Consequently, subgroup percentages regarding this specific demographic (such as their high frequent/weekly usage rate reported in Figure 3) must be interpreted with caution due to the small sample size ($N = 13$), and should be viewed as exploratory indications rather than robust population-level findings. Nevertheless, given the high overall response rate (20% of the workforce) and the broad coverage of all Research Groups (RSN) among permanent and technical staff, these trends provide a robust baseline for understanding the institute's core operational workflows.

4.6 Conclusions

This study provides the first quantitative mapping of Generative AI adoption within the Italian astrophysical community. The picture that emerges is one of a "pragmatic revolution": INAF personnel are not waiting for centralized guidance or official infrastructure to integrate AI into their workflows. They are using it to write code faster, to bridge the linguistic gap in English writing, and to navigate administrative complexities.

Our findings support three main directives for future institutional strategy:

1. Recognize the "Shadow Workforce": The existence of a large cluster of users (The Pragmatists) who rely on free tools proves that the demand is mature. The institute should consider providing access to high-performance models that guarantee data confidentiality. This would bring this workforce into a secure, compliant environment, mitigating the equity gap created by personal subscriptions and preventing the leakage of sensitive internal data into public training sets.
2. Invest in Literacy, Not Just Policy: The correlation between literacy and satisfaction proves that "knowing how to prompt" is a crucial skill. Training programs should mainly cover practical AI usage skills (advanced prompting, RAG, code debugging) and, crucially, foster awareness regarding data privacy. Users must understand that utilizing institutional or enterprise-grade models is not just a perk, but a necessary measure to protect their intellectual property compared to consumer-grade free tiers.
3. Differentiate Governance: Policies cannot be one-size-fits-all. They must accommodate the push for transparency from early-career scientists while addressing the compliance concerns of senior staff.

Encouragingly, INAF has already begun translating these directives into concrete actions. At the time of the survey, INAF already provided its staff with data-protected access to enterprise tools such as Google Gemini. However, our qualitative data revealed a widespread lack of awareness, leading many users to overlook these secure channels and rely on unverified free-tier models. To address this paradox and mitigate the "Shadow IT" phenomenon, the institute recently expanded its enterprise-grade licenses, guaranteeing strict data confidentiality and privacy-compliant AI assistance for a wider portion of its staff.

Furthermore, to bridge the literacy gap, the institution launched a series of targeted technical and ethical seminars, including specialized sessions for administration and public outreach, as well as advanced training on AI Agents and Retrieval-Augmented Generation (RAG). Looking ahead, INAF is actively networking to participate in GenAI-driven research projects and has initiated strategic dialogues with the CINECA national supercomputing consortium to explore the hosting of local, privacy-compliant LLMs. Ultimately, the challenge for research institutions is no longer whether to adopt AI, but how to proactively govern and support its use to ensure it remains a force for equity, security, and scientific excellence.

Data availability statement

The datasets presented in this article are not readily available. To protect participant confidentiality and mitigate the risk of re-identification associated with free-text responses and demographic intersections, the raw qualitative dataset will not be made publicly available. However, the fully de-identified and aggregated quantitative datasets generated and analyzed during the current study are available from the corresponding author upon reasonable request.

Ethics statement

Ethical review and approval were waived for this study, as it consisted of a voluntary survey administered to the staff of the Italian National Institute for Astrophysics (INAF) regarding their professional use of software tools. The survey did not involve patients, clinical data, or vulnerable populations. Informed consent was obtained from all participants prior to starting the questionnaire, with an explicit disclaimer that the institutional platform automatically collected email addresses. However, to guarantee strict privacy and confidentiality in accordance with the European General Data Protection Regulation (GDPR), all records were immediately de-identified. The dataset was completely stripped of personally identifiable information (PII) before any data processing took place, ensuring that the authors responsible for the data analysis and manuscript preparation had no knowledge of the respondents' identities. All data were subsequently analyzed and reported solely in an aggregated format. The study does not contain identifiable human images. The studies were conducted in accordance with the local legislation and institutional requirements. The participants provided their written informed consent to participate in this study.

Author contributions

AC: Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review and editing. MM: Investigation, Visualization, Writing – original draft, Writing – review and editing. GC: Conceptualization, Project administration, Supervision, Writing – original draft, Writing – review and editing. AP: Writing – original draft, Writing – review and editing. AB: Writing – original draft, Writing – review and editing. UL: Writing – original draft, Writing – review and editing. SR: Writing – original draft, Writing – review and editing. RS: Writing – original draft, Writing – review and editing. CA: Writing – original draft, Writing – review and editing. SC: Writing – original draft, Writing – review and editing.

Funding

The author(s) declared that financial support was received for this work and/or its publication. The publication fees for this manuscript were entirely funded by the Central Scientific Unit for Computing (Unità Scientifica Centrale per il Calcolo–USC-C) of the Italian National Institute for Astrophysics (INAF).

Acknowledgements

The authors wish to express their deepest gratitude to all the INAF researchers, technologists, technicians, and administrative

staff who took the time to participate in the survey. Their thoughtful responses, candid perspectives, and detailed qualitative feedback provided the foundational data for this study and offered an invaluable window into the real-world impact of Generative AI within our scientific community. Declaration of Generative AI and AI-assisted technologies in the writing process.

Conflict of interest

The author(s) declared that this work was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declared that generative AI was used in the creation of this manuscript. During the preparation of this manuscript, the authors used Google Gemini (model version March 2026, Google, https://gemini.google.com) primarily to assist with English language editing, stylistic refinement, and structural formatting of the text. After using this tool, the authors thoroughly reviewed and edited the content, checked it for factual accuracy and plagiarism, and take full responsibility for the final content of the publication. The Generative AI is not listed as an author.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/fspas.2026.1829012/full#supplementary-material

References

Amano, T., Ramírez-Castañeda, V., Berdejo-Espinola, V., Borokini, I., Chowdhury, S., Golivets, M., et al. (2023). The manifold costs of being a non-native english speaker in science. *PLOS Biol.* 21, e3002184. doi:10.1371/journal.pbio.3002184

Astarita, S., Kruk, S., Reerink, J., and Gómez, P. (2024). Delving into the utilisation of chatgpt in scientific publications in astronomy. doi:10.5281/zenodo.13885577

Birhane, A., Kasirzadeh, A., Leslie, D., and Wachter, S. (2023). Science in the age of large language models. *Nat. Rev. Phys.* 5, 277–280. doi:10.1038/s42254-023-00581-4

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., et al. (2020). Language models are few-shot learners. *Adv. Neural Inf. Process. Syst.* 33, 1877–1901.

Chugunova, M., Harhoff, D., Hölzle, K., Kaschub, V., Malagimani, S., Morgalla, U., et al. (2026). Who uses AI in research, and for what? Large-scale survey evidence from Germany. *Res. Policy* 55, 105381. doi:10.1016/j.respol.2025.105381

Dell'Acqua, F., McFowland, E., Mollick, E., Lifshitz-Assaf, H., Kellogg, K. C., Rajendran, S., et al. (2024). Navigating the jagged technological frontier: field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Organ. Sci. Forthcom.* doi:10.2139/ssrn.4573321

Fouesneau, M., Momcheva, I. G., Chadayammuri, U., Demianenko, M., Dumont, A., Hviding, R. E., et al. (2024). *What is the role of large language models in the evolution of astronomy research?*

He, Y., and Bu, Y. (2026). Academic journals' ai policies fail to curb the surge in ai-assisted academic writing. *Proc. Natl. Acad. Sci.* 123, e2526734123. doi:10.1073/pnas.2526734123

Hosseini, M., Rasmussen, L. M., and Resnik, D. B. (2023). Using AI to write scholarly publications. *Account. Res.* 31, 1–9. doi:10.1080/08989621.2023.2168535

Iyer, K. G., Yunus, M., O'Neill, C., Ye, C., Hyk, A., McCormick, K., et al. (2024). Pathfinder: a semantic framework for literature review and knowledge discovery in astronomy. *Astrophysical J. Suppl. Ser.* 275, 38. doi:10.3847/1538-4365/ad7c43

Liang, W., Zhang, Y., Cao, H., Wang, B., Ding, D., Yang, X., et al. (2024). "Monitoring AI-modified content at scale: a case study on the impact of ChatGPT on AI conference peer reviews," in *International conference on machine learning*.

MacQueen, J. (1967). "Some methods for classification and analysis of multivariate observations,"*Proc. Fifth Berkeley Symposium Mathematical Statistics and Probability* 1. Oakland, CA, USA, 281–297.

McKinney, W. (2010). Data structures for statistical computing in python. *Proc. of 9th Python Sci. Conf. (Austin, TX)* 445, 51–56. doi:10.25080/Majora-92bf1922-00a

Messeri, L., and Crockett, M. J. (2024). Artificial intelligence and illusions of understanding in scientific research. *Nature* 627, 49–58. doi:10.1038/s41586-024-07146-0

Mohammadi, E., Thelwall, M., Cai, Y., Collier, T., Tahamtan, I., and Eftekhar, A. (2026). Is generative ai reshaping academic practices worldwide? A survey of adoption, benefits, and concerns. *Inf. Process. & Manag.* 63, 104350. doi:10.1016/j.ipm.2025.104350

Nguyen, T. D., Ting, Y.-S., Ciucă, I., O'Neill, C., Sun, Z.-C., Jabłońska, M., et al. (2023). *AstroLLaMA: towards specialized foundation models in astronomy*.

Noy, S., and Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science* 381, 187–192. doi:10.1126/science.adh2586

Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., Mishkin, P., et al. (2022). Training language models to follow instructions with human feedback. *Adv. Neural Inf. Process. Syst.* 35, 27730–27744.

Oxford University Press (2024). Researchers and AI: survey findings. *Tech. Rep*.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: machine learning in python. *J. Mach. Learn. Res.* 12, 2825–2830.

Sharma, P., Vaidya, B., Wadadekar, Y., Bagla, J., Chatterjee, P., Hanasoge, S., et al. (2025). Computational astrophysics, data science and AI/ML in astronomy: a perspective from Indian community. *J. Astrophysics Astronomy* 46, 28. doi:10.1007/s12036-025-10049-9

Stokel-Walker, C. (2023). ChatGPT listed as author on research papers: many scientists disapprove. *Nature* 613, 620–621. doi:10.1038/d41586-023-00107-z

Trotta, R. (2025). The indiscriminate adoption of ai threatens the foundations of academia. *Nat. Astron.* 9, 1748–1749. doi:10.1038/s41550-025-02738-w

Van Noorden, R., and Perkel, J. M. (2023). AI and science: what 1,600 researchers think. *Nature* 621, 672–675. doi:10.1038/d41586-023-02980-0

Vipra, J., and Korinek, A. (2023). Market concentration implications of foundation models: the invisible hand of AI. *Tech. Rep*.

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., et al. (2020). SciPy 1.0: fundamental algorithms for scientific computing in python. *Nat. Methods* 17, 261–272. doi:10.1038/s41592-019-0686-2