



SCIREA Journal of Sociology

ISSN: 2994-9343

<http://www.scirea.org/journal/Sociology>

October 29, 2026

Volume 10, Issue 2, April 2026

<https://doi.org/10.54647/sociology841594>

What Is the “Human” in Human-Centered AI? Restoring the Aesthetic Subject

Byungkil Choi ^{1,*}

¹ Department of Arts, College of Arts and Design, Wonkwang University, Iksan, Republic of Korea

*Corresponding Author

Abstract

Human-centered artificial intelligence (HCAI) seeks to orient AI toward human values, autonomy, oversight, safety, and meaningful control. Yet the term “human-centered” leaves a prior question unresolved: what conception of the human is being centered? Much HCAI discourse approaches the human as user, decision-maker, rights-holder, stakeholder, or bearer of preferences and values. These conceptions are indispensable but incomplete. This article argues that HCAI must also recognize the human as an aesthetic subject: a being whose perception, attention, imagination, interpretation, and judgment are formed through experience. As recommendation, personalization, generative AI, and predictive systems increasingly organize the conditions under which people encounter and evaluate the world, AI can participate in forming rather than merely serving human preferences. The article develops the concept of aesthetic agency to describe the human capacity to participate meaningfully in the formation of one’s perceptual, imaginative, interpretive, and evaluative relation to the world. It distinguishes choice from judgment, preference from taste, generation from imagination, and optimization from experience. On this basis, it proposes an aesthetically expanded HCAI organized around openness, plurality, reflectivity, and transformability. To center the human is

not only to preserve control over AI, but to preserve the capacity to perceive, imagine, judge, and become otherwise.

Keywords: Human-Centered AI; Aesthetic Subject; Aesthetic Agency; Aesthetic Judgment; Human–AI Interaction

1. Introduction: The Missing Human in Human-Centered AI

Artificial intelligence has generated an unusual reversal in the language of technological progress. For much of the history of modern technology, innovation was commonly narrated in terms of what machines could newly accomplish: calculate faster, automate more efficiently, recognize more accurately, or perform tasks previously reserved for human expertise. The rapid development of artificial intelligence, however, has increasingly shifted attention from what machines can do to what should remain human. Questions of human autonomy, oversight, dignity, responsibility, safety, and control now accompany questions of technical performance. The emergence of “human-centered artificial intelligence” (HCAI) is one of the clearest expressions of this shift. Ben Shneiderman’s formulation is representative: rather than treating automation and human control as a simple zero-sum trade-off, he argues that well-designed systems can combine high levels of computer automation with high levels of human control in ways that improve performance while supporting human self-efficacy, mastery, creativity, and responsibility.^[1] This formulation is important for the present article because it makes human control a positive design objective rather than a residual safeguard added after automation has already been maximized.

The problem becomes particularly visible when four distinctions are maintained: choice is not identical with judgment; preference is not identical with taste; generation is not identical with imagination; and optimization is not identical with experience. A choice can be registered computationally as a behavioral event, whereas judgment may involve hesitation, comparison, interpretation, reflection, and the revision of one's own criteria. A preference can be inferred from previous selections, whereas taste can be educated, disrupted, expanded, or transformed through encounters with what one did not previously prefer. A generative system can produce unprecedented combinations, whereas imagination concerns not merely the production of alternatives but a subject's capacity to conceive possibilities and relate them meaningfully to

experience. Optimization can reduce friction and increase predicted satisfaction, whereas meaningful experience may depend precisely upon difficulty, ambiguity, surprise, resistance, or an encounter with something that does not correspond to an existing preference. The philosophical traditions of aesthetic judgment and aesthetic experience are particularly important here because they provide conceptual resources for understanding forms of experience that cannot be adequately described as preference satisfaction or rule-governed cognition. [2]

To articulate this claim, this article introduces the concept of aesthetic agency. Aesthetic agency refers here to the capacity of human subjects to perceive, attend, imagine, interpret, discriminate, and judge without these activities being exhaustively predetermined by computational prediction, personalization, or optimization. The qualification “exhaustively” is important. Human perception and judgment have never existed outside mediation. Language, education, artistic conventions, cultural institutions, technologies, and social environments have always shaped what human beings perceive and how they judge. The relevant goal can therefore not be an impossible recovery of an unmediated or technologically pure human subject. Accounts of technologically mediated perception are useful precisely because they show that technologies need not simply stand between an independently constituted subject and world; they can participate in shaping the relation through which the world is experienced. [3] The problem posed by AI is therefore one of the degree, opacity, direction, and reversibility of mediation.

Aesthetic agency also differs from a simple demand for “human control.” Control ordinarily describes a relation between an agent and a system: the human can intervene, override, authorize, reject, or terminate an automated process. HCAI has rightly emphasized such forms of supervisory control and human governance. [4] Aesthetic agency concerns something prior to and deeper than this relation. It concerns the constitution of the subject who exercises control. A user may possess an override button and yet inhabit an informational environment in which attention has already been extensively directed, preferences repeatedly reinforced, and available possibilities algorithmically delimited. Conversely, an encounter with an AI system may sometimes enlarge aesthetic agency by presenting unfamiliar possibilities, disturbing habitual categories, or enabling forms of experimentation that would otherwise have been inaccessible. The aim, therefore, is not to characterize AI mediation as inherently diminishing. It is to develop

a conceptual framework capable of asking when such mediation expands, redistributes, or contracts the human capacity for aesthetic experience and judgment.

This approach also requires resisting a romantic conception of the human. Restoring the aesthetic subject does not entail claiming that human judgment is inherently authentic, autonomous, creative, or superior to computational processes. Human taste is socially formed; perception is selective; judgment is fallible; imagination is historically and culturally conditioned. Dewey's account of aesthetic experience, for example, situates experience within the continuous transactions between organism and environment rather than within an isolated sovereign consciousness. [5] Nor does the argument depend upon establishing an absolute metaphysical boundary between human and machine creativity. The narrower and more consequential question is how increasingly generative, predictive, and adaptive AI systems alter the conditions under which human beings develop and exercise these capacities. The subject at issue is not a sovereign human standing outside technology, but a technologically mediated subject whose ability to respond to mediation remains normatively and aesthetically significant.

The argument proceeds in six stages. First, the article examines how the human has been constructed within human-centered AI, identifying the dominant figures of the user, autonomous decision-maker, rights-holder, stakeholder, and bearer of preferences and values. Second, it considers why the ethical vocabulary that has productively shaped responsible AI nevertheless provides an incomplete account of human subjectivity. Third, drawing on aesthetic theory and philosophical accounts of perception, experience, imagination, and judgment, it develops the conception of the human as an aesthetic subject. Fourth, it examines how contemporary AI systems increasingly intervene in the formation of attention, visibility, taste, imagination, and judgment rather than merely responding to pre-existing human intentions. Fifth, it analyzes the resulting problem of aesthetic agency through the distinctions between choice and judgment, preference and taste, generation and imagination, and optimization and experience. Finally, it proposes an expanded conception of human-centered AI organized around the preservation and cultivation of aesthetic agency.

The central claim of this article can therefore be stated simply: the human at the center of human-centered AI must be understood not only as a user, chooser, moral agent, or rights-holder, but as an aesthetic subject—a being capable of perception, imagination, interpretation, and judgment whose preferences are not merely given but formed, whose experiences cannot be reduced to optimized outcomes, and whose encounters with the unexpected can transform

the very standards by which judgment occurs. Human-centered AI will remain conceptually incomplete until it recognizes this dimension of human subjectivity. To restore the aesthetic subject is consequently not to add aesthetics to an already complete theory of human-centered AI. It is to reconsider what it means, in the age of artificial intelligence, to place the human at the center at all.

2. How Human-Centered AI Constructs the Human

2.1. From Human-Centered Design to Human-Centered AI

Human-centered artificial intelligence did not emerge in an intellectual vacuum. Its vocabulary inherits a substantial part of its conceptual orientation from the longer tradition of human-centered design, human–computer interaction, participatory design, value-sensitive design, and usability research. These traditions developed in response to a recurrent problem in technological design: systems optimized around technical performance often fail when they encounter the complexity of actual human practices. Human-centered design therefore sought to reverse the direction of technological development by beginning not with the capabilities of a system but with the needs, limitations, contexts, and goals of the people expected to use it. [6]

This enlarged normative role makes the underlying conception of humanity more consequential. If AI is to remain “human-centered,” the term cannot be treated as self-explanatory. What counts as a properly human value? Which human capacities deserve preservation? What forms of agency should not be delegated? Which transformations of human experience are acceptable? Existing HCAI approaches answer these questions in different ways, but they tend to organize the human around a small number of recurring figures. The most prominent among them are the human as user, autonomous decision-maker, and bearer of rights, values, and vulnerabilities.

2.2. The Human as User

The oldest and perhaps most persistent figure within human-centered technological discourse is the user. The concept of the user is attractive because it offers a concrete point of orientation. A system is human-centered insofar as it serves human purposes, accommodates human capacities, and supports tasks that users actually seek to perform. Usability, accessibility, comprehensibility, satisfaction, and efficiency can all be evaluated in relation to the user.

Recommendation systems make this problem particularly visible. They are frequently evaluated according to their ability to predict what users will click, watch, purchase, listen to, or otherwise

engage with. Yet Joseph Konstan and John Riedl explicitly argue that recommender systems should not be assessed only at the level of algorithms: the way an algorithm is embedded in the user experience can substantially determine the value of the recommendation to the user, which in turn requires measures broader than predictive accuracy alone. [7] This shift from algorithmic performance to user experience is significant, but it still tends to represent the person through observable interaction. Previous actions are translated into signals from which future preferences are inferred, so that the human subject appears computationally as a profile, probability distribution, cluster membership, or sequence of predicted responses.

The category of the user remains indispensable, but insufficient. It identifies the recipient and participant of technological interaction without fully accounting for the processes through which that participant becomes a particular kind of subject. For this reason, HCAI has increasingly supplemented the user model with a second conception: the human as autonomous decision-maker.

2.3. The Human as Autonomous Decision-Maker

A central concern in contemporary AI ethics and governance is the preservation of human autonomy. As automated systems acquire greater capacities to classify, recommend, predict, and decide, the question becomes whether human beings retain meaningful authority over actions that affect them and others. The human consequently appears in HCAI not merely as a user but as a decision-maker whose agency must not be displaced by automated processes.

This is especially important in situations of algorithmic decision support. A recommendation can function as an anchor. A ranking can influence attention. A predicted risk score can structure the interpretation of subsequent evidence. A generative system can define the initial range of conceivable responses. Thus, the human may remain formally authorized while becoming epistemically dependent upon the system. Human autonomy can survive at the level of final authorization while being weakened at the level of problem formulation, attention, interpretation, or imagination.

For this reason, meaningful human autonomy cannot be adequately described solely through the preservation of veto power or final decision authority. A stronger conception must consider the conditions under which judgment is formed. It must ask whether humans retain the capacity to understand, reinterpret, reject, reframe, and imagine otherwise. This broader requirement points beyond autonomy understood as choice and toward capacities that are partly aesthetic in character.

2.4. The Human as Bearer of Rights, Values, and Vulnerabilities

A third influential conception of the human emerges from the ethical and regulatory turn in AI. Here, the human is understood primarily as a bearer of rights, values, interests, and vulnerabilities. This approach has become central to responsible AI, trustworthy AI, rights-based AI, and emerging regulatory frameworks.

This represents an essential advance beyond narrow user-centeredness. A system can be highly usable while remaining unjust. It can satisfy user preferences while violating the rights of others. It can be efficient while producing unacceptable distributions of risk. Rights- and values-based approaches therefore correctly refuse to evaluate AI according to user satisfaction alone.

Moreover, the rights-bearing subject is often conceptualized in defensive terms. The central question becomes how the human can be protected against harm, manipulation, discrimination, or loss of autonomy. Such protection is essential, but a complete account of human-centeredness must also ask what forms of human capacity should be positively sustained or cultivated. To avoid harm is not yet to enable meaningful experience. To preserve privacy is not yet to preserve imagination. To prevent manipulation is not yet to cultivate judgment.

The normative subject of AI ethics therefore requires supplementation. The human is certainly someone whose rights must be protected and whose vulnerabilities must not be exploited. But the human is also a subject who learns how to perceive, develops forms of attention, acquires tastes, forms interpretations, confronts ambiguity, and revises judgments. These dimensions cannot be reduced to the language of rights without remainder.

2.5. The Limits of the Functional and Normative Human

The preceding analysis reveals that HCAI has developed through several distinct but overlapping conceptions of the human. The user emphasizes needs, usability, and successful interaction. The autonomous decision-maker emphasizes control, agency, and the preservation of meaningful human authority. The bearer of rights, values, and vulnerabilities emphasizes dignity, fairness, safety, protection, and normative standing. None of these conceptions should be discarded. Each identifies a real and indispensable dimension of human-centered technological design.

What remains underdeveloped is a third axis: the formative or aesthetic dimension of subjectivity. Humans do not simply enter technological systems with fully formed preferences, values, identities, and criteria of judgment. These are continually shaped through encounters

with environments, cultural forms, other persons, institutions, and technologies. Human beings are not only agents who act in the world; they are subjects whose ways of seeing, imagining, feeling, and judging are formed through experience.

The term “aesthetic” becomes necessary at precisely this point. It names a dimension of human existence prior to the narrow question of art or beauty. Human beings encounter the world sensibly and interpretively. They direct attention, experience salience, respond affectively, imagine alternatives, develop taste, and form judgments whose criteria are not always given in advance. These capacities help determine not merely what humans choose, but what appears available for choice; not merely what humans value, but how something comes to matter as valuable.

The next section therefore turns from the ethically protected and functionally empowered human to the aesthetic subject. Its purpose is not to replace the user, autonomous agent, or rights-holder, but to reveal the dimension of human subjectivity that makes those categories possible yet cannot be reduced to any of them. Before a human can choose, something must appear as worth choosing. Before a preference can be expressed, a field of experience must have been encountered. Before judgment can be exercised, a subject must have learned how to perceive differences, imagine alternatives, and respond to meanings that cannot be exhaustively specified in advance.

The aesthetic subject, as the term is used here, names this prior and formative dimension of the human: the level at which possibilities become perceptible, meaningful, and judgeable before they are registered as choices or preferences.

Table 1 summarizes the conceptual progression from the dominant figures of the human in HCAI to the aesthetic expansion proposed in this article.

Table 1. Conceptions of the Human in HCAI and Their Limits

Conception of the human	Central concern	What HCAI protects or supports	Limitation	Aesthetic expansion
User	Usability, needs, satisfaction	Effective interaction, accessibility, task support	Tends to assume relatively given goals and preferences	Subject whose preferences can be formed and transformed

Autonomous decision-maker	Autonomy, control, authority	Choice, veto, oversight, meaningful intervention	Can concentrate on final authorization rather than the formation of judgment	Subject who helps form the perceptual and evaluative conditions of decision
Rights-holder / value-bearer	Fairness, dignity, safety, rights	Normative protection against harm, discrimination, manipulation, and loss of autonomy	Primarily protective and defensive; does not fully explain how values become experientially meaningful	Subject whose perception, imagination, and judgment make valuation possible
Aesthetic subject	Perception, imagination, interpretation, judgment	Aesthetic agency	Not a replacement for the other conceptions, but their formative supplement	Openness, plurality, reflectivity, and transformability

(Note. The table synthesizes the conceptual distinctions developed in Sections 2–4; it is an analytical reconstruction proposed in this article, not a taxonomy attributed to a single prior source.)

3. Why Ethics Alone Does Not Exhaust the Human

3.1. The Ethical Turn in AI Discourse

The rise of artificial intelligence has been accompanied by a remarkable expansion of ethical vocabulary. Terms such as fairness, accountability, transparency, explainability, privacy, autonomy, dignity, responsibility, safety, and trustworthiness have become central to contemporary discussions of AI. Floridi and his collaborators give this ethical turn a particularly influential formulation by adapting the familiar principles of beneficence, non-maleficence, autonomy, and justice to AI and adding explicability, which they understand as joining intelligibility with accountability. [8] Their framework illustrates why AI can no longer be assessed solely according to whether it performs a task accurately or efficiently: technical performance must be evaluated together with the ways in which AI benefits or harms persons, constrains or supports autonomy, distributes advantages and burdens, and remains intelligible and answerable to those affected by it.

The problem is not that AI ethics asks the wrong questions. It asks indispensable questions. The problem is that those questions do not encompass every way in which artificial intelligence may affect human subjectivity. Ethical frameworks are particularly well equipped to ask whether an AI system treats persons justly, respects rights, causes harm, distributes benefits fairly, preserves autonomy, and permits accountability. They are less well equipped to ask how technologies reshape attention, perceptual salience, imagination, taste, interpretive habits, and forms of judgment.

The point, therefore, is not to move “beyond ethics” in the sense of abandoning ethical evaluation. It is to recognize that the human subject to whom ethical protections are owed is also a subject whose capacities for perception, imagination, and judgment are continually formed through experience.

3.2. Fairness, Accountability, Transparency, and Autonomy

The conceptual strengths and limitations of contemporary AI ethics become particularly clear in four of its most influential concerns: fairness, accountability, transparency, and autonomy.

Fairness asks whether the burdens and benefits of algorithmic systems are distributed unjustly and whether particular persons or groups are disadvantaged by biased data, discriminatory classifications, or unequal error rates. The literature cited here is especially important because it resists treating fairness as a single technical quantity. Barocas, Hardt, and Narayanan show that formal fairness criteria embody substantive assumptions about what kinds of equality matter, while Kleinberg, Mullainathan, and Raghavan demonstrate that important statistical fairness conditions can be mutually incompatible under common conditions. [9] Fairness analysis therefore cannot be completed by selecting a metric without also defending the normative conception of equality that makes that metric appropriate.

Accountability addresses another essential problem: whether responsibility can be identified when automated systems produce harmful or objectionable outcomes. Kroll and his co-authors argue that accountability for algorithmic decision systems cannot be reduced to exposing source code; what matters is whether a system can be governed through forms of procedural regularity, auditability, and review that make its decision process answerable. Mittelstadt and his collaborators likewise map algorithmic ethics around interconnected problems of opacity, bias, transformative effects, and traceability, emphasizing that increased transparency at one level does not by itself resolve questions of fairness or responsibility at another. [10] These

arguments are particularly relevant to HCAI because algorithmic decision-making can diffuse agency among developers, organizations, data suppliers, deployers, operators, and users, making accountability a problem of institutional structure as well as technical explanation.

The same difficulty arises with autonomy. Among the central principles of AI ethics, autonomy comes closest to the concern developed in this article because it explicitly addresses the human capacity for self-determination. Ethical approaches rightly seek to prevent manipulation, coercion, and the displacement of meaningful human choice. The preservation of human agency is consequently prominent in both responsible AI and HCAI. [11]

The limitation becomes most visible at this point. Autonomy can protect the formal capacity for self-direction, yet it does not by itself explain how the perceptual and evaluative capacities presupposed by self-direction are formed. It is in this more formative sense that aesthetics becomes philosophically necessary to the argument.

3.3. The Difference Between Moral Evaluation and Aesthetic Judgment

Moral evaluation and aesthetic judgment are related forms of evaluation, but the fact that both involve reflection does not make them interchangeable. Moral evaluation ordinarily asks whether an action, institution, or state of affairs can be justified by reference to norms concerning rightness, obligation, harm, justice, or the good. Aesthetic judgment, by contrast, becomes philosophically distinctive where the criterion of evaluation cannot simply be supplied in advance as a determinate rule. In the Kantian tradition, a judgment of taste arises from the subject's reflective relation to a particular presentation and cannot be mechanically deduced from a concept that already determines what must count as beautiful. [12] Its relevance to HCAI therefore lies less in a doctrine of beauty than in the structure of a judgment whose criteria are not exhausted by prior classification or preference data.

This distinction also explains why aesthetic judgment should not be reduced to private liking. Kant's account is more demanding: although the judgment of taste is grounded in feeling, it makes a claim to communicability that exceeds the report that an object happens to please a particular individual. The difference matters in AI-mediated environments because a system may predict what a person is likely to like without thereby reproducing the reflective process through which that person interprets an encounter, compares possibilities, revises a standard, or comes to regard an initially unfamiliar object as significant. Prediction can therefore approximate a pattern of response while leaving open the question of how judgment itself is

formed.

None of this establishes a rigid opposition between ethical and aesthetic judgment. Aesthetic experience can have moral significance, and moral judgment can depend upon cultivated perception and imagination. Martha Nussbaum's account of literary imagination and Iris Murdoch's emphasis on disciplined attention both show that ethical life itself may depend upon modes of seeing that cannot be reduced to the mechanical application of rules. [13] The appropriate relation is therefore one of overlap without reduction: ethical analysis can determine whether persons are treated justly or their rights respected, while aesthetic analysis can clarify how the capacities of perception, imagination, and reflective judgment through which situations become meaningful are themselves formed. An adequate HCAI needs both levels because protecting the freedom to judge is not yet the same as understanding or sustaining the human capacity for judgment.

3.4. Experience Cannot Be Reduced to Preference Satisfaction

One of the most persistent assumptions underlying contemporary digital systems is that human welfare can be approximated through the satisfaction of preferences. If systems can infer what users want and supply it more efficiently, the interaction appears successful. Recommendation, personalization, optimization, and predictive modeling are frequently organized around this logic.

Preferences are often represented as prior dispositions that can be elicited, inferred, ranked, or predicted. Experience, by contrast, can transform the subject who enters it. John Dewey's account of aesthetic experience is particularly instructive here. Dewey rejects the conception of experience as a sequence of isolated sensations consumed by an already constituted subject. Experience arises through an ongoing transaction between organism and environment and becomes aesthetically significant when perception, action, tension, development, and consummation acquire a meaningful unity. [14]

On such an account, human experience cannot be modeled adequately as a succession of preference-satisfying events. The subject does not simply consume outcomes. The subject is altered through engagement.

This dynamic becomes aesthetically significant because taste is not simply a collection of existing likes and dislikes. Taste develops historically. It can be educated, expanded, refined, challenged, corrupted, standardized, or transformed. Pierre Bourdieu famously emphasized the

social formation of taste, while Dewey's account of experience equally resists any conception of preference as an isolated internal fact. [15]

A genuinely human-centered system should therefore preserve what might be called experiential openness: the possibility that an encounter will not merely satisfy the subject who enters it but alter that subject's sensibility, interpretation, or judgment.

3.5. From the Ethically Protected Human to the Aesthetically Experiencing Human

The argument of this section should not be understood as a choice between ethics and aesthetics. Human beings need ethical protection. Their privacy matters. Their rights matter. Fairness matters. Accountability matters. Autonomy matters. No theory of human-centered AI can responsibly diminish these concerns.

The distinction is crucial because artificial intelligence increasingly occupies environments in which formation occurs. It accompanies writing, viewing, listening, searching, selecting, learning, creating, and communicating. In generative AI especially, technological mediation can enter into activities previously understood as closely connected with imagination and expression. The system may suggest the sentence before the writer has formulated it, produce the visual alternative before the artist has imagined it, or frame the question before the user has determined what the problem is.

This capacity is closely related to autonomy but is not identical with it. Autonomy emphasizes self-direction. Aesthetic agency emphasizes the formation of the sensibility through which self-direction becomes possible. It concerns the ability not merely to choose but to experience alternatives as meaningful enough to be chosen; not merely to express preferences but to acquire and revise taste; not merely to receive generated possibilities but to judge them; not merely to optimize outcomes but to undergo experiences whose value may not have been predictable in advance.

The following section therefore turns directly to the philosophical foundations of this subject. Beginning from *aisthesis* as perception, and moving through Kantian judgment, Deweyan experience, and phenomenological accounts of embodied perception, it asks what it means to understand the human not merely as a chooser or rights-holder, but as a being for whom the world is continually constituted through sensing, imagining, interpreting, and judging.

4. The Human as an Aesthetic Subject

4.1. *Aisthesis*: Before Beauty, There Is Perception

To restore the aesthetic subject to human-centered AI, it is necessary first to recover a broader meaning of the aesthetic than the one commonly assumed in contemporary usage. “Aesthetic” is often associated with beauty, artistic quality, style, or visual appeal. In technological discourse, it may refer to the pleasing appearance of an interface or the expressive qualities of AI-generated images. Such meanings are legitimate but historically and philosophically narrower than the conceptual field from which aesthetics emerged.

The term ultimately derives from the Greek *aisthēsis*, meaning sensation, perception, or sensible apprehension. When Alexander Gottlieb Baumgarten established aesthetics as a distinct philosophical discipline in the eighteenth century, he defined it in relation to *cognitio sensitiva*, sensible cognition. [16] Aesthetics, in this original modern philosophical formulation, was therefore not simply a theory of beauty. It concerned a mode of knowing and experiencing that differs from abstract logical cognition while possessing its own order and significance.

The aesthetic subject is therefore first of all a perceptual subject, because agency presupposes a field in which something can appear, acquire salience, and become available for interpretation before it can become an object of choice.

This claim has immediate relevance for artificial intelligence. Contemporary AI increasingly operates at the level of perceptual access. Search systems decide which information becomes readily findable. Ranking systems order visibility. Recommender systems determine what appears next. Computer vision systems classify objects and persons, while generative systems produce images and representations that enter into the perceptual environment. In each case, technological mediation affects not only what users choose after perception but what becomes available for perception in the first place.

Before beauty, then, there is perception. And before autonomous choice, there is a world that has become visible in a particular way.

4.2. Kant: Judgment Beyond Determinate Concepts

If *aisthesis* identifies the perceptual foundation of aesthetic subjectivity, Kant's aesthetics provides a second crucial dimension: judgment.

In the *Critique of the Power of Judgment*, Kant distinguishes judgments of taste from judgments determined by concepts. To judge something aesthetically is not simply to recognize that an object belongs to a category or satisfies a pre-established rule. The judgment of taste arises from the subject's reflective relation to a particular presentation and cannot be mechanically deduced from a determinate concept of what the object ought to be. [17]

This point is often obscured when aesthetic judgment is reduced to subjective preference. Kant's argument is more demanding. A judgment of taste is subjective in that it is grounded in feeling, yet it nevertheless claims a kind of communicability that exceeds the statement of private liking. To say “this pleases me” is not identical to judging “this is beautiful.” The latter implicitly asks others to see or feel something that cannot simply be demonstrated through a concept.

Generative AI brings this problem into sharp focus. Suppose a system generates twenty visual possibilities for a designer. The system can increase the speed and range of available alternatives. Yet the decisive aesthetic activity may lie not in generation but in determining which possibility is compelling, why it is compelling, whether an initially unattractive option deserves reconsideration, and what criteria should govern further revision.

Kant also helps clarify why judgment cannot be reduced to preference data. Preferences can be modeled through patterns of previous selection. Reflective judgment, by contrast, may interrupt such patterns. A person may encounter something that does not fit previous tastes and nevertheless judge it significant. The ability to revise one's standards in response to the particular is precisely what predictive systems find difficult to represent when they infer future responses from prior behavior.

That capacity becomes especially important in a technological environment increasingly organized around prediction.

4.3. Dewey: Aesthetic Experience as a Mode of Being in the World

If Kant clarifies the structure of reflective judgment, John Dewey broadens aesthetics by relocating it within the continuity of ordinary experience. In *Art as Experience*, Dewey rejects the isolation of aesthetic experience from everyday life and argues that the aesthetic emerges from intensified forms of interaction between organism and environment. [18]

This shift is especially important for the present argument because it prevents the aesthetic subject from being confined to museums, artworks, or specialized cultural practices. For Dewey,

experience becomes aesthetic when its elements acquire a particular qualitative unity through rhythms of doing and undergoing, tension and resolution, resistance and consummation. The aesthetic is therefore not merely a property of objects. It is a mode of relational experience.

Aesthetic experience, on this account, cannot be understood as one-way consumption, because the significance of an encounter lies partly in the reciprocal process through which the subject acts upon an environment and is altered by undergoing it.

This insight complicates the preference model discussed in Section 3. If human-centered systems seek only to maximize what users already prefer, they risk overlooking experiences whose value emerges only retrospectively. A difficult book may become meaningful after sustained engagement. An unfamiliar artwork may initially repel but later reshape perception. A creative failure may contribute more to artistic development than an immediately successful output.

This notion has important consequences for AI design. The user model often assumes that the system should adapt to the human. Dewey suggests a more reciprocal relation: environments shape humans while humans shape environments. Human-centered AI must therefore evaluate not merely how effectively AI adapts to users, but how users themselves are being formed through repeated interaction with adaptive systems.

Dewey contributes a third dimension to the aesthetic subject: the human is not merely a perceiver or judge but a being constituted through temporally extended transactions with environments.

Human-centered AI should therefore ask not only, “Did the system produce the desired result?” but also, “What kind of experience occurred in reaching that result, and what kind of subject did that experience help to form?”

4.4. Phenomenology: Embodiment, Perception, and Situated Experience

Phenomenology adds a fourth dimension by challenging any conception of the human as a detached information processor standing outside the world. Human experience is embodied, situated, perspectival, and relational.

Maurice Merleau-Ponty's phenomenology is particularly important in this respect. In *Phenomenology of Perception*, perception is not treated as the internal representation of an external world by a disembodied mind. The body is itself the medium through which a world becomes practically and perceptually available. [19] We do not first receive neutral

information and subsequently decide how to act. We inhabit a field of possibilities structured through bodily capacities, habits, orientations, and contexts.

The embodied dimension also becomes important in generative AI. Text-based interaction can create the impression that intelligence consists predominantly in linguistic exchange. Yet much human understanding is bodily, practical, tacit, and environmentally situated. Artists learn through gesture, material resistance, spatial orientation, bodily rhythm, and repeated practice. A sculptor's knowledge of material cannot be reduced to propositions about sculpture. A musician's phrasing depends partly upon bodily timing and kinaesthetic awareness.

The aesthetic subject is thus not merely a mind making judgments about representations. It is an embodied subject whose sensing and acting are intertwined.

Postphenomenology further extends this insight by emphasizing technological mediation. Don Ihde and Peter-Paul Verbeek have shown that technologies can transform the relations through which humans encounter the world.[20] A microscope does not merely add information to an unchanged perceiver; it makes a different perceptual relation possible. A camera reorganizes seeing. A navigation system alters how space is encountered.

Phenomenology thus contributes a fourth principle to aesthetically human-centered AI: the human must be understood as embodied, situated, and technologically mediated, rather than as an abstract bundle of preferences and decisions.

4.5. The Aesthetic Subject as Sensing, Interpreting, Imagining, and Judging

The preceding traditions can now be brought together into a more precise account of the aesthetic subject. From aisthesis comes the recognition that agency begins with sensing and perception: before a person can choose among alternatives, something must first become perceptually available and salient enough to invite interpretation. Kant adds the reflective dimension of judgment, showing why evaluation cannot always be reduced to the application of an antecedently fixed concept. Dewey, in turn, situates perception and judgment within temporally extended transactions between organism and environment, while phenomenology emphasizes that these transactions are embodied, situated, and technologically mediated. Taken together, these traditions describe the human not as a collection of isolated faculties but as a subject whose relation to the world is continually organized through sensing, interpreting, imagining, and judging.

These capacities are analytically distinguishable but experientially intertwined. Perception is

never a merely neutral reception of data, because what becomes salient is already encountered from within histories, habits, bodily orientations, and interpretive contexts. Imagination can reorganize perception by making absent or unrealized possibilities available, while judgment evaluates those possibilities and, in doing so, can alter the standards that shape subsequent perception. The human is therefore aesthetically significant not because human beings uniquely produce art, but because sensibility, imagination, interpretation, and judgment participate in the formation of who they become. Behavioral representations may model aspects of these activities, but they cannot simply be identified with the lived and developing relation through which a subject experiences the world.

This account also clarifies why aesthetic agency cannot be equated with maximal human control. A person may retain formal control over a system while remaining deeply shaped by the informational and perceptual environment the system has organized. Conversely, a person may deliberately surrender a measure of control in order to encounter an unexpected result, unfamiliar perspective, or generative variation that enlarges imagination. The relevant criterion is therefore whether mediation remains sufficiently open to reflective response: whether the subject can reinterpret what appears, resist habitual patterns, imagine alternatives, and revise judgments rather than merely confirm statistical continuities in prior behavior. Understood in this way, the aesthetic subject is a subject of openness whose future relation to the world cannot be exhausted by the profile through which it is predicted.

Once the human is understood in these terms, the central question for HCAI changes. It is no longer sufficient to ask whether AI performs tasks for people, respects their rights, or leaves them a final decision. It becomes necessary to ask how AI intervenes in the perceptual, imaginative, and evaluative processes through which a world becomes available and a judgment becomes possible. The next section therefore examines artificial intelligence not simply as a tool or decision system, but as an increasingly powerful mediator of aesthetic subjectivity.

5. What AI Changes in Aesthetic Subjectivity

5.1. AI as Mediator of Perception and Attention

If human beings are aesthetic subjects in the sense developed in the preceding section, artificial intelligence acquires significance not merely when it replaces a decision or generates an artifact, but whenever it intervenes in the conditions under which a world becomes perceptually available. AI increasingly performs precisely this function. It selects, ranks, filters, recommends,

summarizes, labels, predicts, and generates. Through these operations, it participates in organizing not only what users *do* but what they are likely to *notice* before acting.

A search engine does not merely retrieve information existing elsewhere. It ranks results, thereby distinguishing what is encountered immediately from what requires further effort to discover. A social-media feed does not simply contain available content. It organizes temporal and visual priority. A recommender system does not merely respond to a request for more content; it actively proposes what the next object of attention should be. An AI assistant can summarize a large body of information and thereby determine which distinctions, themes, and conclusions become perceptually and cognitively prominent.

The critical distinction is therefore not between mediated and unmediated perception, since the latter is not a plausible ideal, but between forms of mediation that keep the perceptual field open to revision and forms that progressively narrow it around prior prediction.

An aesthetically significant AI system may broaden or narrow perceptual plurality; encourage or discourage exploration; expose or conceal the logic of selection; stabilize existing interests or facilitate unexpected encounters. The central HCAI question should therefore not be whether AI mediates perception—it inevitably does—but whether that mediation preserves sufficient openness, plurality, intelligibility, and revisability for aesthetic agency to remain active.

Accordingly, AI increasingly mediates not merely what humans know after information has been encountered, but what becomes available as an object of attention in the first place; this upstream position is what gives algorithmic ordering its specifically formative significance.

5.2. Recommendation and the Organization of Visibility

Recommendation systems make this transformation particularly visible because their principal function is precisely to select from an abundance of possible objects and determine which should be presented to particular users. They have become central to streaming services, social media, e-commerce, news distribution, search, cultural consumption, and increasingly educational and informational environments.

Recommendation therefore occupies an unstable position between prediction and formation: it infers likely relevance from prior behavior, yet the resulting exposure can itself contribute to the preferences from which later predictions are made.

Such recursion does not entail that the users are passively determined. People reject recommendations, search independently, share unexpected material, abandon platforms, and

deliberately cultivate interests contrary to algorithmic prediction. Nor is personalization necessarily narrowing. Empirical research on recommender systems shows that the effects of personalization, filter bubbles, diversity, and exposure are more complex than a simple deterministic account permits.[21]

For this reason, the familiar concept of the “filter bubble” should be used cautiously rather than as a universal description of algorithmic recommendation. The philosophical point developed here does not depend on proving that every recommender system traps every user within an informational bubble. It depends on the more modest and defensible observation that recommender systems systematically allocate visibility, and that repeated allocations of visibility can become formative conditions of experience.

Aesthetic development has historically depended upon institutions and practices that introduce subjects to objects they would not necessarily have selected independently: teachers assign unfamiliar literature; critics advocate difficult works; museums juxtapose historically distant objects; friends recommend music outside one's usual taste; libraries permit accidental discovery. These forms of mediation are not neutral either. They embody institutional values, cultural hierarchies, and exclusions. Nevertheless, they illustrate that mediation can be oriented toward something other than the confirmation of already observed preference.

The organization of visibility consequently becomes an organization of possible aesthetic formation, not because recommendation determines taste in a simple causal sense, but because repeated patterns of exposure condition what can become familiar, difficult, attractive, or worthy of sustained attention.

5.3. Generative AI and the Outsourcing of Imagination

Recommendation systems select among existing alternatives. Generative AI introduces a more radical development because it can participate in producing the alternatives themselves.

Large language models, image generators, music-generation systems, and multimodal models can rapidly supply texts, images, designs, compositions, narratives, concepts, variations, and solutions in response to relatively limited instructions. Their significance for aesthetic subjectivity therefore differs from that of conventional recommendation. Recommendation intervenes principally in encounter; generative AI can intervene directly in ideation.

It would be inaccurate to suggest that the use of generative AI necessarily replaces human imagination. Empirical research already indicates more complicated effects. In several

experiments, interaction with ChatGPT has been found to improve performance on certain creative ideation tasks, particularly through the combination of remotely related concepts, while other studies emphasize that human–AI creative effects vary according to task and interaction design. [22] [23] Lee and Chung, for example, found across five experiments that ChatGPT could increase the creativity of ideas compared with no technological assistance or conventional web search, although its advantage was stronger for incremental than radically novel ideas.

This transformation can be enormously enabling. It can lower barriers to creative participation, assist users with limited technical skills, facilitate experimentation, and enable rapid exploration of alternatives. Michelle Vaccaro, Abdullah Almaatouq, and Thomas Malone’s meta-analysis makes the point more precisely: across more than one hundred experimental studies, human–AI combinations did not, on average, outperform the better of the human or AI working alone, but the results varied sharply by task, with more favorable effects in open-ended creation tasks than in decision tasks involving predefined alternatives. [24] Their finding cautions against general claims that AI assistance either uniformly augments or uniformly diminishes human performance. For the present argument, its significance is that the aesthetic consequences of collaboration must be examined in relation to how imaginative and evaluative labor is actually distributed within a particular task.

These performance gains, however, do not settle the aesthetic question. Evaluation of output quality asks whether assistance improves what is produced; an analysis of aesthetic agency asks how the distribution of imaginative labor changes when the system begins to generate possibilities before the human has articulated them independently.

The human-centered question should therefore concern neither prohibition nor maximal automation, but the design of interactions in which generative assistance remains compatible with the continued exercise of human imaginative agency. What matters is not that humans generate every possibility unaided, but that they retain a reflective awareness of when generation is being delegated and a meaningful capacity to redirect, reject, transform, or originate possibilities themselves.

The critical risk is thus not assistance as such, but a habitual displacement of the interval in which a subject struggles to formulate what is not yet clear. If that interval disappears unnoticed, greater productive fluency may coexist with a diminished exercise of imagination.

5.4. Prediction, Personalization, and the Anticipation of Taste

Generative AI intervenes in the production of possibilities. Predictive personalization introduces another transformation: AI can increasingly act in anticipation of the subject.

This changes the temporal relation between preference and environment. Rather than merely responding after a preference has been formed, predictive systems can infer a likely preference and organize the field of subsequent exposure in advance.

A system trained on previous selections can infer that a user usually prefers particular genres, styles, topics, formats, or aesthetic characteristics. It can then organize future options around that inference. The more accurate the model becomes, the more the environment may appear to conform naturally to the user's taste.

Yet taste cannot be adequately represented as continuity with previous selection, because its development also includes discontinuity.

People become tired of what they once loved. They acquire interests through education or relationships. They encounter an unfamiliar work that restructures previous standards. They sometimes deliberately cultivate difficult tastes. They may value an object despite not enjoying it, or enjoy something while judging it aesthetically trivial.

Again, this should not be exaggerated into a deterministic “filter bubble” thesis. Areeb and his co-authors’ systematic review is useful precisely because it finds evidence that recommender systems can produce filter-bubble effects while also emphasizing that these effects arise through multiple biases and can be mitigated through diversity-oriented recommendation, transparency, and explanatory mechanisms. [25] Recommender systems can therefore be designed to promote novelty, diversity, and serendipity, and users themselves actively shape their informational trajectories. The relevant philosophical claim is consequently not that personalization inevitably imprisons users within fixed tastes, but that the design of personalized exposure can either narrow or enlarge the conditions under which taste develops.

A maximally personalized environment might appear maximally human-centered because everything has been adapted to the individual. But if personalization becomes too complete, human-centeredness risks becoming model-centeredness: the environment is organized not around the irreducible subject but around the system's representation of that subject.

For this reason, aesthetic agency requires the possibility that users may surprise the system and, through unfamiliar encounters, surprise themselves; prediction must remain corrigible by experiences that exceed the profile from which it was generated.

The central aesthetic question is therefore whether algorithmic anticipation remains provisional rather than becoming a self-confirming organization of the user's future environment.

Aesthetically human-centered AI should treat predictions as hypotheses about subjects, not definitions of them.

5.5. From Assisting Judgment to Pre-Structuring Judgment

The developments considered so far—algorithmic attention, organized visibility, generative ideation, and predictive personalization—converge on a deeper transformation in the position that AI occupies within human judgment. In a conventional model of decision support, the human first encounters a problem, gathers information, imagines alternatives, and begins to form criteria; technology then assists one or more stages of an already developing judgment. Contemporary AI can operate earlier. The same system may infer preferences from prior behavior, determine which evidence becomes salient, summarize or frame that evidence, generate possible responses, and adapt subsequent presentation to the user's reactions. AI can thus move from assisting judgment to participating in the organization of the conditions from which judgment begins.

To describe this process as pre-structuring is not to equate it automatically with manipulation. Human judgment has always depended upon prior structures of selection and interpretation: educational institutions determine curricula, museums select and juxtapose objects, editors order texts, search indexes organize information, and language itself supplies categories through which experience becomes intelligible. The distinctive issue raised by AI is the conjunction of scale, personalization, prediction, opacity, and recursive adaptation. Because an AI system can repeatedly observe a subject, revise its model, and reorganize the next environment of encounter in response, mediation can become dynamically tailored to the very behavior that it helps to elicit. The system is therefore capable of functioning not merely as one source of advice among others but as a recursively adaptive environment of judgment.

This shift exposes a limitation in accounts of human oversight that focus primarily on review or override. The capacity to reject an automated recommendation remains indispensable, especially in high-stakes contexts, but it does not by itself reveal how much of the preceding field of judgment has already been computationally arranged. The relevant distinction concerns

the human's position within the process. When AI assists a judgment already being formed, the human retains substantial responsibility for framing the problem, identifying relevant evidence, imagining alternatives, and developing criteria. When AI pre-structures judgment, those elements may already have been organized before reflective evaluation begins. The appropriate HCAI question is therefore not only whether a human remains formally authorized to decide, but whether the human still participates meaningfully in forming the perceptual and imaginative horizon within which the decision becomes intelligible.

This argument does not require excluding AI from early stages of inquiry or creativity. AI can broaden evidence, expose neglected patterns, challenge assumptions, and generate unfamiliar alternatives, thereby enlarging rather than contracting aesthetic agency. The concern arises when such assistance habitually occupies the interval in which perception, imagination, preference, and criteria would otherwise be formed without leaving adequate room for reflective revision. Artificial intelligence increasingly acquires the capacity to predict what a subject is likely to notice, want, imagine, and choose; aesthetically human-centered AI must therefore preserve the complementary capacity of the subject to encounter what was not anticipated, imagine what was not generated in advance, revise what was previously preferred, and judge according to criteria that emerge through experience. The following section develops this structural tension as a crisis of aesthetic agency and explains why keeping a human formally 'in the loop' cannot, by itself, resolve it.

6. The Crisis of Aesthetic Agency in Human–AI Interaction

The preceding section argued that artificial intelligence increasingly operates upstream of explicit human judgment. It organizes attention, distributes visibility, generates alternatives, anticipates preferences, and structures the informational environments within which choices are made. These developments do not establish that human agency has disappeared. Nor do they justify a simple opposition between autonomous humans and determining machines. They reveal, rather, a structural tension within human–AI interaction: the systems designed to augment human agency may increasingly participate in forming the conditions under which that agency is exercised.

These distinctions are not semantic refinements. Together they identify the difference between a human being who remains formally involved in an AI-mediated process and a human being

whose capacities for perception, imagination, interpretation, and judgment remain actively exercised through that process.

6.1. Choice Is Not Identical with Judgment

Choice occupies a privileged position within contemporary accounts of human agency. To preserve human choice appears to preserve autonomy: if the human retains the ability to select, reject, approve, or override, the system remains under meaningful human control.

This distinction has precedents in philosophical accounts of judgment. Kant's reflective judgment is significant precisely because judgment is not always the subsumption of a particular under a determinate concept already available in advance. [26] Hannah Arendt's appropriation of Kant similarly emphasizes the capacity to orient oneself among particulars by considering perspectives rather than mechanically applying a predetermined rule. [27]

The problem is particularly acute when AI output becomes the default starting point. The first generated answer may function as an anchor, making subsequent thought responsive to the machine's formulation rather than genuinely prior to it. Raja Parasuraman and Dietrich Manzey's review of automation research distinguishes automation complacency from automation bias but argues that both are rooted in overlapping attentional mechanisms: users may insufficiently monitor automated systems or accept erroneous automated advice even when the system is imperfect. [28] Their analysis also shows that these effects are not confined to inexperienced users and are not eliminated simply by training. Although this literature predates generative AI, it provides a defensible basis for the narrower claim made here: formal human authority does not ensure reflective independence when automated outputs structure attention and become the default starting point for evaluation. Generative AI extends this concern beyond decisions among established alternatives to the generation and framing of alternatives themselves.

Aesthetic agency requires more than the first capacity. Choice selects among alternatives that have already become available, whereas judgment can interrogate the adequacy of those alternatives, the framing of the problem, and the criteria by which the alternatives are compared.

A person who retains the final click or veto therefore remains a chooser, but whether that person has exercised judgment depends on the extent to which the preceding field of alternatives and criteria remained open to reflective scrutiny.

6.2. Personalization Is Not Identical with Aesthetic Autonomy

Personalization is frequently presented as one of the clearest expressions of human-centered technology. Instead of requiring all users to adapt to a uniform system, personalized systems adapt themselves to individual needs, histories, behaviors, and preferences. The more accurately a system understands the user, the more relevant and convenient the interaction can become.

It would be tempting to describe increasingly accurate personalization as maximal autonomy, since the environment appears ever more responsive to the individual. That inference is nevertheless too quick.

Aesthetic autonomy includes the ability not merely to obtain what one already prefers but to develop preferences through encounters that were not selected because they fit the present self. Taste depends partly upon exposure to unfamiliarity. A subject may discover that an initially difficult object deserves attention, that a previously disliked form has become compelling, or that established preferences no longer satisfy.

This is why the concept of aesthetic autonomy should not be romanticized as independence from mediation. Taste has never been formed in isolation. Pierre Bourdieu demonstrated powerfully that taste is socially conditioned and implicated in structures of education, class, cultural capital, and distinction. [29] Museums, schools, critics, markets, families, peer groups, and media institutions have always participated in aesthetic formation.

The relevant normative aim is therefore not independence from mediation but productive non-conformity between the user model and the user's developing experience. A system may learn from past behavior without treating that past as the normative template for future exposure.

The human subject must, in other words, retain the possibility of becoming aesthetically different from the subject represented in the data; otherwise, personalization risks converting a descriptive model of prior behavior into a prescriptive architecture of future experience.

6.3. Generation Is Not Identical with Imagination

Generative AI makes the distinction between production and imagination increasingly urgent.

The claim that generation is not identical with imagination does not require denying creativity to AI in advance. Nor does it depend upon proving that every human-generated artifact is more original than every machine-generated artifact. Empirical comparisons on divergent-thinking

tasks already show that contemporary generative systems can equal or exceed average human performance on some conventional measures of creative output. [30] The philosophical argument should therefore not rest upon a competition over who produces more novel artifacts. Human imagination is not merely an output-producing mechanism. It is bound to memory, perception, embodiment, affect, anticipation, desire, cultural inheritance, and lived experience. To imagine is not simply to produce something statistically or formally novel; it is to make an absent possibility meaningful from within a situated relation to the world.

Such external generation can certainly enhance imagination when it supplies unfamiliar materials, prompts recombination, or enables experimentation that the user could not otherwise undertake. The relevant issue is therefore not whether AI-generated possibilities are intrinsically opposed to imagination, but how their availability redistributes the activity through which possibilities become meaningful to the subject.

This distinction parallels one familiar from physical technology. A calculator increases the ability to obtain mathematical results without necessarily strengthening mental arithmetic. A navigation system improves the ability to reach destinations while potentially changing the cognitive practices through which spatial orientation is learned. The analogy should not be pushed too far, but it illustrates the general point that improved task performance and strengthened underlying capacity are not always identical.

The appropriate question is consequently relational rather than exclusionary: which moments of imaginative formation remain actively exercised by the human, which are delegated to the system, and how does that distribution affect the capacity to imagine beyond the alternatives already supplied?

An aesthetically human-centered AI should not prescribe a single correct distribution. But it should make that distribution visible and preserve opportunities for the human to remain more than a selector of machine-generated possibilities.

6.4. Optimization Is Not Identical with Experience

Optimization is among the foundational logics of computational systems. Given an objective, a system seeks an outcome that performs better according to relevant criteria. Optimization makes technological systems useful precisely because it can reduce waste, increase accuracy, shorten time, personalize services, and improve performance.

Dewey's account is again instructive. An experience possesses aesthetic quality not because every moment maximizes pleasure but because tensions, resistances, actions, and undergoings become organized into a meaningful temporal development. [31] The path matters to the quality of the experience.

A potential conflict consequently arises between frictionless interaction and formative experience. Removing difficulty can be genuinely liberating when difficulty is merely instrumental, but it can also remove forms of resistance through which perception, skill, or judgment would otherwise develop.

The distinction is particularly relevant to generative AI in education and creative practice. If a system supplies a polished answer immediately, the output may be better than what the user could initially have produced. Yet the educational or aesthetic value of the activity may have resided partly in developing the capacity required to produce, criticize, or revise such an answer.

Human-centeredness cannot mean satisfying every immediate preference if doing so systematically weakens capacities that enable richer future agency. This is already recognized outside aesthetics. Education does not simply give students whichever task they find easiest. Physical training does not eliminate exertion. Democratic deliberation cannot be reduced to presenting citizens only with views they already prefer. This is an aesthetic question because it concerns the qualitative structure of experience rather than efficiency alone.

6.5. When AI Begins to Shape the Conditions Under Which Humans Judge

The four distinctions developed above converge upon a single problem: formal human participation is not a sufficient measure of aesthetic agency when AI increasingly structures the conditions from which participation acquires its content. The crisis of aesthetic agency therefore does not begin only when AI completely replaces human judgment. It can arise earlier, when a person remains visibly involved in a process but the perceptual field, available alternatives, likely preferences, and evaluative criteria have already been organized in ways that leave little room for independent reinterpretation or revision. What matters is consequently not merely whether the human appears somewhere in the decision chain, but where and how the human remains active within the formation of judgment.

This formulation does not imply that judgment should ever be free from mediation. Such a requirement would be incoherent. Teachers frame problems, editors select evidence, museums organize objects, critics propose interpretations, languages provide categories, and institutions

establish standards. The relevant issue is whether mediation remains contestable and whether the subject can still participate in transforming the terms through which experience is organized. Aesthetic agency is therefore better understood as a temporally extended capacity than as an isolated outcome: it does not require every choice to be original, unpredictable, difficult, or resistant to recommendation, nor does it require constant refusal of AI assistance. It requires that repeated interaction leave meaningful room for the subject to perceive differently, imagine alternatives, revise preferences, and alter criteria rather than merely enact the continuity predicted by a computational model.

The crisis becomes clearest when the computationally legible subject begins to substitute for the aesthetically developing subject. A profile can model, predict, personalize, and optimize aspects of behavior, but the person represented by that profile remains a subject of experience whose future perception and judgment may change. The gap between these two perspectives is not simply an error to be eliminated by better prediction. It is also the space in which aesthetic development remains possible. Human-centered AI becomes aesthetically inadequate when a descriptive representation of the person begins to determine the experiential environment so thoroughly that the predicted subject becomes the template for the possible subject.

For this reason, the crisis of aesthetic agency is not fundamentally a contest over which capacities belong exclusively to humans and which may be acquired by machines. Such a strategy would become increasingly fragile as AI systems improve. The constructive question is instead whether AI-mediated environments preserve perceptual openness, imaginative plurality, reflective judgment, and experiential transformation. The next section develops these requirements as principles for an aesthetically expanded HCAI. Its objective is not to place the human outside technological mediation, but to make mediation compatible with a subject who remains capable of becoming otherwise.

7. Restoring the Aesthetic Subject to Human-Centered AI

The preceding sections have argued that contemporary human-centered AI remains incomplete when it identifies the human primarily as a user, decision-maker, rights-holder, or bearer of preferences and values. These figures are indispensable, but they do not sufficiently account for the human as a subject whose capacities for perception, imagination, interpretation, and judgment are formed through ongoing interaction with technologically mediated environments.

The constructive task is therefore not to reject HCAI but to revise its philosophical center so that human-centeredness refers not only to control, welfare, and preference satisfaction but also to the preservation of capacities through which preferences and judgments are formed.

Such a revision requires a shift from treating human control as the endpoint of the analysis toward treating aesthetic agency as a complementary and more formative dimension of human–AI relations.

7.1. From Human Control to Human Aesthetic Agency

Human control is one of the strongest organizing concepts in HCAI. Shneiderman’s two-dimensional HCAI framework is especially important because it rejects the assumption that more automation must necessarily mean less human control: depending on context, systems can be designed for high automation and high human control simultaneously, while some situations may legitimately require stronger control on one side or the other. [32] This responds to a genuine concern that, as automated systems become more capable, humans should not lose the ability to understand, intervene in, override, or terminate processes that affect them. Human control is therefore central to safety, accountability, responsibility, and legitimacy. The present argument accepts that premise but asks whether control, even when genuinely preserved, is sufficient to capture the more formative dimensions of human agency.

This reveals a distinction between operational control and formative agency. Operational control concerns whether a person can intervene in or terminate a system, whereas formative agency concerns whether that person continues to participate meaningfully in the processes through which perception, imagination, preference, and judgment are formed.

This reframing avoids a common mistake in AI discourse. If human control is treated as the ultimate goal, then any system that leaves final authorization to a person may appear adequately human-centered. But aesthetic agency requires attention to what happens before authorization: what the human has encountered, which possibilities have been made available, how preferences have been anticipated, and whether criteria of judgment remain open.

Aesthetically human-centered AI must care not only about who decides at the end, but also about how the human subject develops throughout the process.

7.2. Preserving Spaces of Indeterminacy and Discovery

The first condition of aesthetic agency is the preservation of indeterminacy, understood not as randomness, opacity, or disorder for their own sake, but as the continued availability of encounters whose significance has not been fully specified by prior data or predicted preference. Predictive systems are designed to infer likely interests, rank likely relevance, anticipate likely choices, and optimize likely engagement. In many contexts these functions are useful precisely because they reduce uncertainty. Yet aesthetic experience often depends upon a different relation to uncertainty: an object, idea, or possibility may become valuable only after sustained attention, even though its relevance could not have been established beforehand. If prediction becomes the dominant architecture of experience, such indeterminacy risks appearing merely as a technical failure rather than as a condition of discovery.

This distinction becomes clearer when novelty and serendipity are considered. A recommender system may deliberately introduce unfamiliar material and even optimize a measurable degree of surprise. Such strategies can increase diversity, but engineered novelty is not identical with aesthetic openness. In recommender-system research, serendipity commonly denotes a useful or relevant surprise; aesthetically, however, the concept must remain wider because some consequential encounters are initially neither useful nor obviously relevant. Their significance emerges only through the subject's willingness to attend, interpret, and revise an earlier response. A difficult artwork, an unfamiliar style, or an unexpected conceptual association may matter precisely because its value was not predictable from the subject's established profile.

The design implication is not that AI should maximize randomness, but that it should avoid making every future encounter answerable to a model of the user's past. This requirement is especially important in educational, cultural, and creative contexts, where discovery can depend upon exposure to what initially appears marginal or unpromising. Systems can support such openness by allowing exploratory modes, alternative pathways, non-personalized views, or forms of curation that are not wholly derived from behavioral prediction. The deeper principle is that human-centeredness should preserve a difference between what a system can infer will matter and what a subject may come to find meaningful through experience.

7.3. Designing for Judgment Rather Than Merely Preference

The second condition is the preservation of judgment, because an aesthetically human-centered system should support not only the expression of preference but the reflective formation and revision of the criteria through which preference is evaluated.

Designing for judgment therefore means creating environments in which users are not simply presented with what they are predicted to prefer but are given meaningful opportunities to evaluate why something appears significant and whether the criteria guiding selection deserve reconsideration.

This could include presenting multiple defensible alternatives rather than a single authoritative answer, making uncertainty and disagreement visible, inviting comparison across divergent framings, allowing users to inspect how a result changes under different assumptions, or explicitly distinguishing recommendation from justification.

More fundamentally, however, designing for judgment means treating the user not merely as a source of preference data but as a subject capable of changing the standards according to which preferences are formed.

If a user repeatedly selects a particular style, the system may reasonably infer a preference. But an aesthetically human-centered system should not assume that the optimal response is always to reinforce that preference.

Instantaneous systems tend to compress reflection. Yet some forms of evaluation benefit from delay. Aesthetic judgment can change after repeated encounter, contextual information, or comparison.

A human-centered AI organized around judgment should sometimes support suspension rather than immediate decision.

Accordingly, AI should support the formation of judgment rather than merely accelerate the expression of preference; speed and convenience are human-centered only when they do not systematically displace the reflective capacities that give evaluation its meaning.

7.4. Protecting Perceptual and Imaginative Plurality

The third condition is plurality, understood as the availability of genuinely heterogeneous perceptual and imaginative possibilities rather than the mere numerical abundance of outputs.

If AI increasingly structures visibility and generates possible forms, then human-centeredness must address the range of perceptual and imaginative possibilities made available through those systems.

The relevant question is whether the alternatives presented by a system genuinely open distinct conceptual or aesthetic possibilities, rather than multiplying variations that remain confined within the same assumptions, styles, or statistical center.

Aesthetically, plurality matters because human imagination develops through encounters with difference.

A system may satisfy formal representational diversity while still producing aesthetically standardized outputs. Conversely, stylistic plurality alone does not guarantee justice.

The aesthetic principle therefore adds something distinct: it asks whether the mediated environment preserves enough heterogeneity for subjects to develop discriminating perception and non-redundant imagination.

This is especially important because generative systems are trained on already existing cultural materials. They can recombine and transform those materials in impressive ways, but their outputs remain shaped by distributions within the training corpus and by subsequent alignment, filtering, and product design.

Aesthetic agency therefore requires space for what initially appears implausible, marginal, difficult, or poorly fitted to dominant patterns, since aesthetic development often depends upon learning to perceive significance where immediate prediction would not assign it.

Even failure can contribute to imaginative plurality when it exposes the limits of an established pattern, forces revision, or makes a different possibility perceptible.

A recommendation environment that continually converges on a stable profile may reduce the effective diversity of experience even if the platform contains enormous amounts of content.

Together they preserve one of the most important conditions of aesthetic subjectivity: the possibility that the world exceeds the categories through which it has previously been represented.

7.5. Aesthetic Agency as a Principle of Human-Centered AI

The preceding discussion allows aesthetic agency to be formulated not merely as a philosophical description of subjectivity but as a principle of human-centered AI. Such a principle would require particular attention in systems involved in recommendation, personalization, generation, education, cultural mediation, and creative assistance, because in these domains perception, interpretation, imagination, and judgment are themselves among the capacities being exercised. The principle does not demand that AI withdraw from these processes. Rather, it asks whether technological mediation preserves the subject's capacity to form and transform the sensibility through which the world is encountered and judged.

Aesthetic agency should not be treated as an overriding value that displaces safety, fairness,

accessibility, or other established concerns. Its normative weight is context-dependent. In safety-critical settings, predictability may properly take precedence over exploratory openness; in accessibility contexts, personalization may substantially enlarge rather than restrict agency; and in routine tasks, users may reasonably prefer extensive automation. The principle becomes most demanding where AI participates in informational, cultural, educational, expressive, or creative environments, because there the quality of human-centeredness cannot be assessed solely by accuracy, convenience, or final control. It must also be assessed by whether repeated interaction leaves human capacities for perception, imagination, interpretation, and judgment meaningfully active.

Aesthetic agency can therefore be defined as the preservation of a human subject's capacity for open perception, plural imagination, reflective judgment, and experiential self-transformation within AI-mediated environments. This definition implies a recursive rather than linear model of human–AI interaction: human perception and imagination shape the use of AI; AI mediation reorganizes the field of encounter and interpretation; reflective judgment responds to that field; and the resulting experience may transform the human subject who enters subsequent interactions. The importance of this model is that it refuses to treat the human as a static set of preferences waiting to be served. A genuinely human-centered AI must respect present choices while also preserving the possibility that those choices, tastes, and standards will change through encounters that prediction could not fully anticipate.

Figure 1 visualizes this recursive structure: AI mediation affects the field of encounter and judgment, while the resulting experience can in turn transform the subject who enters subsequent interactions.

A Recursive Model of Aesthetic Agency in Human-AI Interaction




Figure 1. A Recursive Model of Aesthetic Agency in Human–AI Interaction

AI can therefore enlarge aesthetic agency when it opens unfamiliar possibilities, supports experimentation, increases access to cultural difference, reveals alternative perspectives, or helps users articulate judgments they could not previously formulate. It can contract aesthetic agency when responsiveness hardens into repetitive confirmation and when the modeled user becomes the implicit standard governing future exposure. Restoring the aesthetic subject consequently changes the meaning of human-centeredness itself: to center the human is not only to privilege human interests over technical efficiency, but to preserve the perceptual, imaginative, interpretive, and evaluative capacities through which humans continue to form those interests, revise those values, and transform their relation to the world. Human-centered AI must protect not only the human who chooses, but also the human who can become different from the chooser the system has learned to predict.

8. Conclusion: After Human-Centeredness

8.1. The Human Cannot Be Reduced to the User

Human-centered AI has become one of the most influential frameworks for thinking about the relation between artificial intelligence and human values. Its central intuition is compelling: AI should remain oriented toward human purposes rather than becoming an autonomous technological end. Yet this article has argued that the phrase “human-centered” contains an unresolved philosophical problem. The human is not exhausted by the functions performed in relation to a technological system. A person is more than a user because human beings do not simply enter interaction with stable goals, preferences, and values already fully formed. Their ways of perceiving, attending, imagining, interpreting, and judging develop through experience. The human at the center of AI should therefore be conceived not merely as a user of technology but as a developing subject of experience.

8.2. Human Values Presuppose Human Modes of Experiencing and Judging

This reconstruction also has implications for the language of human values, because values become operative only through subjects capable of perceiving situations, interpreting conflicts, imagining alternatives, and judging what matters. Contemporary AI ethics rightly insists that artificial intelligence should respect values such as fairness, autonomy, dignity, privacy, responsibility, and well-being. These principles remain essential to any defensible account of human-centered technology.

The ethical subject therefore cannot be separated from the aesthetic subject: the person whose rights are protected is also the person whose capacities of attention, interpretation, imagination, and judgment make those rights and values experientially meaningful. This is why aesthetics should not be added to HCAI merely as another value alongside fairness, privacy, safety, or autonomy. Human judgment requires perceptual and imaginative conditions within which something can appear meaningful enough to be judged. The aesthetic dimension therefore reveals a precondition that purely procedural accounts of human-centeredness can easily overlook. Before the human can express a value, choose among alternatives, or exercise control over an AI system, a world must already have become available as something worth perceiving, interpreting, and evaluating.

8.3. Aesthetic Agency as a Condition of Meaningful Human–AI Relations

The concept developed throughout this article to describe this deeper dimension is aesthetic agency: the capacity of human subjects to participate meaningfully in the formation of their own perceptual, imaginative, interpretive, and evaluative relation to the world. It includes the ability to encounter what has not been selected in advance, to resist habitual patterns, to remain with ambiguity, to recognize or generate alternatives, to revise taste, and to judge beyond computational prediction. Because technological mediation is unavoidable and can be enabling as well as constraining, the normative question is not whether AI influences these capacities, but whether the form of influence preserves sufficient openness, plurality, reflectivity, and transformability for the subject to develop beyond the profile through which the system models that subject.

This is why final authority alone cannot establish meaningful agency. A person may retain the power to accept or reject an AI recommendation while the system has already organized what is visible, plausible, or imaginable. The earlier distinctions between choice and judgment, preference and taste, generation and imagination, and optimization and experience identify different aspects of the same problem. A system can multiply choices without strengthening judgment, predict preference without supporting the development of taste, generate possibilities without ensuring that human imagination remains actively exercised, and optimize an outcome while diminishing the formative value of the process through which it is reached.

Meaningful human–AI relations therefore require more than the presence of a human at the endpoint of a computational loop. They require forms of interaction in which the human remains capable of participating in the formation of the horizon within which technological assistance

acquires meaning. Aesthetic agency names this requirement by shifting attention from whether a human is formally included to which human capacities remain active throughout the relation. Its preservation is consequently a condition of human-centeredness whenever AI participates not only in executing choices but also in shaping the perceptual and imaginative conditions from which choices and judgments emerge.

8.4. To Center the Human Is to Preserve the Human Capacity to Perceive, Imagine, and Judge

The question with which this article began can now be answered by returning to the ambiguity of the term 'human-centered.' The central challenge is that AI increasingly enters the processes through which humans come to perceive problems, imagine possibilities, form preferences, and exercise judgment. Human-centeredness must therefore be evaluated at a formative as well as an instrumental level. It is not enough that a system responds accurately to expressed preferences or leave a person a final veto if the environment within which those preferences and decisions arise has already been organized around a computational representation that is treated as exhaustive of the person.

An aesthetically expanded HCAI should instead preserve a gap between the predicted self and the possible self. That gap is not a celebration of unpredictability for its own sake, nor a denial that prediction can be useful. It marks the possibility that future experience may exceed the behavioral regularities from which a system learns. A human subject must remain capable of encountering what was not anticipated, imagining what was not generated in advance, revising what was previously preferred, and judging according to criteria that remain open to reflection. When these capacities remain active, AI can function as a medium of exploration and transformation rather than merely as an increasingly precise mechanism for returning the subject to a modeled version of the past.

To center the human in artificial intelligence is therefore not merely to protect human interests, preserve human control, or align systems with existing human preferences and values. It is to preserve the conditions under which human beings can perceive what they were not directed to notice, imagine possibilities that were not supplied for them in advance, revise the standards by which they judge, and become different from the subjects their data predict them to be. The theoretical contribution of the aesthetic subject is precisely to make this formative dimension visible. Human-centered AI should not aspire only to know humans better; it should also preserve the conditions under which humans can revise, exceed, and sometimes contradict what

systems have learned about them. The human at the center of AI must remain not merely the person a system can model, assist, and protect, but the subject who can still encounter what was not predicted, imagine what was not supplied, judge what was not already optimized, and become otherwise.

Footnotes

- [1] Ben Shneiderman, *Human-Centered AI* (New York: Oxford University Press, 2022), especially the discussion of the two-dimensional HCAI framework; Ben Shneiderman, “Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy,” *International Journal of Human–Computer Interaction* 36, no. 6 (2020): 495–504, especially 495–498.
- [2] Immanuel Kant, *Critique of the Power of Judgment*, ed. Paul Guyer, trans. Paul Guyer and Eric Matthews (Cambridge: Cambridge University Press, 2000), §§1–9; John Dewey, *Art as Experience* (New York: Minton, Balch & Company, 1934), chs. 2–3. Kant is cited for judgment that is not mechanically derived from determinate concepts; Dewey for experience as temporally and environmentally constituted rather than reducible to preference satisfaction.
- [3] Don Ihde, *Technology and the Lifeworld: From Garden to Earth* (Bloomington: Indiana University Press, 1990), especially the chapters on human–technology relations; Peter-Paul Verbeek, *What Things Do: Philosophical Reflections on Technology, Agency, and Design*, trans. Robert P. Crease (University Park: Pennsylvania State University Press, 2005), especially Part II on technological mediation.
- [4] Shneiderman, *Human-Centered AI*, discussion of the two-dimensional HCAI framework; Shneiderman, “Human-Centered Artificial Intelligence,” 495–498.
- [5] Dewey, *Art as Experience*, chs. 2–3, especially the discussion of organism–environment interaction and “an experience.”
- [6] International Organization for Standardization, *ISO 9241-210:2019, Ergonomics of Human-System Interaction—Part 210: Human-Centred Design for Interactive Systems* (Geneva: ISO, 2019).
- [7] Joseph A. Konstan and John Riedl, “Recommender Systems: From Algorithms to User Experience,” *User Modeling and User-Adapted Interaction* 22 (2012): 101–123, especially 101–103 and the discussion of user-experience evaluation. The cited claim—that embedding the algorithm in the user experience affects its value and that evaluation

requires measures beyond predictive accuracy—is stated explicitly at the outset of the article.

- [8] Luciano Floridi et al., “AI4People—An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations,” *Minds and Machines* 28 (2018): 689–707, especially 696–700; Virginia Dignum, *Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way* (Cham: Springer, 2019), especially the chapters developing responsibility and human-centered governance. Floridi et al. formulate beneficence, non-maleficence, autonomy, justice, and explicability, with explicability joining intelligibility and accountability.
- [9] Solon Barocas, Moritz Hardt, and Arvind Narayanan, *Fairness and Machine Learning: Limitations and Opportunities* (Cambridge, MA: MIT Press, 2023), especially the chapters introducing group fairness criteria and their normative assumptions; Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan, “Inherent Trade-Offs in the Fair Determination of Risk Scores,” *Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*, Article 43, 43:1–43:23, especially the impossibility results.
- [10] Joshua A. Kroll et al., “Accountable Algorithms,” *University of Pennsylvania Law Review* 165, no. 3 (2017): 633–705, especially 660–682 on transparency, procedural regularity, and technical mechanisms of accountability; Brent Mittelstadt et al., “The Ethics of Algorithms: Mapping the Debate,” *Big Data & Society* 3, no. 2 (2016): 1–21, especially the sections “Map of the Ethics of Algorithms,” “Transformative Effects,” and “Traceability.”
- [11] High-Level Expert Group on Artificial Intelligence, *Ethics Guidelines for Trustworthy AI* (Brussels: European Commission, 2019), especially the sections on human agency and oversight; Shneiderman, *Human-Centered AI*, discussion of human control and human responsibility.
- [12] Immanuel Kant, *Critique of the Power of Judgment*, ed. Paul Guyer, trans. Paul Guyer and Eric Matthews (Cambridge: Cambridge University Press, 2000), §§1–9, especially the account of judgments of taste as grounded in feeling rather than determinate concepts while nevertheless laying claim to universal assent.
- [13] Martha C. Nussbaum, *Love’s Knowledge: Essays on Philosophy and Literature* (New York: Oxford University Press, 1990), especially the essays on literary imagination and moral

- perception; Iris Murdoch, *The Sovereignty of Good* (London: Routledge & Kegan Paul, 1970), especially “The Idea of Perfection.”
- [14] John Dewey, *Art as Experience* (New York: Minton, Balch & Company, 1934), chs. 2–3, especially the distinction between ordinary experience and “an experience” and the account of organism–environment interaction.
- [15] Pierre Bourdieu, *Distinction: A Social Critique of the Judgement of Taste*, trans. Richard Nice (Cambridge, MA: Harvard University Press, 1984), especially Part I on the social conditions of taste; Dewey, *Art as Experience*, chs. 2–3.
- [16] Alexander Gottlieb Baumgarten, *Aesthetica*, 2 vols. (Frankfurt an der Oder: Johann Christian Kleyb, 1750–1758), §1. See also Alexander Gottlieb Baumgarten, *Metaphysics: A Critical Translation with Kant’s Elucidations, Selected Notes, and Related Materials*, trans. and ed. Courtney D. Fugate and John Hymers (London: Bloomsbury, 2013). Baumgarten famously formulated aesthetics as *scientia cognitionis sensitivae*, the science of sensible cognition.
- [17] Immanuel Kant, *Critique of the Power of Judgment*, ed. Paul Guyer, trans. Paul Guyer and Eric Matthews (Cambridge: Cambridge University Press, 2000), §§1–9, 18–22. For the broader distinction between determining and reflecting judgment, see the First Introduction and the published Introduction, especially §§IV–V.
- [18] John Dewey, *Art as Experience* (New York: Minton, Balch & Company, 1934), chs. 2–3, especially the account of “an experience,” doing and undergoing, and the temporal integration of experience.
- [19] Maurice Merleau-Ponty, *Phenomenology of Perception*, trans. Donald A. Landes (London: Routledge, 2012), Preface and Part One, especially the analysis of the lived body as the medium of perceptual access to the world.
- [20] Don Ihde, *Technology and the Lifeworld: From Garden to Earth* (Bloomington: Indiana University Press, 1990), especially the analysis of human–technology relations; Peter-Paul Verbeek, *What Things Do: Philosophical Reflections on Technology, Agency, and Design*, trans. Robert P. Crease (University Park: Pennsylvania State University Press, 2005), Part II on technological mediation.
- [21] Dietmar Jannach, Markus Zanker, Alexander Felfernig, and Gerhard Friedrich, *Recommender Systems: An Introduction* (Cambridge: Cambridge University Press, 2010), especially the chapters on evaluation and user-oriented recommendation; Joseph A. Konstan and John Riedl, “Recommender Systems: From Algorithms to User Experience,”

User Modeling and User-Adapted Interaction 22 (2012): 101–123, especially 101–103 and the discussion of user-experience evaluation.

- [22] Byung Cheol Lee and Jaeyeon (Jae) Chung, “An Empirical Investigation of the Impact of ChatGPT on Creativity,” *Nature Human Behaviour* 8 (2024): 1906–1914, especially the Results and Discussion. Across five experiments, the authors report that ChatGPT assistance increased the creativity of ideas relative to no technology or conventional web search, with stronger effects for incremental than radically new ideas.
- [23] Jack McGuire, David De Cremer, and Tim Van de Cruys, “Establishing the Importance of Co-Creation and Self-Efficacy in Creative Collaboration with Artificial Intelligence,” *Scientific Reports* 14 (2024): Article 18525, especially the Results and Discussion. The study shows that the creative effects of generative-AI collaboration depend importantly on the organization of human–AI interaction and perceived co-creation.
- [24] Michelle Vaccaro, Abdullah Almaatouq, and Thomas W. Malone, “When Combinations of Humans and AI Are Useful: A Systematic Review and Meta-Analysis,” *Nature Human Behaviour* 8 (2024): 2293–2303, especially 2295–2299. The meta-analysis of 106 experimental studies reports negative average synergy for decision tasks but significantly more favorable effects for open-ended creation tasks.
- [25] Qazi Mohammad Areeb et al., “Filter Bubbles in Recommender Systems: Fact or Fallacy—A Systematic Review,” *WIREs Data Mining and Knowledge Discovery* 13, no. 6 (2023): e1512, especially the review findings and the sections on mitigation and explainable recommender systems. The review finds evidence of filter-bubble effects while emphasizing multiple contributing biases and the potential of diversity, transparency, explainability, and user control to mitigate them.
- [26] Immanuel Kant, *Critique of the Power of Judgment*, ed. Paul Guyer, trans. Paul Guyer and Eric Matthews (Cambridge: Cambridge University Press, 2000), especially the Introduction, §§IV–V, and §§1–9. Kant's distinction between determining and reflecting judgment is important here because reflective judgment confronts particulars for which an adequate universal is not simply given in advance.
- [27] Hannah Arendt, *Lectures on Kant's Political Philosophy*, ed. Ronald Beiner (Chicago: University of Chicago Press, 1982), especially 42–45, 70–72. Arendt develops Kantian judgment through the ideas of enlarged mentality and the capacity to consider possible standpoints of others.

- [28] Raja Parasuraman and Dietrich H. Manzey, “Complacency and Bias in Human Use of Automation: An Attentional Integration,” *Human Factors* 52, no. 3 (2010): 381–410, especially the sections synthesizing automation complacency, automation bias, and their attentional mechanisms; Mary L. Cummings, “Automation Bias in Intelligent Time Critical Decision Support Systems,” in *AIAA 1st Intelligent Systems Technical Conference* (Reston, VA: American Institute of Aeronautics and Astronautics, 2004). These studies concern automation rather than generative AI specifically and are used only to support the broader claim that formal human authority does not eliminate overreliance on automated recommendations.
- [29] Pierre Bourdieu, *Distinction: A Social Critique of the Judgement of Taste*, trans. Richard Nice (Cambridge, MA: Harvard University Press, 1984), especially Part I on the social conditions and class structuring of taste. Bourdieu is used here as a corrective to any conception of aesthetic autonomy as the expression of a socially unconditioned individual subject.
- [30] Kent F. Hubert, Kim N. Awa, and Darya L. Zabelina, “The Current State of Artificial Intelligence Generative Language Models Is More Creative Than Humans on Divergent Thinking Tasks,” *Scientific Reports* 14 (2024): Article 3440, especially the Results and Discussion. The study supports the limited methodological point that aesthetic agency should not be grounded in an assumption of categorical human superiority on conventional divergent-thinking measures.
- [31] John Dewey, *Art as Experience* (New York: Minton, Balch & Company, 1934), chs. 2–3, especially the account of “an experience,” doing and undergoing, and temporal development. Dewey is cited here for the claim that aesthetic quality concerns the organization of an experience rather than isolated moments of immediate satisfaction.
- [32] Ben Shneiderman, *Human-Centered AI* (New York: Oxford University Press, 2022), discussion of the two-dimensional HCAI framework; Ben Shneiderman, “Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy,” *International Journal of Human–Computer Interaction* 36, no. 6 (2020): 495–504, especially 495–498. Shneiderman’s framework is relevant because it explicitly rejects a simple trade-off between automation and human control.