

Rethinking Foundation Courses in Design Education under AIGC: A Perception-Judgment-Articulation Framework

Dong Wang^{1*}

¹ College of Art and Design, Southwest Forestry University, Kunming, China

*Correspondence: Dong Wang

Received: 30 June 2026; Accepted: 31 July 2026; Published: 17 August 2026

Abstract:

Generative tools have changed the economics of visual production, and foundation courses in design feel it first, since what they aim to produce — competent visual form — is now what such tools supply on demand. The risk is not poor work, but work that is perfectly adequate and increasingly alike, made with less and less judgment, while the usual responses of adding the tools or banning them leave the goals and assessment of the course unexamined. Building on constructive alignment, this paper proposes an integrated Perception–Judgment–Articulation (P-J-A) framework for developing design judgment, treated as four trainable competencies: perceptual discrimination, problem reframing, intent externalization, and solution convergence. Generative tools bear on these unevenly, wearing away the first and third as production becomes effortless while pushing harder on the second and fourth, and the framework is built to protect the one pair and extend the other. A Prompt–Variant–Rationale (P-V-R) unit ties indicators, tasks, and evidence into a single teaching–learning–assessment loop, with a generation provenance portfolio carrying the evidence of process. The curriculum runs across three modules, worked through in the typesetting of a graduation thesis.

Keywords:

Design Foundation Courses, AIGC, Design Judgment, P-J-A Framework, Process Evidence, Teaching Reform

1. Introduction

1.1. Research Background and Problem Statement

Generative tools have changed the basic economics of making things visual. A concept sketch that used to take an afternoon now arrives in seconds, and the number of options a student can put on the table is no longer capped by the hours available to draw them. These changes reach the classroom through the students first, well before any syllabus has caught up. National policy has moved the same way [1,2], but it can only point a direction; it does not tell us what a foundation course should now teach.

Foundation courses feel this first. Composition, graphic ideation, color, layout are usually the first professional courses a student meets, long organized around producing competent visual form, with the marking settling on whether the hierarchy

holds, the spacing is resolved, the conventions observed. So students put most of their effort into getting the work correct.

The trouble is that competent form is no longer scarce. Once an acceptable piece can be had without exercising the capability it was meant to demonstrate, it stops being evidence that anyone has learned anything: a student can meet every formal requirement and still not say what problem the work solves, what alternatives were weighed, or what the chosen version gave up. Generative help tends to lift the quality of the individual piece even as it flattens the variety across a class [3], so the likely result is work that is perfectly adequate, increasingly alike, and made with steadily less judgment.

The real task of reform, then, is neither to bolt generative tools onto the syllabus nor to ban them, but to rebuild the course itself so that content, teaching, and assessment all work toward judgment, and that judgment leaves a trail that can be gathered up, looked at, and fed back into the teaching.

1.2. Research Gap

It is well established that design judgment is built up rather than innate, and that externalizing ideas in visible form is central to building it [4,5,6], and a newer strand of work records how generative tools are entering design curricula [7,8,9]. Even so, three gaps remain once the focus narrows to foundation courses.

The first is that most discussion gathers around studio and applied courses, where the expected outcome is a design proposal. Foundation courses draw less notice, though the problem bites harder there: what they aim to produce, competent visual form, is exactly what generative tools now hand over for free. What such a course is even for, once production is cheap, has not been worked through.

The second is that reform proposals often reach for phrases like critical judgment or aesthetic discrimination but leave them as broad hopes. What is missing is a breakdown into competencies that fit foundation-level learning, concrete enough to guide how tasks are set and work is judged.

Third, even where tool experiments exist, no classroom-operable framework places judgment at the center and integrates content, pedagogy, and assessment as a system. Assessment in particular remains unreconstructed: the dominant response addresses integrity through detection and restriction [10] rather than the validity of what product-focused assessment can still demonstrate.

1.3. Research Aims and Questions

To close these gaps, the paper puts forward an integrated Perception–Judgment–Articulation (P-J-A) framework for developing design judgment in foundation courses, one that runs a course through goal specification, task organization, and evidence-based feedback. It works through three questions. How can design judgment be defined and broken down into competencies that are observable and can be trained? How can content, teaching, and assessment be made to pull together? And how can the framework be made to run in practice, through the design of modules and tasks, so that the competencies develop in steps?

2. Foundations and Prior Attempts

2.1. Design Judgment and Visual Externalization

Research on design cognition has long held that drawing is part of thinking, not just a record of it: Schön described design as a reflective conversation in which a drawing states an intention and answers back with something unplanned [4], Goldschmidt noted how even a rough sketch carries an inference [5], and Cross gathered these habits under the heading of designerly ways of knowing [11].

Two points matter for teaching: putting an idea into visible form is a condition for thinking, not its last step [6], and working several alternatives in parallel tends to produce stronger, more varied results than polishing one [12]. Both concern how judgment is formed rather than how a finished piece is produced — a difference that matters far more once the piece is cheap. Goldschmidt once argued that manual sketching keeps affordances digital tools do not replace [13]; this paper asks the same under generative conditions.

2.2. Generative AI in Design Education: Current Attempts

Recent studies have begun to record how generative tools are entering design curricula, in poster workshops, in support for student creativity, and in curricula that carry the tools from foundational into applied courses [7,8,9].

Meanwhile evidence has been building about what these tools do to the work itself. Doshi and Hauser found that generative help made individual pieces more creative but left a group's output less diverse [3], and other studies report the same pull toward homogeneity [14,15]. For design education the danger is thus not poor work, but work that is perfectly acceptable and increasingly alike, made with less and less judgment.

2.3. Limitations of Existing Attempts

Constructive alignment asks that outcomes, activities, and assessment pull in the same direction, and that assessment yield traceable evidence for planning [16,17]. In the creative disciplines this has always been hard, since judgments of creative work lean on the finished piece and an experienced eye rather than on stated criteria [18]. Three problems keep recurring. New tools arrive while the logic of the tasks stays as it was, so a course looks current while assessing what it always assessed. Judgment is named as a goal but never built into the task, so the quickest path to an acceptable submission still calls for the least of it. And most seriously, assessment itself has been left untouched: the response to generative tools has been about integrity — detecting, restricting, setting policy [10] — rather than validity. Yet a product-focused assessment no longer supports the conclusion it once did, since a submission not asked to show its criteria, its discarded alternatives, or its path of revision leaves little on which to judge whether judgment took place [19]. The field has long known assessment to be its weak point [20] and called for evaluation tied to process as much as outcome [21]; what it lacks, for foundation teaching in particular, is a course-level mechanism linking goals, tasks, and evidence.

3. Conceptual and Theoretical Foundations

3.1. Repositioning Design Foundation Courses

Foundation courses in design have usually been treated as technical groundwork. Yet the research on design cognition points the other way, suggesting that putting something in visible form is part of how thinking happens, not merely a way of writing it down after the fact [4,5].

3.2. The Principle of Constructive Alignment in This Study

The paper builds on constructive alignment [16] in three moves: outcomes are stated as observable performances of competency, which organize the content; activities are arranged around tasks that pull students' grounds for judgment into the open while they work; and assessment gathers evidence of whether that judgment is happening, feeding it back to adjust content and tasks [17].

Applying the principle directly is not straightforward, since the same making task can serve different competency goals depending on the instructor's intent; unless that intent is written into the task requirements and the evidence collected, a course piles up finished assignments without a clear picture of which competencies improved. A further complication, taken up in Section 4.2, concerns what a finished artifact can still be taken to prove.

3.3. A Four-Dimensional Model of Design Judgment

Drawing on the work above and on what foundation-level learning involves, this paper treats design judgment as four competencies that work together.

The first is perceptual discrimination: a student picks out differences in quality among closely related options and can say in concrete terms what they are, the point being the specificity, not whether the preference is the "right" one. A student who can tell you the weight contrast between two levels of text is too slight to hold them apart has discriminated; one who can only report that a version looks better has not.

The second is problem reframing. Rather than take a brief at face value, the student digs into what it leaves unsaid and redraws the edges of the problem. An instruction to typeset a document spells out the visible result but not the harder requirement — that it keep working when the text is edited, so citations, notes, and tables go on behaving as the document changes underneath them. Reframing brings that to the surface, and with it the gap between a document that merely looks right and one built to hold.

The third is intent externalization, turning a still-forming intention into something that can be looked at and tested — a written note of what one is after and under what constraints, an annotated sketch, a logged sequence of instructions to a generative tool — with each move tied to an assumption the student can name. Someone who reworks an instruction thirty times but cannot say which assumption changed from one try to the next has generated output without ever externalizing an intent.

The fourth is solution convergence, treating a solution as something worked toward through comparison. The student sets up criteria, generates options that genuinely differ, settles on one, and can say what settling on it gave up. Picking the most appealing image from a large batch of machine-generated candidates is not convergence; it hands the judgment to whichever option catches the eye first.

In practice the four run together: discrimination gives comparison something to stand on, reframing settles what it is for, externalization keeps the reasoning open to inspection, and convergence carries it to a defensible decision. What is worth noticing is that generative tools do not touch them evenly. Discrimination and externalization

tend to wither when production costs nothing, since a student can get an acceptable result without either; reframing and convergence are pushed harder, since a bigger field of candidates raises the demand for criteria and trade-offs. A framework built around judgment therefore has two jobs: to protect the first pair and stretch the second.

4. Building the Integrated P-J-A Framework for Design Foundation Courses

4.1. Design Principles

Four principles hold the framework together. Instructional fit keeps the training of judgment anchored to specific visual concepts and to work students can produce within a class. A judgment orientation insists on performances that can be observed and assessed rather than letting “judgment” stand as a slogan. Integrated linkage lines content, teaching, and assessment up behind the same goals. And operability guides module and evidence design without a fixed script, treating differences in students’ preparation and tool access as constraints to design around.

4.2. Overall Structure

Figure 1 shows the P-J-A framework as a system in which perception, judgment, and articulation work in concert, not a line of steps in order. Perception supplies the discriminations judgment works on, judgment produces the choices articulation must account for, and articulation turns those choices into evidence that feeds back into content and task design.




Figure 1. Overall structure of the integrated P-J-A framework for design foundation courses.

At the center sits design judgment, split into the four competencies, which organize the content and give task design and assessment a shared set of targets. The arrows carry two ideas: all three layers aim at the same goals, and how students perform flows back to shape which exercises are kept and where assessment puts its weight. A

course run this way stops being a fixed sequence and becomes something that recalibrates around the competencies it is trying to build.

4.3. Design Logic of the Three Layers

At the level of content, the framework keeps or drops a manual exercise by a single test: it stays when it trains a capability that is still scarce, not because it is traditional [13]. On that test, exercises that sharpen fine perceptual discrimination gain value, since they train the eye that will later weigh what a machine produces, while those whose whole point was a clean rendering of a fixed result lose their warrant. Part of the old abstract sequence gives way to contexts along a gradient of difficulty, each with explicit constraints, so bare production commands become problems calling for a judgment: the instruction shifts from what to make toward what to decide.

At the level of teaching, tasks are arranged so students' choices become visible, across three phases. The tools are first held back while instruction works on discrimination; once that is steady, students make one solution by hand, obtain several by generative means, and set them side by side [12], with the instructor's questions pushing preference toward something checkable — not which is better, but on what criterion and at what cost. In the last phase the briefs open up and draw the competencies together.

At the level of assessment, the focus moves off whether an artifact is competent and onto whether students can say what problem it solves, where alternatives were weighed, why one was chosen, and what it gave up. Judging creative work has long leaned on the finished piece and the tacit read of an expert eye [18], and generative tools have mostly prompted worry about integrity rather than validity [10]. The framework keeps the finished work but adds a generation provenance portfolio to carry the evidence of process — the intent statement, the variants with sources marked, the comparison criteria, the discarded options with reasons, the traces of manual intervention, and a short reflection — showing whether the path from problem to revision rests on grounds that can be followed. Section 5 fills in each layer in turn.

4.4. Core Mechanism: The P-V-R Unit

The three layers show how content, teaching, and assessment fit together in principle. To run them in a classroom, the framework uses a smaller building block: the Prompt–Variant–Rationale (P-V-R) unit, where competency goals become something a teacher can carry out.

Here prompt is not restricted to input to a generative model but denotes a written statement of intent and constraints that precedes making, whatever the medium, so the unit applies where generative tools are withheld and where they are central. It treats a single activity as three linked elements: the prompt externalizes intent; the variant set requires several substantively different solutions — differing in structure, hierarchy, or problem definition, not detail or style; and the rationale states the comparison criteria, the selected option, and the discarded ones with reasons.

Take a unit in Layout Design set in the typesetting of a thesis, where the brief asks not just for a document that looks right but one that stays right after editing — citations still pointing to the correct reference once the list is reordered, notes staying at the foot of their page as pagination shifts. Table 1 lays out how indicator, task, and evidence connect across the three parts.

Table 1. *Structure of the P-V-R unit in a design foundation task.*

Element	Competency indicator	Task requirement	Evidence produced	What the evidence shows
Prompt	States intent and constraints in checkable terms	Write a brief before making: users, viewing condition, the constraint in tension	Written intent statement	Whether the problem is reframed beyond the brief as given
Variant	Generates substantively distinct options	Produce four options differing in hierarchy or problem definition; mark the source of each	Annotated variant set	Whether difference is structural rather than decorative
Rationale	Sets criteria and accounts for trade-offs	State comparison criteria; select one; explain what each discarded option gave up	Comparison notes and discard record	Whether the choice rests on stated grounds

Several P-V-R units with related themes form a module; modules are sequenced from lower to higher cognitive demand and together constitute one semester, as detailed in Section 5.

5. Instructional Elaboration: Operational Examples of the Framework

5.1. Teaching Context and Course Positioning

The worked example is planned for design-major students in Layout Design and uses the typesetting of a graduation thesis, across three settings of growing scope: a single data table, the cross-references and notes running through the document, and the report cover. The context comes with constraints one can check, it is part of students' immediate experience, and generative tools are already heavily used here, so the gap this paper is about shows up in its sharpest form. These students have visual training but no background in programming or structured documents, so an AI-generated document correct on the surface opens a systematic gap in their judgment. The tasks stay foundation tasks — reading a problem in the layout, producing evidence, comparing — not finished publications.

5.2. Module Structure and Coordinated Implementation

The course splits into three modules, each built around the four competencies and coordinated across the three layers (Table 2).

Table 2. *Modular coordination logic of the P-J-A framework in a design foundation course.*

Core module	Competency focus	Content	Pedagogy	Assessment
M1: Perceptual training and controlled comparison	Perceptual discrimination	Alignment, gray value, reading flow; a thesis data table; controlled variant sets, generative tools withheld	Demonstration; the squint-and-step-back method; peer checking	Annotated fault-finding sheets; the student's own controlled set
M2:	Problem	Static appearance	Workshop; the	Intent

Interrogating a generated document	reframing; intent externalization	versus maintained structure; cross-references and notes in a thesis; an AI-produced document that breaks on edit	questioning chain; edit-and-observe	statements; break-point annotations; repair logs
M3: Choosing among defensible options	Solution convergence, integrating the preceding three	Matching form to content; a report cover with several defensible layouts	Project-based learning; defend-and-discard critique	Generation provenance portfolio; the discard record with reasons; reflection

5.3. Representative Teaching Contexts and Task Examples

5.3.1. Module 1 Example: Reading a Data Table

The first example works on discrimination before any generative tool comes into play, using a single data table from a thesis. An untrained eye, and the tools built to serve it, judge such a table geometrically — if the left edges line up, it looks clean. The module trains a functional standard instead: whether the eye can move down the table without friction.

5.3.1.1. Competency indicators.

The indicators come out of perceptual discrimination. A student should name a specific fault in terms of reading — that a column is left-aligned so the decimal points fail to line up, making the reader re-judge every number row by row — and, asked to squint and step back, should say where a page reads too heavy or too loose. What the unit assesses is the naming of the fault, not the preference.

5.3.1.2. Task design and classroom enactment

The instructor hands over a table lifted from an AI-typeset draft — a column of figures like 12.3456, 128.9, 0.567, 99.0, left-aligned into a tidy block — and asks only what is wrong with it. Most students say nothing is. The instructor then switches the column to decimal alignment, the figures drop onto a single vertical line, and the size of each becomes readable from the width of the number block alone, before a digit is read. That shift, from “looks clean” to “reads wrong,” is the discrimination the exercise is after. Students then squint and step back from a whole page to hunt more faults of the same family — margins asymmetric at the binding gutter, gray value patchy where mixed Chinese and Latin type has not been spaced — saying in each case what it does to the reading, not just what it looks like.

The evidence is the annotated fault-finding sheet and the corrected table, weighed against Table 3. The moment that tells you the unit is working is the student who leans toward the screen and asks how they had not seen the fault before — that turn, from looking at the content to looking at the reading, is what the unit is after.

5.3.2. Module 2 Example: Interrogating a Generated Document

The focus now moves to problem reframing and intent externalization, using the apparatus of an AI-typeset thesis — in-text citations and footnotes — which looks correct at first glance. The module turns on a recurring failure: the tool produces a document right as a static picture and students take it as finished, but it is really a

system meant to hold up while edited, and the moment anything moves the picture comes apart. The reframing it trains turns “a result that looks correct” into “a structure I must understand well enough to change.”

5.3.2.1. Competency indicators.

Students write a statement of intent for what the document must still do after editing — a citation that keeps pointing to the right reference once the list is reordered, a note that stays at the foot of its page as pagination moves — then predict where the AI version will break. The indicator is met not where a student fixes the document, but where they can say what a citation or note actually is: a relation that gets maintained, not a number that got typed.

5.3.2.2. Task design and classroom enactment.

Working in small groups, students edit an AI-typeset draft on purpose — reorder two references, drop in a paragraph that forces the pages to reflow — and watch what gives way: the in-text numbers stop matching, and a footnote the tool wrote as superscript text at the end of the file stays stuck there instead of moving to the foot of its page. The instructor does not hand over the fix. When a student who has a working result cannot say why a step was needed, the instructor works through a chain: what is this mark; what is it meant to do; go find out, come back and tell me, and then I will ask you something else. The rule underneath is simple — having got the result to work, the student is now more obliged to know why it works, not less. From there students rebuild the apparatus with the maintained versions — a cross-reference field that follows its target, a header row set to repeat, a footnote anchored to the page — logging which faked relation each break had exposed.

The evidence is the statement of intent, the predictions of where things would break, and the repair log, weighed against Table 3. As the term goes on, this repair stops being something students carry to the instructor and becomes something they settle themselves — and that change, from copying a result to operating the structure behind it, is the evidence the unit is after.

5.3.3. Module 3 Example: Choosing Among Defensible Covers

The last example brings the earlier competencies together on a problem with no standard answer: a report cover — say, for a survey of the incremental renovation of an aging residential district — A4 portrait carrying a title, subtitle, department, and name. Several layouts can be defended, and students circle back to the same three: a left-aligned grid that looks professional, a centered arrangement that feels safe, and a rotated, color-blocked composition that looks creative. The module shifts them from choosing by how a cover looks to how well form fits content, accounting for what each dropped option cost.

5.3.3.1. Competency indicators.

The indicators rest on three things a student does: generate several layouts that genuinely differ; for each, name not just its appeal but its blind spot — the grid picked only because it reads as high-end, the centered version because it feels safe, the rotated version only to look unlike everyone else’s; and converge by throwing options out for reasons of fit rather than taste, so that a ninety-degree rotation obstructing first reading goes on functional grounds and a dark color block that flattens under the

black-and-white printing many archives require goes on grounds of reproduction. What gets assessed is whether the final choice rests on a defensible relation between the layout and what the report is about.

5.3.3.2. *Task design and classroom enactment.*

Each student lays the options out as overlaid drafts, then settles on one and argues, out loud, for dropping the rest. The telling case is the student who backs into the safe centered option, dismissing the grid as “too empty” and the rotated version as “too busy” without being able to say where. The instructor leans on exactly those words — empty where? busy where? — and insists the discards be defended before the choice, so the most comfortable option cannot win just by asking the least of the student. Only then does the reasoning get concrete: the centered layout reads as closed and authoritative, which sits badly against a subject about intervention and repair, while the grid reads as planning and reason, its white space as room for what might yet happen. Each student hands in a short reflection. The aim is what one student put into words at the end — that it had never really been about which looked better, but which was right for the content.

The evidence is a generation provenance portfolio built around the discard record: for every layout let go, the reason it was let go, in terms of fit rather than taste. Feedback happens through a critique in which students defend their discards to peers, since hearing how someone else gave up an appealing option resets a student’s own sense of the criteria. Table 3 gives the reference criteria used across all three modules.

Table 3. *Assessment reference for the generation provenance portfolio.*

Competency	Primary evidence	Adequate performance	Common shortfall
Perceptual discrimination	Annotation and ranking records	Names a specific relation and its effect	Records preference without grounds
Problem reframing	Intent statement	Identifies a need the brief omits	Restates the brief in other words
Intent externalization	Iteration log	Each round changes a stated assumption	Repeated rewording without stated purpose
Solution convergence	Retained options and discard record	Options genuinely distinct; discards explained by fit	Options differ in detail; discards unexplained

6. Discussion

6.1. *Theoretical Contributions*

The paper argues that foundation courses can move beyond producing competent form to become a setting in which judgment is developed, made visible, and corrected. Its contribution is to carry that idea, long present in the design-cognition literature [4,5,13], into course design: it defines judgment as four competencies that generative tools bear on unevenly — two eroded by effortless production, two amplified by a larger field of candidates — and, because a finished artifact no longer proves the capability was exercised, makes articulation a layer of its own, run through the P-V-R unit.

6.2. *Implications for Teaching Practice*

The framework brings out a pattern: keep the output format the same but put it inside a different task structure, and what students do changes a great deal. The data-table unit moves attention off whether content is correct and onto whether it reads; the document-apparatus unit uses deliberate editing to expose the gap between a faked appearance and a maintained structure; the cover unit makes students set their own criteria and answer for what each dropped option cost. Reform, then, turns not on swapping out content or bringing in tools but on redesigning the task, so exercises a course already runs do more cognitive work.

What assessment attends to shapes how students work: check only whether the artifact is competent and they take the route asking the least judgment of them; look at the evidence of process and they keep their revision traces and spell out their trade-offs. It also puts integrity in a different light, since a student who has exercised no judgment cannot produce a convincing provenance record — so integrity becomes a byproduct of assessing the right thing rather than a separate problem to be policed [10].

6.3. Implementation Challenges and Strategies

Two practical concerns come up at once. Time can be reorganized rather than added to, since the hours once spent demonstrating technique free up once practice runs alongside real problems; and the sharper diagnostic judgment the framework asks of the instructor is supported by the criteria in Table 3 and the evidence checklists, which anchor feedback to explicit criteria.

A third difficulty belongs to the generative setting specifically: students do not have equal access to the tools, which increasingly sit behind subscriptions, so that rewarding the variant set as such would favor those who can pay for reasons unrelated to competency. The framework works against this by withholding the tools in Module 1, so discrimination is trained on level ground; by assessing the intent statement and rationale rather than how polished the variants are; and by defining a substantive difference at the level of structure, which more spending cannot buy. Even so, instructors should state the tool baseline they assume, allow hand-made variants throughout, and try the framework on a small scale first.

7. Conclusions

Reforming a design foundation course is less about adding generative tools or forbidding them than about redesigning tasks and assessment, so students produce — and keep — a traceable record of the judgments behind their work. The paper has set out an integrated Perception–Judgment–Articulation framework, defined design judgment as four trainable competencies, and introduced the Prompt–Variant–Rationale unit as the means of putting it into practice, worked through in three modules set in the typesetting of a graduation thesis.

The elaboration is built on framework logic and teaching experience; it shows internal coherence but cannot guarantee effective operation once implemented. The principal limitation is the absence of classroom-based empirical data: whether students demonstrate the intended discriminations, whether the four competencies develop in sequence, and whether the provenance portfolio in fact distinguishes students who have exercised judgment from those who have not all require verification through classroom evidence. Future work therefore needs small-scale pilots — selecting one or two P-V-R units, implementing them across two course

iterations, and tracking competency development through pre- and post-comparisons together with process records. Beyond piloting, scaling will depend on the practicality of assessment: further work is needed to develop structured tools — standardized rubrics, peer-review guidelines, a portfolio scoring framework — and to test their reliability, validity, and fairness across cohorts with differing access to generative tools.

Conflicts of Interest

The author declares that there is no conflict of interest regarding the publication of this article.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

References

- [1] Ministry of Education of the People's Republic of China, et al. Guideline on Accelerating the Digitalization of Education. 2025. Available online: http://www.moe.gov.cn/srcsite/A01/s7048/202504/t20250416_1187476.html (accessed on 29 July 2026).
- [2] Ministry of Education of the People's Republic of China, et al. "AI + Education" Action Plan. 2026. Available online: http://www.moe.gov.cn/srcsite/A16/s3342/202604/t20260410_1433240.html (accessed on 29 July 2026).
- [3] Doshi, A.R.; Hauser, O.P. Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Sci. Adv.* 2024, 10(28), eadn5290, DOI: 10.1126/sciadv.adn5290.
- [4] Schön, D.A. *The Reflective Practitioner: How Professionals Think in Action*; Basic Books: New York, NY, USA, 1983; ISBN 9780465068784.
- [5] Goldschmidt, G. The dialectics of sketching. *Creat. Res. J.* 1991, 4, 123-143, DOI: 10.1080/10400419109534381.
- [6] Novica, D.R.; Wianto, E.; Andrade Campos, S. Drawing and ideation process in design education: A systematic literature review. *Cogent Arts Humanit.* 2023, 10(1), 2219487, DOI: 10.1080/23311983.2023.2219487.
- [7] Hwang, A.H.-C. Graphic design education in the era of text-to-image generation: Transitioning to contents creator. *Int. J. Art Des. Educ.* 2025, DOI: 10.1111/jade.12558.
- [8] Lee, C.W. Application of generative artificial intelligence in design education: An exploration and analysis to enhance student creativity. *Eng. Proc.* 2025, 98(1), 29, DOI: 10.3390/engproc2025098029.
- [9] Strategic applications of generative AI in design education. *Eng. Proc.* 2025, 120(1), 56, DOI: 10.3390/engproc2025120056.

- [10] Evangelista, E.D.L. Ensuring academic integrity in the age of ChatGPT: Rethinking exam design, assessment strategies, and ethical AI policies in higher education. *Contemp. Educ. Technol.* 2025, 17(1), ep559.
- [11] Cross, N. *Designerly Ways of Knowing*; Springer: London, UK, 2006; ISBN 9781846283000.
- [12] Dow, S.P.; Glassco, A.; Kass, J.; Schwarz, M.; Schwartz, D.L.; Klemmer, S.R. Parallel prototyping leads to better design results, more divergence, and increased self-efficacy. *ACM Trans. Comput.-Hum. Interact.* 2010, 17(4), 18, DOI: 10.1145/1879831.1879836.
- [13] Goldschmidt, G. Manual sketching: Why is it still relevant? In *The Active Image: Architecture and Engineering in the Age of Modeling*; Ammon, S., Capdevila-Werning, R., Eds.; Springer: Cham, Switzerland, 2017; pp. 77-97, DOI: 10.1007/978-3-319-56466-1_4.
- [14] Anderson, B.R.; Shah, J.H.; Kreminski, M. Homogenization effects of large language models on human creative ideation. In *Proceedings of the 16th Conference on Creativity & Cognition (C&C '24)*, Chicago, IL, USA, 23-26 June 2024; ACM: New York, NY, USA, 2024; pp. 413-425, DOI: 10.1145/3635636.3656204.
- [15] Agarwal, D.; Naaman, M.; Vashistha, A. AI suggestions homogenize writing toward Western styles and diminish cultural nuances. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*, Yokohama, Japan, 26 April-1 May 2025; ACM: New York, NY, USA, 2025; pp. 1-21, DOI: 10.1145/3706598.3713564.
- [16] Biggs, J.B.; Tang, C. *Teaching for Quality Learning at University: What the Student Does*, 4th ed.; Open University Press: Maidenhead, UK, 2011; ISBN 9780335242757.
- [17] Loughlin, C.; Lygo-Baker, S.; Lindberg-Sand, Å. Reclaiming constructive alignment. *Eur. J. High. Educ.* 2021, 11(2), 119-136, DOI: 10.1080/21568235.2020.1816197.
- [18] Onsmann, A. Assessing creativity in a “New Generation” architecture degree. *Think. Skills Creat.* 2016, 19, 210-218, DOI: 10.1016/j.tsc.2015.07.001.
- [19] Razzouk, R.; Shute, V. What is design thinking and why is it important? *Rev. Educ. Res.* 2012, 82, 330-348, DOI: 10.3102/0034654312457429.
- [20] Brosens, L.; Raes, A.; Octavia, J.R.; Emmanouil, M. How future proof is design education? A systematic review. *Int. J. Technol. Des. Educ.* 2023, 33(2), 663-683, DOI: 10.1007/s10798-022-09743-4.
- [21] Oo, T.Z.; Kadyirov, T.; Kadyjrova, L.; Józsa, K. Design-based learning in higher education: Its effects on students’ motivation, creativity and design skills. *Think. Skills Creat.* 2024, 53, 101621, DOI: 10.1016/j.tsc.2024.101621.



© 2026 by the author(s); licensee International Technology and Science Publications (ITS), this work for open access publication is under the Creative Commons Attribution International License (CC BY 4.0). (<http://creativecommons.org/licenses/by/4.0/>)