MULTIMODAL SIGNS, CLINICAL MEANING, AND THE LIMITS OF ARTIFICIAL CREATIVITY: A THEORETICAL EXPLORATION

George Damaskinidis
University of Western Macedonia, Greece
damaskinidis@hotmail.com

Abstract

Published medical discourse has explicitly used the language of "creative AI" and described some model outputs as creative. This study offers a conceptual contribution from semiotics to that debate rather than an account based on first-hand clinical practice. It synthesises work on creativity, multimodality, medical semiotics and clinical reasoning to ask what kind of novelty medical AI can produce and what remains specific to accountable human interpretation. Using Boden's distinction between combinational, exploratory and transformational creativity, the study argues that current systems can support combinational and exploratory novelty within learned

Digital Age in Semiotics & Communication, Vol. IX, 2026, 143–158
Under CC BY 4.0
https://doi.org/10.33919/dasc.26.9.7

representational spaces, including spaces acquired through self-supervised and contrastive pretraining. They cannot, by their output alone, legitimate transformations of the categories and evidential norms that make clinical signs actionable. Primary studies of radiology report generation and of diagnostic reasoning with large language models confirm that fluent outputs do not settle clinical value and that capable models do not automatically improve human reasoning. Clinical reasoning is therefore framed as interpretive creativity: the construction and revision of defensible meanings under uncertainty and responsibility. Three criteria follow: accountability and answerability, indeterminacy tolerance, and the dialogicity and negotiability of sign value. A documented chest-radiograph case, followed by a hypothetical AI application, shows how these criteria can be operationalised conceptually. Since the framework remains conceptual, it requires empirical testing in clinical settings.

Keywords: artificial creativity, clinical reasoning, generative AI, interpretive creativity, medical semiotics

1. Introduction

Generative and multimodal AI has entered a medical discourse in which model outputs are sometimes described as creative or creativity-like. Shaban-Nejad et al. (2023), for example, explicitly discuss "creative and generative artificial intelligence" in personalised medicine, while Falahee et al. (2025) describe some large language model medication choices and rationales as creative and reasonable. Creativity conventionally implies more than novelty: it also implies evaluation, agency, and responsibility for whether a novel proposal is appropriate (Boden 2004; Sawyer 2012). In medicine, those implications are unusually consequential because interpretations can alter investigation, treatment, and patient risk.

The term is also unstable. In technical contexts, "generative" describes a class of models and objectives; in cultural discourse, "creative" can suggest originality, imagination, authorship, or autonomous insight. Treating these senses as interchangeable risks importing human-centred expectations into systems whose outputs are produced through statistical learning and constrained evaluation. The point of a semiotic analysis is not to deny AI novelty, but to separate the production of a novel sign-combination from the

social and clinical processes that make an interpretation valid, consequential, and attributable.

This research is a conceptual contribution from semiotics to the clinical-AI debate. It does not claim first-hand clinical expertise. Its account of clinical reasoning is synthesised primarily from medical-education and clinical-reasoning literature (Bleakley 2021; Chartash et al. 2021; Deschênes et al. 2025; Kassirer 2010), supplemented by empirical studies of medical AI. The proposed framework therefore requires testing with clinicians, patients, and real workflows before it can be treated as an operational clinical instrument.

The central question is not whether AI can generate outputs that are statistically new, fluent, or useful. It is what kind of creativity such outputs instantiate. Computational-creativity research treats novelty as generation or search within a structured possibility space and asks how value is assigned to the result (Colton & Wiggins 2012; Wiggins 2006). By contrast, clinical reasoning involves situated interpretation: heterogeneous signs must be related to a particular patient, temporal course, institutional setting, and field of possible action. This difference becomes clearer when "clinical data" are treated not as self-interpreting signals but as signs whose relevance emerges through interpretation (Eco 1976; Peirce 1998).

Multimodality sharpens the problem. Images, laboratory values, physiological traces, written notes, and patient narratives do not merely accumulate; each can reframe the meaning of the others (Jewitt 2014; Kress 2010). Medical AI may learn cross-modal regularities and produce convincing syntheses, yet clinical meaning also depends on contextual weighting, uncertainty, dialogue, and professional accountability. The article develops this distinction through Boden's (2004) three forms of creativity – combinational, exploratory, and transformational – and uses them as the organising spine of the analysis.

The argument proceeds in five steps. Section 2 establishes the minimal theoretical vocabulary needed for creativity, multimodality, and sign interpretation. Section 3 describes clinical reasoning as interpretive creative labour, while explicitly distinguishing cultural-semiotic "medical semiotics" from clinical semeiotics or semiology, the older clinical discipline concerned with eliciting and interpreting signs and symptoms. Section 4 examines what current multimodal systems can create, including supervised pipelines and foundation models. Section 5 derives three evaluative criteria and applies them to a worked example. Section 6 states the limits and implications of the framework.

2. Theoretical background and literature review

2.1. Creativity as novelty within and beyond a conceptual space

Creativity is usually defined through a conjunction of novelty and appropriateness: an outcome must be new in some relevant sense and also count as valuable, meaningful, or fit for purpose within a domain (Boden 2004; Sawyer 2012). This makes creativity irreducibly evaluative. A surprising output is not creative merely because it differs from previous outputs; its status depends on the norms by which a community recognises a contribution as worthwhile.

Boden's (2004) trichotomy is especially useful here. Combinational creativity produces unfamiliar arrangements of familiar elements. Exploratory creativity searches a structured conceptual space and discovers possibilities that its rules already permit. Transformational creativity changes one or more dimensions of that space, altering what can count as a possible or valuable outcome. Wiggins (2006) formalises related questions for computational creativity by distinguishing the space of possible artefacts, the rules for traversing it, and the evaluation criteria applied to outputs. Colton and Wiggins (2012) further stress that generation alone is insufficient if evaluation remains entirely external to the system.

These distinctions also prevent creativity from being reduced to an individual mental property. Creative work is distributed across people, artefacts, routines, and evaluative institutions (Literat & Glăveanu 2016). In medicine, this point is decisive: a proposal becomes clinically meaningful not simply because it is generated, but because it is interpreted, tested, documented, challenged, and made accountable within a professional system.

For medical AI, the analytic advantage of this model is that it separates three questions that are often collapsed: What is new in the output? What rules or learned regularities made that output reachable? Who or what determines that the output is appropriate? A system may be highly productive at the first two levels while depending on external institutions for the third. Conversely, a clinician's contribution may contain little surface novelty yet still be creative in the strong domain-relative sense if it reorganises evidence in a way that makes a previously obscure case intelligible and defensible. The relevant contrast is therefore not machine repetition versus human inspiration, but different locations of novelty, evaluation, and responsibility within a distributed system.

2.2. Multimodality and semiotic interpretation

For the present argument, multimodality contributes one essential premise: meaning emerges from relations among modes rather than from isolated tokens. Written notes, speech, images, numbers, layout, and embodied signs are socially shaped resources whose salience depends on how they are combined in a particular activity (Jewitt 2014; Kress 2010). A radiographic feature, for example, is interpreted through its relation to a report, prior imaging, symptoms, and clinical history rather than as a free-standing visual fact.

General semiotics adds a second premise. A sign does not carry a complete meaning within itself; it becomes meaningful through an interpretant and through codes, purposes, and contexts of use (Eco 1976; Peirce 1998). For AI, the relevant issue is therefore not only whether a system detects or generates a pattern, but how that pattern enters a chain of interpretation in which relevance, uncertainty, and action are negotiated.

2.3. Medical semiotics, clinical semeiotics, and institutional visibility

The term medical semiotics is used here in the cultural and communicative sense developed by Colaguori and Danesi (2017) and Danesi and Zukowski (2019): the study of how bodily, verbal, technological, and institutional signs acquire medical meaning. This usage is distinct from clinical semeiotics/semiology, the long-standing clinical discipline concerned with the systematic elicitation and interpretation of signs and symptoms. The two overlap in their concern with signification but operate at different analytic levels.

Medical meaning is dialogic and institutionally organised. What counts as evidence depends on patient narrative, professional genre, measurement practice, and the categories through which findings become visible and actionable (Danesi & Zukowski 2019). Foucault's ([1963] 1975) "medical gaze" directs attention to regimes of visibility: records, templates, taxonomies, and diagnostic conventions do not merely store observations; they help determine what can be noticed, compared, and legitimised. Telemedicine makes this mediation especially visible because camera framing, patient-operated devices, and platform constraints alter the signs available for interpretation (Brizio et al. 2022).

3. Clinical reasoning as interpretive creative labour

3.1. Narrative, inference, and accountability

The clinical-reasoning literature describes diagnosis as iterative explanatory work rather than simple classification. Clinicians develop provisional hypotheses, seek discriminating evidence, revise their accounts, and justify why a diagnosis is plausible for this patient at this time (Bleakley 2021; Kassirer 2010). Narrative organisation matters because temporal sequence, negative findings, treatment response, and contextual detail can change the inferential force of the same observation (Chartash et al. 2021). Deschênes et al. (2025) similarly describe reasoning as transformation and articulation of heterogeneous clinical data into problem representations and justificatory accounts.

From a semiotic perspective, this is interpretive creativity. It is creative not because clinicians are licensed to innovate freely, but because they must construct a defensible configuration of signs when the relevant pattern is incomplete, contested, or evolving. The novelty lies in the case-specific organisation of evidence; appropriateness is constrained by clinical knowledge, patient consequences, professional standards, and the obligation to remain revisable when new evidence appears.

The expression "the clinician" should not obscure important differences among specialties and encounters. A radiologist may reason primarily through image morphology, comparison studies, and the clinical question supplied in a referral; a psychiatrist may depend heavily on dialogue, temporal narration, behaviour, and the pragmatic conditions of the interview; a general practitioner managing a chronic condition may interpret small changes against a longitudinal relationship and knowledge of the patient's circumstances. These are different semiotic ecologies. What they share is not a single reasoning sequence but the need to organise incomplete signs into an account that can guide action and remain open to revision.

Clinical reasoning also involves repeated transformations of form. A patient's account may become a problem list; a physical or visual finding may become a coded descriptor; a series of measurements may become a trend; an imaging appearance may become a differential diagnosis. Each transformation gains economy and comparability but can also suppress details that were meaningful in the preceding mode. Interpretive creativity therefore includes deciding what can be compressed, what must remain explicit, and when a representation should be reopened because the case no longer fits the working frame.

3.2. A worked multimodal scenario

A published case reported by Saliba and Tack (2023) provides a concrete example of multimodal clinical interpretation. A 60-year-old woman presented to the emergency department with dyspnoea. As part of the initial investigation, a chest radiograph was obtained. Rather than showing a nonspecific respiratory abnormality, the image displayed a prominent pulmonary artery, focal peripheral hyperlucency caused by oligemia, and abrupt tapering of the pulmonary artery, known respectively as Fleischner's sign, Westermark's sign, and the knuckle sign. These findings raised the possibility of central pulmonary embolism and prompted contrast-enhanced CT pulmonary angiography, which confirmed emboli within the enlarged right intermediate pulmonary artery. The patient was subsequently treated with anticoagulation.

The case illustrates why clinical meaning cannot be reduced to an isolated visual feature. Dyspnoea provided the presenting narrative sign; the chest radiograph reorganised the initial diagnostic possibilities by making several otherwise uncommon but relevant visual signs salient; and CT angiography subsequently stabilised the interpretation by demonstrating the embolus directly. The diagnostic meaning therefore emerged sequentially across symptom narrative, conventional radiography, and confirmatory cross-sectional imaging. The chest radiograph did not independently constitute a diagnosis: its significance depended on the relation between the patient's presentation, the configuration of radiographic signs, and the subsequent evidential test. This is the kind of multimodal and abductive process in which a finding acquires clinical force through its relation to other signs rather than through intrinsic meaning alone.

In terms of creativity theory, the case also clarifies the distinction developed here. Recognising a familiar conjunction of dyspnoea and established radiographic signs corresponds to combinational novelty, while considering pulmonary embolism among competing explanations and pursuing confirmatory imaging exemplifies exploratory novelty within an established diagnostic space. Transformational novelty would require a more fundamental alteration of the categories or evidential relations through which the case is rendered clinically intelligible. The example therefore anchors the argument in documented practice without implying that the present paper itself derives from clinical observation.

4. What multimodal AI creates in medicine

4.1. Multimodal architectures and the three forms of novelty

Multimodal medical AI does not have a single training or fusion architecture. Supervised systems may learn from labelled images, coded outcomes, or annotated records, whereas foundation models can learn representations through self-supervised and contrastive objectives with unlabelled or weakly paired data. EyeCLIP, for example, combines self-supervised reconstruction, multimodal image contrastive learning, and image–text contrastive learning across 2.77 million ophthalmic images and partial clinical text (Shi et al. 2025). Multimodal medical AI is therefore not necessarily dependent on pre-labelled inputs.

Fusion also occurs at different stages. Early fusion combines modalities near the input level; intermediate fusion learns joint or cross-attentive representations; late fusion combines modality-specific predictions or embeddings (Judge et al. 2024; Simon et al. 2025). These choices have semiotic consequences because they determine when relations among modes are learned and which modality-specific distinctions remain available for later interpretation. No single architecture, however, removes the need for a representational and evaluative space in which outputs become possible.

The architecture matters because multimodal "integration" is not a neutral technical step. Early fusion can allow cross-modal relations to shape representation from the outset but may make it difficult to recover which modality drove a particular inference. Late fusion preserves modality-specific pathways longer but can reduce interaction to the combination of already-formed outputs. Intermediate fusion creates shared representations at selected stages and therefore makes the timing and direction of cross-modal influence part of the design. These differences do not map mechanically onto clinical quality, but they matter for a semiotic account because they specify where one mode can reweight, anchor, or suppress another.

Boden's (2004) categories clarify what follows. First, multimodal systems exhibit combinational novelty when they assemble learned visual, textual, numerical, or temporal relations into a new report, summary, or hypothesis. Second, they exhibit exploratory novelty when they search or sample a learned space to produce alternatives that are not stored as fixed responses. Foundation-model pretraining expands that space beyond manually labelled classes, but it does not make the space norm-free: training corpora, reporting practices, data selection, prompt conventions, and downstream evaluation still structure what is easy to represent and reward.

Transformational creativity is the limiting case. A model may generate a proposal that appears to challenge an established category, but generation is not equivalent to legitimating a transformation of the clinical conceptual space. Changing what counts as a disease category, a salient sign, an acceptable threshold, or a warranted action requires evidential and institutional ratification. In terms of Foucault's (1975) clinical visibility, taxonomies and documentation genres are part of the infrastructure through which a clinical object becomes recognisable. AI can perturb or expose those structures, but an output cannot by itself authorise their replacement.

The distinction is therefore not between "labelled AI" and "human interpretation". It is between novelty generated within computationally learned spaces and the responsibility for deciding whether the boundaries, categories, or evaluative criteria of those spaces should govern a particular case. That distinction survives the shift from supervised pipelines to self-supervised foundation models.

This formulation also avoids a second overstatement: foundation models do not escape institutional mediation simply because their pretraining is self-supervised. The corpus still has a provenance; clinical text and images still reflect decisions about what was recorded, how it was described, and which populations and institutions produced the data. Downstream fine-tuning, prompting, retrieval, and evaluation then introduce further constraints. Self-supervision changes the mechanism by which a representational space is learned; it does not make that space socially or clinically unstructured.

4.2. Primary evidence: fluent synthesis, error, and human–AI interaction

Primary studies show why this distinction matters. Tanno et al. (2025) evaluated Flamingo-CXR, a vision-language foundation-model system for chest-radiograph report generation, with board-certified radiologists. Preferences varied across readers and settings, and clinically significant errors occurred in both human and AI reports; in the in/outpatient set, 22.8% of cases contained clinically significant errors in the AI report alone. The result is not simply that AI can be wrong. It shows that output quality is partly a matter of situated expert judgement and that fluent report generation does not settle the clinical value of the interpretation.

Tanno et al. (2025) also show why purely automatic evaluation is insufficient. Their study used specialist review because lexical similarity between generated and reference reports cannot determine whether a difference is clinically important. Radiologists disagreed in some cases about preference and significance, indicating that the "value" component of creativity remains partly embedded in professional judgement. The model can generate a well-formed report and still leave unresolved whether a finding was omitted, overcalled, or framed with an inappropriate degree of certainty. Semiotic adequacy is therefore not reducible to textual resemblance.

A different mechanism appears in Goh et al.'s (2024) randomised clinical trial. Physicians given access to a large language model did not significantly outperform physicians using conventional resources, even though the model alone scored highly on the study rubric. This finding shifts the analytic focus from model competence to human–AI orchestration: a capable output does not automatically improve the reasoning process into which it is inserted. Falahee et al. (2025) similarly found that some medication suggestions were judged creative or useful while hallucinations and unsafe choices remained a concern. Novelty and appropriateness can therefore diverge within the same system.

The Goh et al. (2024) trial is equally important because it prevents the argument from locating failure solely inside the model. If access to a high-performing model does not automatically improve physician reasoning, then clinical AI must be analysed as part of an interactional system: what the model makes salient, how a user interrogates it, when suggestions are accepted or resisted, and whether the tool changes the sequence of information seeking. These are precisely the relations through which a sign acquires practical force.

These studies support a semiotic account of failure. Hallucination, premature closure, or context mismatch is not only a technical error at the level of tokens or labels; it can also be a failure to preserve the conditions under which a sign remains contestable, revisable, and answerable to consequences. The issue is not whether AI participates in creative work – it clearly can – but which forms of novelty it supplies and where evaluative authority resides.

5. Discussion

5.1. Synthetic and interpretive creativity

The study therefore distinguishes synthetic creativity from interpretive creativity. Synthetic creativity names the combinational and exploratory

novelty produced through learned associations, search, and generation. It may be technically sophisticated and clinically useful, and foundation models can widen its representational scope. Its limitation is not an absence of novelty but the fact that appropriateness is initially mediated through training data, optimisation objectives, benchmarks, prompts, and downstream human evaluation (Colton & Wiggins 2012; Wiggins 2006).

Interpretive creativity names the responsibility-bearing work of constructing and revising clinical meaning under uncertainty. It includes deciding what counts as a relevant sign, how modalities should be weighted, when an apparently coherent account should remain provisional, and whether a new framing deserves to alter action. Here, the term "interpretive" is used deliberately rather than "hermeneutic": the argument concerns situated, revisable interpretation but does not claim a substantive derivation from philosophical hermeneutics. The two forms are not mutually exclusive. AI may extend the range of combinations a clinician considers, and clinicians may use model outputs to explore a case differently. The boundary concerns evaluative authority. Transformational novelty in medicine becomes clinically meaningful only when changes to categories or evidential norms can be defended within practices of validation, governance, dialogue, and responsibility.

Boden's (2004) trichotomy can now be applied across the distinction. Synthetic creativity is strongest in combinational and exploratory activity: systems can recombine modalities and sample large possibility spaces with speed and breadth that may be clinically useful. Interpretive creativity becomes most visible when the task requires deciding whether the current space is adequate at all. A clinician can reject the framing supplied by a model, reopen a history, request a different test, or treat an apparently anomalous sign as evidence that the working category is wrong. Such acts are potentially transformational because they alter the terms in which the case is being represented. They remain constrained, however, by evidence and by collective professional validation rather than by individual preference.

5.2. Three criteria for evaluating creativity in AI-augmented care

The criteria below are derived from the creativity framework rather than added as a separate ethics checklist. Each addresses a point at which novelty must be converted into warranted clinical meaning. Accountability concerns the evaluation of appropriateness; indeterminacy tolerance concerns the quality of exploration within a space; dialogicity concerns whether the space and its sign-values remain open to revision. Together they ask whether

AI extends creative inquiry while preserving the practices that make clinical interpretation responsible.

The first criterion is accountability and answerability. Creativity theory asks not only whether an output is novel but who can warrant its appropriateness. In clinical care, this criterion therefore asks whether a human decision-maker can reconstruct and defend the reasons for accepting, modifying, or rejecting an AI output, and whether responsibility remains visible when the output affects action. This criterion supplements rather than replaces existing machinery for accountability, including clinical governance, validation, medical-device regulation, post-market surveillance, and liability. In the European context, medical-device software is subject to Regulation (EU) 2017/745 where its intended purpose brings it within the medical-device definition, while high-risk AI requirements may also apply under Regulation (EU) 2024/1689 (European Parliament and Council 2017, 2024). Regulation is relevant but does not exhaust answerability. A system may satisfy formal requirements and still be used in a local workflow in ways that blur responsibility, for example when generated text is copied into a record without clear verification or when a recommendation is treated as authoritative because its provenance is opaque. Conversely, a clinically responsible use may involve rejecting a technically sophisticated output because the reasons for action cannot be reconstructed. The criterion therefore connects formal governance to the micro-practices of attribution and justification at the point of care.

The second criterion is indeterminacy tolerance. Exploratory creativity is valuable only if the search space does not force a premature single answer where the evidence supports alternatives. An AI system should therefore represent uncertainty, missing information, and plausible competing interpretations rather than convert absence into apparent evidence or fluency into confidence. This is a creativity criterion because it governs whether exploration remains open or collapses into an unjustified result. Indeterminacy is not equivalent to indecision. Clinical work often requires action before certainty is complete; the point is that the representation of uncertainty should match the evidence available. A useful system may narrow a differential or rank possibilities, but it should not erase unresolved alternatives merely to produce a fluent answer. In creativity terms, productive exploration preserves the possibility of further search when the current space does not yet justify closure.

The third criterion is dialogicity and negotiability of sign value. If transformational creativity concerns changes to the conceptual space, then the

meanings and categories involved must remain revisable through interaction. Patient narrative, clinician questioning, new tests, comparison imaging, and institutional review should be able to reshape the interpretation rather than merely decorate a fixed model output (Brizio et al. 2022; Danesi & Zukowski 2019). Dialogicity is also a safeguard against treating model outputs as final representations of the patient. A generated summary can become a powerful anchor: once entered into documentation, subsequent clinicians may inherit its categories and emphases. Negotiability requires that later evidence can alter not only the final diagnosis but the salience structure of the case: what is treated as central, peripheral, absent, or uncertain. This is where semiotics connects most directly with transformational creativity.

Applied to the pulmonary-embolism scenario in Section 3.2, the three criteria become concrete. Consider, hypothetically, an AI system that interprets the chest radiograph as suggestive of pulmonary embolism on the basis of Fleischner's sign, Westermark's sign, and the knuckle sign. Accountability asks whether the clinician can reconstruct why those signs were treated as diagnostically relevant and justify the decision to proceed to CT pulmonary angiography. Indeterminacy tolerance asks whether the output distinguishes suspicion from confirmation and preserves plausible alternatives rather than presenting pulmonary embolism as established from radiography alone. Dialogicity asks whether additional history, examination findings, prior imaging, or the subsequent CT result can revise the generated interpretation. A system that satisfies all three criteria would support exploratory reasoning without displacing the clinical processes through which meaning becomes warranted and actionable.

6. Conclusion

Calling medical AI "creative" is defensible only if the term is specified. Current multimodal and generative systems can produce genuine combinational and exploratory novelty, including through foundation-model architectures that learn from unlabelled or weakly paired data. Their limit is not that they merely copy labelled examples. It is that a generated output cannot, by itself, establish the clinical legitimacy of the categories, evidential thresholds, or actions through which that output becomes consequential. The semiotic distinction developed here separates synthetic creativity from interpretive creativity. The former expands and searches learned representational spaces; the latter constructs defensible meaning under uncertainty and remains answerable to patients, professional standards, and institutions.

Transformational novelty sits at their boundary: AI may propose a reframing, but its adoption requires human and institutional validation.

The three criteria – accountability and answerability, indeterminacy tolerance, and dialogicity and negotiability of sign value – operationalise that distinction at a conceptual level. They are intended as questions for evaluating how AI outputs enter clinical interpretation, not as validated clinical metrics. Their next test is empirical: observation of how clinicians and patients accept, revise, contest, or refuse multimodal AI outputs in real encounters and how those practices interact with formal governance and regulatory requirements.

References

Bleakley, A. 2021. "Re-visioning Clinical Reasoning, or Stepping out from the Skull". *Medical Teacher*, Vol. 43, No. 4, 456–462. Available at: https://doi.org/10.1080/0142159X.2020.1859098.

Boden, M. A. 2004. *The Creative Mind: Myths and Mechanisms*. (2nd ed.). London and New York: Routledge.

Brizio, A., M. Tirassa and M. Castiglioni. 2022. "Medical Semiotics in the Telemedicine Era: The Birth of Telesemiotics". *International Journal of Medical Informatics*, Vol. 157, Article 104573. Available at: https://doi.org/10.1016/j.ijmedinf.2021.104573.

Chartash, D., M. B. Rosenman, J. Bollen, M. Dickinson and S. M. Downs. 2021. "A Complex Model of Clinical Narrative Information for the Diagnostic Act". medRxiv preprint. Available at: https://doi.org/10.1101/2021.02.21.21252158.

Colaguori, R., M. Danesi. 2017. "Medical Semiotics: A Revisitation and an Exhortation". *International Journal of Semiotics and Visual Rhetoric*, Vol. 1, No. 1, 11–18. Available at: https://doi.org/10.4018/IJSVR.2017010102.

Colton, S., G. A. Wiggins. 2012. Computational Creativity: The Final Frontier? In De Raedt, L., C. Bessière, D. Dubois, P. Doherty, P. Frasconi, F. Heintz and P. J. F. Lucas (eds.). *ECAI 2012: 20th European Conference on Artificial Intelligence*. Amsterdam: IOS Press, 21–26. Available at: https://doi.org/10.3233/978-1-61499-098-7-21.

Danesi, M., N. Zukowski. 2019. *Medical Semiotics: Medicine and Cultural Meaning*. Munich: LINCOM Europa.

Deschênes, M.-F., N. Fernandez, K. Lechasseur, M.-È. Caty, B. M. Uctu, Y. Bouzeghrane and P. Lavoie. 2025. "Transformation and Articulation of Clinical

Data to Understand Students' Clinical Reasoning: A Scoping Review". *BMC Medical Education*, Vol. 25, Article 52. Available at: https://doi.org/10.1186/s12909-025-06644-7.

Eco, U. 1976. *A Theory of Semiotics*. Bloomington: Indiana University Press.

European Parliament and Council. 2017. Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on Medical Devices. *Official Journal of the European Union*, L 117, 5 May 2017, 1–175.

European Parliament and Council. 2024. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, L, 2024/1689, 12 July 2024.

Falahee, B., H. Lee, T. Tran, T. Deasy, C. Cristancho, C. M. Apovian and D. W. Kim. 2025. "Comparison of Antiobesity Medication Decisions: Large Language Models and Obesity Medicine Expert Clinicians". *Diabetes Technology and Obesity Medicine*, Vol. 1, No. 1, 71–78. Available at: https://doi.org/10.1089/dtom.2025.0021.

Foucault, M. 1975 [1963]. *The Birth of the Clinic: An Archaeology of Medical Perception*. (transl. A. M. Sheridan Smith). New York: Vintage Books.

Goh, E., R. Gallo, J. Hom, E. Strong, Y. Weng, H. Kerman, J. A. Cool, Z. Kanjee, A. S. Parsons, N. Ahuja, E. Horvitz, D. Yang, A. Milstein, A. P. J. Olson, A. Rodman and J. H. Chen. 2024. "Large Language Model Influence on Diagnostic Reasoning: A Randomized Clinical Trial". *JAMA Network Open*, Vol. 7, No. 10, Article e2440969. Available at: https://doi.org/10.1001/jamanetworkopen.2024.40969.

Jewitt, C. (ed.). 2014. *The Routledge Handbook of Multimodal Analysis*. (2nd ed.). London and New York: Routledge.

Judge, C. S., F. Krewer, M. J. O'Donnell, L. Kiely, D. Sexton, G. W. Taylor, J. A. Skorburg and B. Tripp. 2024. "Multimodal Artificial Intelligence in Medicine". *Kidney360*, Vol. 5, No. 11, 1771–1779. Available at: https://doi.org/10.34067/KID.0000000000000556.

Kassirer, J. P. 2010. "Teaching Clinical Reasoning: Case-Based and Coached". *Academic Medicine*, Vol. 85, No. 7, 1118–1124. Available at: https://doi.org/10.1097/ACM.0b013e3181d5dd0d.

Kress, G. 2010. *Multimodality: A Social Semiotic Approach to Contemporary Communication*. London and New York: Routledge.

Literat, I., V. P. Glăveanu. 2016. "Same but Different? Distributed Creativity in the Internet Age". *Creativity: Theories – Research – Applications*, Vol. 3, No. 2, 330–342. Available at: https://doi.org/10.1515/ctra-2016-0020.

Peirce, C. S. 1998. *The Essential Peirce: Selected Philosophical Writings*, Vol. 2 (1893–1913). (ed. Peirce Edition Project). Bloomington: Indiana University Press.

Saliba, T., D. Tack. 2023. "Central Pulmonary Embolism Detected on a Chest X-Ray: A Case Report". *Journal of the Belgian Society of Radiology*, Vol. 107, No. 1, Article 11. Available at: https://doi.org/10.5334/jbsr.3052.

Sawyer, R. K. 2012. *Explaining Creativity: The Science of Human Innovation*. (2nd ed.). New York: Oxford University Press.

Shaban-Nejad, A., M. Michalowski and S. Bianco. 2023. "Creative and Generative Artificial Intelligence for Personalized Medicine and Healthcare: Hype, Reality, or Hyperreality?". *Experimental Biology and Medicine*, Vol. 248, No. 24, 2497–2499. Available at: https://doi.org/10.1177/15353702241226801.

Shi, D., W. Zhang, J. Yang, S. Huang, X. Chen, P. Xu, K. Jin, S. Lin, J. Wei, M. Yusufu, S. Liu, Q. Zhang, Z. Ge, X. Xu and M. He. 2025. "A Multimodal Visual–Language Foundation Model for Computational Ophthalmology". *npj Digital Medicine*, Vol. 8, Article 381. Available at: https://doi.org/10.1038/s41746-025-01772-2.

Simon, B. D., K. B. Ozyoruk, D. G. Gelikman, S. A. Harmon and B. Türkbey. 2025. "The Future of Multimodal Artificial Intelligence Models for Integrating Imaging and Clinical Metadata: A Narrative Review". *Diagnostic and Interventional Radiology*, Vol. 31, No. 4, 303–312. Available at: https://doi.org/10.4274/dir.2024.242631.

Tanno, R., D. G. T. Barrett, A. Sellergren, S. Ghaisas, S. Dathathri, A. See, J. Welbl, C. Lau, T. Tu, S. Azizi, K. Singhal, M. Schaekermann, R. May, R. Lee, S. W. Man, S. Mahdavi, Z. Ahmed, Y. Matias, J. Barral, S. M. A. Eslami, D. Belgrave, Y. Liu, S. R. Kalidindi, S. Shetty, V. Natarajan, P. Kohli, P.-S. Huang, A. Karthikesalingam and I. Ktena. 2025. "Collaboration between Clinicians and Vision–Language Models in Radiology Report Generation". *Nature Medicine*, Vol. 31, No. 2, 599–608. Available at: https://doi.org/10.1038/s41591-024-03302-1.

Wiggins, G. A. 2006. "Searching for Computational Creativity". *New Generation Computing*, Vol. 24, No. 3, 209–222. Available at: https://doi.org/10.1007/BF03037332.