Mansour *et al. Journal of Artificial Intelligence for Materials Science* *(2026) 1:103*
https://doi.org/10.51847/j862290346

ORIGINAL RESEARCH **Open access**

Default Assumptions as Hidden Parameters: A Theory of Implicit Priors in Materials AI

Ahmed Mansour[1*], Omar Saeed[1]

Abstract

In the rapidly expanding domain of artificial intelligence applied to materials science, default assumptions embedded within machine learning pipelines—ranging from software library choices and architectural presets to data preprocessing routines and evaluation protocols—are routinely treated as neutral, inconsequential background elements that require no explicit justification. Yet these defaults operate as hidden parameters, subtly yet powerfully constraining the hypothesis space, directing optimization trajectories, and ultimately shaping the predictive behavior of models in ways that rival or even exceed the influence of explicitly tuned parameters, as theoretical analyses of deep networks have long emphasized. This paper advances the theoretical claim that default assumptions in materials AI function as implicit priors, encoding unacknowledged inductive biases that propagate through every stage of a pipeline and determine what counts as a valid or reliable prediction about material properties. Building directly on foundational examinations of inductive bias, we distinguish defaults from both explicit parameters and tunable hyperparameters, develop a taxonomy of four primary default types specific to materials informatics, and derive corollaries concerning the epistemic consequences of unexamined defaults for model comparison, reproducibility, and knowledge transfer. We further examine why such defaults persist—owing to cognitive convenience, historical path dependence, and systematic attribution errors—and clarify their subtle yet critical relation to formal Bayesian priors, while noting that understanding deep learning requires rethinking generalization when defaults remain hidden. The analysis culminates in concrete implications for practice, proposing that defaults must be elevated to first-class objects of documentation, justification, and sensitivity analysis if materials AI is to achieve genuine epistemic transparency and scientific robustness. By theorizing defaults as hidden parameters, this work identifies an overlooked dimension of model epistemology in materials science and offers a conceptual framework for making the invisible visible.

Keywords Materials AI, Default assumptions, Implicit priors, Hidden parameters, Inductive bias, Epistemology of machine learning

*Correspondence:
Ahmed Mansour
ahmed.mansour@gmail.com
[1] Department of Materials Engineering and AI Applications, Cairo University, Cairo, Egypt

Introduction

Materials AI systems now routinely predict crystal structures, electronic properties, mechanical responses, and phase stability across vast chemical spaces. Yet, the pipelines powering these predictions contain countless default assumptions that remain largely invisible to both developers and end-users. Architecture choices such as the number of message-passing layers in graph neural networks, hyperparameter defaults inherited from general-purpose libraries, data preprocessing decisions like normalization schemes or cutoff distances, and even evaluation metrics are treated as neutral infrastructure rather than active theoretical commitments. The problem is not merely practical; it is epistemological. These defaults shape model behavior as profoundly as any explicitly optimized weight or learning rate, yet they escape the same level of scrutiny applied to explicit parameters. As a result, claims of model superiority or generalization often conflate the effects of deliberate design with the cumulative influence of unexamined presets.

This oversight is particularly acute in materials informatics because the domain inherits defaults from both general machine learning frameworks and domain-specific conventions in solid-state physics and chemistry [1-5], as surveys of machine learning for molecular and materials science have documented without full epistemic critique. Periodic boundary conditions, symmetry enforcement, and local atomic environment cutoffs, for instance, enter pipelines not through explicit scientific hypothesis but through the path of least resistance in established codebases. The present work, therefore, theorizes default assumptions as an overlooked class of hidden parameters in materials AI systems. We argue that choices made by software libraries, inherited from prior literature, or set without justification function as implicit priors that constrain the space of learnable functions, bias gradient-based optimization, and define the very criteria by which predictions are judged successful.

The theoretical foundation for this analysis draws upon earlier examinations of inductive bias and implicit regularization in deep networks. Deep learning relies on implicit regularization [2], but the specific form is rarely examined when models are deployed to materials problems. Similarly, graph neural network architectures embed assumptions about local connectivity [4], which may not match material structure, and recent advances in solid-state materials science further illustrate how such defaults propagate unnoticed [6]. By foregrounding these defaults, the paper moves beyond surface-level performance discussions to interrogate the epistemic structure of materials AI pipelines. It identifies how unstated assumptions in data representation, such as those surrounding periodic boundary conditions or rotational invariance, silently encode scientific commitments that later manifest as model behavior.

What Are Default Assumptions?

Default assumptions are the unexamined, pre-set configurations that govern the behavior of a computational pipeline without requiring—or even inviting—explicit justification at the point of use. They differ fundamentally from explicit parameters, which are deliberately chosen and optimized during training, and from hyperparameters, which, although often tuned via search procedures, are at least acknowledged as tunable. Defaults, by contrast, are inherited from software libraries, community conventions, or the initial configuration of a modeling framework and are accepted precisely because they appear neutral or "sensible." Definition 1 formalizes this concept:

Default assumptions are any configuration choices in a machine learning pipeline—architectural, algorithmic, data-related, or evaluative—that are instantiated automatically by the modeling environment, inherited from prior literature without re-examination, or set by library developers without domain-specific justification, thereby functioning as hidden parameters that shape model behavior equivalently to explicit parameters while remaining invisible to standard reporting practices.

Table 1 clarifies the epistemic distinction between defaults, hyperparameters, and explicit parameters, highlighting the unique invisibility and influence of defaults as hidden parameters.

Table 1. Distinguishing defaults, hyperparameters, and explicit parameters as epistemic objects in materials AI

Dimension	Default assumptions	Hyperparameters
Visibility	Invisible/implicit	Explicit but tunable
Origin	Libraries, conventions, and inherited code	User-defined search space
Requirement of Justification	None (typically omitted)	Partial (via tuning rationale)
Tunability	Rarely modified	Actively tuned
Epistemic Status	Hidden parameters (implicit priors)	Semi-explicit design choices
Influence on Hypothesis Space	Strong (pre-training constraint)	Moderate

Influence on Optimization	Strong (via initialization, optimizer defaults)	Moderate	
Role in Reproducibility	Critical but undocumented	Important and reported	
Typical Reporting Practice	Omitted	Included	
Failure Mode	Unrecognized bias	Overfitting/underfitting	r

The concept builds directly on earlier theoretical work showing that deep networks exhibit strong implicit regularization even when explicit regularization terms are absent. In search of the real inductive bias, analyses have demonstrated that optimization dynamics themselves impose structure on solutions [2]. Yet, those dynamics are heavily influenced by the default settings of the optimizer, the initialization scheme, and the loss landscape geometry. When these defaults are carried into materials AI without scrutiny, they import inductive biases that may be appropriate for natural images or language but are ill-suited to the discrete, periodic, and symmetry-rich nature of atomic configurations. Theoretical issues in deep networks further underscore that architecture itself imposes inductive bias that cannot be overcome by data alone [1].

Furthermore, defaults encode scientific assumptions that escape critical examination precisely because they are not labeled as assumptions. A default distance cutoff in an atomic graph construction, for instance, implicitly asserts that interactions beyond that radius are negligible—an assertion that is scientific rather than merely technical. The gap between explicit model specification and actual model behavior, therefore, widens when defaults are left unexamined, [3] as rethinking generalization has shown in broader machine learning contexts. In philosophical terms, defaults represent unacknowledged priors that structure inference without ever entering the formal Bayesian machinery. They are not absent from the model; they are rendered transparent by convention.

Distinguishing defaults from formal priors is essential for clarity. Formal priors are stated, debated, and sometimes updated; defaults are enacted without statement. Yet both constrain the posterior. The present section, therefore, establishes defaults as a distinct epistemic category—one that demands its own theoretical treatment if materials AI is to mature beyond black-box empiricism. By making defaults visible, we open them to the same standards of justification applied to explicit scientific hypotheses.

Default Assumptions in Materials AI

Materials AI pipelines are especially dense with defaults because the domain bridges general machine learning frameworks with the specialized representational demands of crystalline and molecular systems. Periodic boundary handling in graph neural networks offers a paradigmatic case. Most implementations default to minimum-image conventions or fixed supercell replications without requiring the user to declare the underlying assumption about translational invariance. These defaults are not neutral; they embed the scientific premise that the material of interest is perfectly periodic—an assumption that fails for defects, surfaces, or amorphous phases. Graph networks as a universal machine learning framework for molecules and crystals inherit such defaults [7], yet the implicit priors they encode are rarely articulated.

Symmetry assumptions constitute another pervasive class. Rotational and translational invariance are frequently imposed by default through equivariant architectures or data augmentation schemes that assume perfect symmetry preservation. Recent symmetry-aware graph neural networks formalize this [8-15]. Still, the default enforcement of such symmetries in standard libraries quietly rules out certain symmetry-breaking phenomena that are physically relevant in real materials. The unstated symmetry assumptions in materials data representations, therefore, function as hidden filters on the hypothesis space.

Default distance cutoffs for atomic interactions provide a third illustration. Many featurization routines default to a 5–6 Å radius for neighbor lists, implicitly asserting that longer-range interactions contribute negligibly to the property of interest. This algorithmic default is carried forward across publications without re-evaluation, even though the appropriate cutoff is itself a material-specific scientific question. Similarly, default pooling operations—mean versus sum versus max—encode different assumptions about extensivity and locality. Mean pooling, for instance, implicitly normalizes across system size in ways that may conflict with extensive thermodynamic properties.

Loss-function defaults add yet another layer. Mean-squared error remains the unchallenged default for regression tasks in materials property prediction, thereby embedding Gaussian error assumptions that may not hold for properties with heavy-tailed distributions or physical constraints. Machine learning for molecular and materials science surveys these pipelines [5], yet leaves the default character of such choices largely implicit. Recent advances and applications of machine learning in solid-state materials science similarly document the proliferation of such presets without subjecting them to epistemic critique [6].

The cumulative effect is a pipeline in which scientific assumptions enter not through deliberate hypothesis formulation but through the path of least resistance in code execution. Untrained neural network priors and architecture choices as implicit regularization further illustrate how defaults propagate. In materials-specific contexts, these defaults become hidden scientific commitments whose influence is masked by the apparent generality of the underlying library. The present section, therefore, maps the terrain of materials AI defaults, demonstrating that they are not peripheral technicalities but central, if invisible, determinants of what models can learn about matter.

Theoretical Claim: Defaults as Hidden Parameters

Default assumptions in materials AI pipelines operate as implicit priors whose influence rivals, and often exceeds, that of explicitly specified parameters. Their effect is not confined to initialization or convenience; rather, they permeate the entire learning process by delimiting the hypothesis space, shaping the trajectory of optimization, and quietly defining the standards by which predictive success is judged. In this sense, defaults function as a latent architecture of inference, structuring what the model can represent, how it arrives at solutions, and which outcomes are ultimately deemed acceptable.

A central mechanism through which this influence manifests lies in the prior restriction of representational possibilities. Architectural presets, introduced before any data are encountered, embed inductive biases that cannot be fully neutralized through training. Foundational analyses of deep networks have shown that the choice of architecture imposes constraints on the class of functions that can be learned, independent of dataset size or diversity [1]. Within materials informatics, this becomes particularly consequential when model structures encode assumptions about locality or symmetry. For example, a graph neural network configured with a default message-passing scheme implicitly privileges short-range interactions, thereby predisposing the model toward representations that may inadequately capture long-range order. The resulting errors are not incidental but structurally induced, reflecting the boundaries imposed by the initial design rather than deficiencies in the data.

This structural conditioning extends beyond representation into the dynamics of learning itself. Optimization in deep learning is not a neutral search over parameter space but a guided trajectory shaped by algorithmic defaults such as initialization strategies, learning rates, and normalization schemes. Empirical and theoretical work on implicit regularization demonstrates that these factors influence which solutions are reachable under gradient-based training, even when the explicit loss function remains unchanged [13]. In materials AI, where models often navigate high-dimensional and sparsely sampled landscapes, such defaults can steer the optimizer toward regions that align with embedded priors rather than with the full complexity of the underlying data. Further analyses of generalization reinforce the idea that the path taken during optimization encodes its own form of bias, independent of the objective being minimized [14]. What emerges is a learning process in which the trajectory itself becomes a carrier of prior assumptions.

The influence of defaults culminates in the evaluation stage, where criteria of success are codified in ways that are rarely interrogated. Choices of metrics and validation protocols implicitly define what constitutes a “good” prediction, thereby shaping both model development and comparative assessment. As highlighted in discussions of generalization in deep learning, the construction of test sets and the selection of evaluation metrics play a decisive role in determining which models appear superior [3]. In materials science applications, the widespread reliance on mean absolute error as a default loss function embeds the assumption that all deviations are equally consequential, an assumption that may conflict with domain-specific priorities where certain property thresholds carry disproportionate importance. Under these conditions, evaluation is not merely a neutral checkpoint but an extension of the prior structure embedded throughout the pipeline.

When considered as a whole, the propagation of defaults reveals a cumulative process through which implicit assumptions are layered and amplified across stages of the workflow. Initial choices in data ingestion—such as featurization schemes or normalization procedures—establish the first set of constraints, which are then reinforced by architectural decisions that determine how information is structured and propagated. Optimization settings further channel the learning dynamics, while evaluation protocols ultimately crystallize these influences into judgments of performance. The result is a compound prior that emerges not from any single parameter but from the interaction of multiple defaults operating in concert. This composite structure exerts a formative influence on model behavior, underscoring the need to treat defaults not as benign conveniences but as integral components of epistemic design in materials AI systems.

Figure 1 formalizes the hierarchical propagation of default assumptions as hidden parameters, illustrating how implicit priors enter at the configuration level and systematically shape model behavior through distinct pipeline mechanisms.

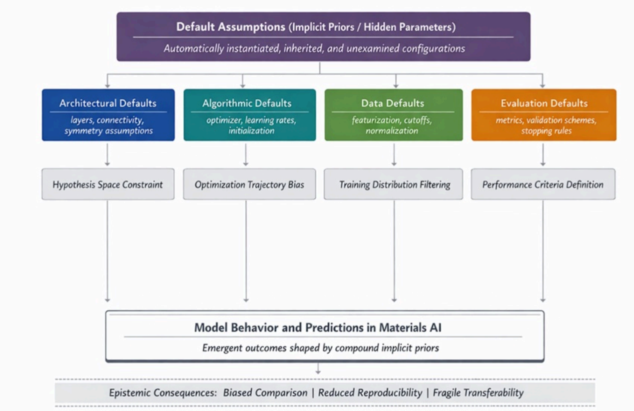

Figure 1. The hierarchical propagation of default assumptions as hidden parameters.

This theoretical claim reframes defaults not as background noise but as hidden parameters whose causal role must be acknowledged if materials AI is to achieve epistemic maturity. By treating defaults as implicit priors, the framework aligns with broader calls to examine the gap between explicit model specification and actual model behavior while extending those insights specifically to the materials domain.

A Taxonomy of Default Types

A more nuanced understanding of defaults in materials AI requires distinguishing the different ways in which they operate across the pipeline, as their influence is neither uniform nor interchangeable. What appears superficially as a single category of “preset choices” in fact spans multiple layers of the modeling process, each encoding a distinct form of implicit prior that shapes how knowledge is represented, learned, and evaluated.

At the level of model construction, architectural choices exert the most immediate and structurally binding influence. Decisions about layer depth, activation functions, attention mechanisms, or the specific form of graph convolution effectively determine the space of functions the model can express before any data are encountered. In materials applications, this is particularly consequential. The widespread reliance on established graph neural network templates, for example, embeds a preference for locality by design, privileging short-range interactions and symmetry-constrained message passing [4]. While often justified on computational or chemical intuition grounds, such defaults can become limiting when the target phenomena—such as long-range electronic correlations or anharmonic lattice effects—fall outside the representational bias encoded in the architecture.

This structural conditioning is then reinforced at the level of optimization, where algorithmic presets shape how the model navigates its parameter space. Choices such as optimizer configuration, learning-rate schedules, normalization schemes, and stopping criteria do more than stabilize training; they implicitly regularize the solution space. Analyses of deep learning dynamics have shown that convergence behavior is strongly influenced by these settings, even when the loss function remains unchanged [2]. In materials informatics, where loss landscapes are often high-dimensional and irregular, such defaults can steer learning toward regions that reflect optimization bias rather than intrinsic structure in the data. What emerges is not simply a fitted model, but a trajectory-dependent solution shaped by the assumptions embedded in the training procedure.

A related layer of influence originates earlier in the pipeline, within the construction of the data itself. Featurization strategies, normalization schemes, and preprocessing conventions introduce assumptions about what aspects of the material system are relevant and how they should be represented. These decisions are rarely neutral. Choices such as cutoff distances in local environments or feature-wise scaling implicitly define the statistical structure that the model will encounter. Studies examining representation choices in materials datasets highlight how such defaults can filter or distort the underlying distribution, particularly in the presence of periodic boundary conditions or heterogeneous sampling regimes [16, 17]. This effect becomes even more pronounced in transfer-learning contexts, where misalignment between source and target preprocessing conventions can dominate performance differences, overshadowing architectural or algorithmic refinements [10].

The influence of defaults reaches its most subtle yet consequential form at the stage of evaluation, where they define what counts as success. Metric selection, validation strategies, and baseline comparisons function as implicit utility functions, encoding normative judgments about acceptable error and model quality. In materials AI, commonly used regression metrics often assume error structures that do not reflect the physical realities of the properties being predicted. As a result, models may be optimized—and ultimately ranked—according to criteria that privilege statistical convenience over domain relevance. The divergence between formal specification and actual scientific utility is particularly visible at this stage, where evaluation defaults silently determine which approaches are recognized as state-of-the-art [18-22].

Taken together, these layers do not operate independently. Their effects accumulate and interact, producing a compound prior that permeates the entire modeling pipeline. What begins as a set of seemingly innocuous defaults ultimately shapes not only how models learn, but also what they are capable of knowing and how their outputs are interpreted. Recognizing this layered structure is therefore essential if defaults are to be treated as objects of design rather than as invisible background conditions. **Table 2** integrates the taxonomy of defaults with their pipeline roles, revealing how distinct implicit priors operate through different mechanisms to produce specific epistemic risks.”

Table 2. Mapping default types to pipeline stages, embedded priors, and epistemic consequences in materials AI

Default type	Pipeline stage	Embedded implicit prior
Architectural defaults	Model construction	Locality, symmetry, and structural constraints
Algorithmic defaults	Optimization	Smoothness and convergence behavior
Data defaults	Preprocessing/representation	Statistical structure of atomic environments
Evaluation defaults	Validation/testing	Error distribution, utility assumptions

The taxonomy reveals that defaults are not monolithic; each type operates at a different pipeline stage and encodes priors of differing granularity. Architectural defaults set global capacity, algorithmic defaults shape local search dynamics, data defaults filter the input distribution, and evaluation defaults adjudicate outcomes. Together they form a compound implicit prior whose total effect is greater than the sum of its parts. By distinguishing these types, the framework enables targeted scrutiny and eventual reform of materials AI practice.

Derived Properties and Corollaries

The recognition that default assumptions operate as implicit priors entails a set of consequences that extend beyond model construction into the evaluation, reproducibility, and transfer of materials AI systems. Once defaults are understood as active contributors to model behavior rather than passive background settings, they must be treated as coequal with explicit parameters in shaping outcomes. This shift reveals that many standard practices rest on an incomplete accounting of the factors that determine performance, and it exposes a series of epistemic distortions that arise when defaults remain unexamined.

A first implication becomes apparent in the interpretation of comparative studies. Benchmarking exercises that hold defaults constant across models are often assumed to isolate the effect of explicit architectural or algorithmic innovations. In practice, however, such comparisons evaluate the interaction between those innovations and a shared, underlying configuration of implicit priors. Foundational analyses of deep networks have demonstrated that architecture alone introduces inductive bias that data cannot neutralize [1], and when this bias is coupled with fixed presets for graph construction, optimization, and loss functions, the resulting performance differences reflect compatibility with the inherited defaults as much as any deliberate design choice. In materials AI, where graph neural network pipelines for crystal property prediction frequently rely on identical neighbor-cutoff schemes and pooling strategies [4], the apparent superiority of one model over another may arise from alignment with these hidden assumptions rather than from intrinsic methodological advancement. What appears as progress in leaderboard rankings or ablation studies is therefore entangled with a constant background of unreported priors, complicating claims of genuine improvement.

This entanglement carries directly into the domain of reproducibility, where the omission of defaults from documentation practices introduces a persistent source of divergence. Traditional reproducibility protocols emphasize explicit parameters such as learning rates, initialization schemes, and random seeds. Yet, they overlook the equally influential role of automatic configurations embedded within software libraries and preprocessing pipelines. Work examining the origins of inductive bias has shown that optimization dynamics themselves encode structure [2], and these dynamics are governed by defaults that are rarely surfaced in published accounts. As a result, nominally identical implementations can yield materially different outcomes when executed under slightly different software environments, where subtle changes in default behavior—whether in pooling operations, normalization routines, or symmetry handling—alter the effective prior imposed on the model. Reproducibility, under these conditions, cannot be secured through partial disclosure. It requires a comprehensive accounting in which defaults are treated as first-class variables, documented alongside explicit parameters to ensure that the full configuration of assumptions can be reconstructed.

A related consequence emerges in the context of transfer learning, where the portability of pretrained models is often attributed to shared statistical structure between source and target domains. While distributional alignment is undoubtedly important, it is insufficient to explain many observed failures of transfer in materials AI. The implicit priors encoded in default settings frequently carry assumptions that are valid in one domain but invalid in another. A model trained on bulk crystalline systems, for instance, may incorporate periodic-boundary conditions and mean-pooling operations as default representations of structure; when applied to surface phenomena or defect-dominated regimes, these assumptions become misaligned with the underlying physics, limiting the effectiveness of fine-tuning. Studies surveying machine learning applications in solid-state materials science have documented both successes and limitations of transfer [6]. Yet, they rarely foreground the role of default compatibility in determining these outcomes. Reframed through the present lens, transfer becomes less a matter of adapting learned weights and more an exercise in harmonizing implicit assumptions, where inherited defaults must be critically evaluated and, when necessary, reconfigured to match the epistemic structure of the target task [23-27].

Taken together, these implications alter the practical interpretation of results across the field. Defaults introduce systematic influences that permeate comparison, reproduction, and transfer, shaping conclusions in ways that remain largely invisible under current reporting conventions. Recognizing their role as hidden parameters elevates default management from a technical afterthought to a central theoretical concern. By making these influences explicit, the framework provides a basis for rethinking evaluation protocols and documentation standards, ensuring that materials AI systems are assessed not only in terms of their explicit design but also in light of the implicit assumptions that govern their behavior.

Why Defaults Persist

Default assumptions persist within materials AI pipelines not because their influence is negligible, but because they are sustained by a set of reinforcing dynamics that operate at cognitive, historical, and social levels. Their endurance reflects a structural equilibrium in which convenience, institutional inertia, and systematic misattribution align to keep these implicit priors largely invisible, even as they exert measurable effects on model behavior.

At the cognitive level, defaults offer an immediate reduction in complexity that is difficult to relinquish in practice. Materials AI workflows are already characterized by high dimensionality and interdisciplinary demands, and preset configurations provide a stable entry point that enables rapid experimentation without requiring exhaustive re-specification of every design choice. The appeal is pragmatic: researchers can direct attention toward scientific questions rather than toward infrastructural decisions that appear, at least initially, to be peripheral. This tendency is reinforced by the empirical reality that many default configurations produce acceptable or even strong baseline performance. As discussions of generalization in deep learning have shown, practitioners often converge on setups that “just work,” not because they are theoretically optimal, but because they reliably produce usable results under typical conditions [3]. Over time, repeated reliance on such configurations transforms convenience into tacit endorsement, embedding defaults as unexamined components of standard practice.

This cognitive stabilization is then amplified through path-dependent processes that operate at the level of the research community. Once particular defaults are adopted in influential studies and disseminated through widely used codebases, they acquire the status of de facto standards. Subsequent work builds upon these implementations, not only for efficiency but also to maintain comparability with established benchmarks. The cost of deviation becomes nontrivial: altering a default may require re-running extensive experiments, revalidating performance claims, and potentially disrupting alignment with prior results. Analyses of implicit regularization in deep learning highlight how optimization trajectories themselves become entrenched through repeated use [2], and a similar inertia governs the persistence of defaults. In materials informatics, frameworks that have standardized representations of molecular and crystalline systems have effectively canonized certain assumptions about symmetry handling and interaction structure [7]. Under these conditions, defaults are not simply convenient—they are institutionally reinforced, and their modification entails both technical and epistemic risk.

A further mechanism sustaining this persistence lies in patterns of attribution that obscure the role of defaults in shaping outcomes. When models perform well, explanatory narratives tend to emphasize explicit innovations—architectural refinements, novel features, or carefully tuned hyperparameters—while the contribution of inherited configurations remains unacknowledged. Conversely, when performance is unsatisfactory, explanations typically invoke data limitations or task complexity, leaving defaults outside the scope of scrutiny. This asymmetry is reinforced by reporting conventions that foreground deliberate design choices while omitting the automatic settings that accompany them. Surveys of machine learning applications in materials science often describe pipelines in terms of their explicit components without interrogating the default assumptions embedded within those components [5], thereby perpetuating the impression that observed behavior arises solely from visible decisions. In the absence of established norms or journal requirements for documenting defaults, there is little incentive to expose these latent influences, allowing them to operate as hidden parameters that shape results without being held accountable.

The interaction of these dynamics produces a self-reinforcing cycle in which defaults remain both pervasive and underexamined. Convenience encourages their initial adoption, path dependence stabilizes their continued use, and attribution patterns obscure their effects. Breaking this cycle requires more than incremental adjustment; it demands a shift in how defaults are conceptualized within the field. Until they are treated as integral elements of model design—subject to the same scrutiny, variation, and documentation as explicit parameters—they will continue to function as powerful yet largely invisible determinants of behavior in materials AI systems.

Relation to Formal Priors

Default assumptions in materials AI bear a close conceptual relationship to formal Bayesian priors, yet the comparison reveals as much divergence as it does similarity. Both operate as mechanisms that structure inference before data exerts its full influence, shaping what can be learned and how learning unfolds. In Bayesian formulations, priors explicitly delimit the hypothesis space through probability distributions; in practice, defaults impose analogous constraints through architectural presets, preprocessing conventions, and optimization settings that restrict representational capacity and admissible solutions. This parallel extends into the dynamics of updating. Just as a prior influences the posterior once evidence is incorporated, defaults guide gradient-based optimization toward regions of parameter space that are consistent with their embedded inductive biases, thereby shaping outcomes even in data-rich regimes. A further alignment appears at the level of scientific commitment. Bayesian priors often encode domain knowledge—assumptions about smoothness, periodicity, or sparsity—while defaults embed comparable commitments, such as privileging local atomic environments or symmetry constraints, albeit without being formally articulated [24-29].

Despite these structural affinities, the epistemic status of defaults differs in ways that are consequential for both analysis and practice. A central distinction lies in their visibility. Formal priors are declared components of a model; they are specified, justified, and subject to revision through established techniques such as hierarchical modeling or empirical Bayes. Defaults, in contrast, enter the pipeline implicitly. They are loaded automatically and frequently remain unaltered precisely because they are not recognized as choices requiring justification. This invisibility also affects how they are evaluated. Bayesian priors invite scrutiny through sensitivity analysis and probabilistic reasoning, encouraging explicit examination of how assumptions influence inference. Defaults, by comparison, are often insulated from such critique, sustained by the convention that library presets are neutral starting points rather than epistemic commitments.

The distinction deepens when one considers the scope at which each operates. Formal priors are confined to the statistical model, entering directly into the likelihood–posterior relationship that defines inference. Defaults, by contrast, permeate the entire computational workflow. They shape data representation, govern model architecture, influence optimization trajectories, and ultimately define evaluation criteria, all without being incorporated into a unified probabilistic framework. Their effects are therefore distributed and cumulative, emerging across stages rather than residing in a single formal component.

Viewed through this lens, defaults can be understood as unacknowledged priors—functionally equivalent in their capacity to constrain inference, yet lacking the transparency and critical apparatus associated with Bayesian formalism. Work examining prior distributions as default assumptions in materials AI points toward this overlap [23], but the present perspective underscores a more general point: defaults achieve comparable epistemic influence while remaining largely outside the domain of explicit analysis. Recognizing them as hidden parameters bridges the divide between formal statistical reasoning and everyday modeling practice, offering a unified framework in which both explicit priors and implicit defaults can be examined, compared, and ultimately governed.

Implications for Materials AI Practice

The theoretical analysis carries direct implications for how materials AI should be practiced, shifting defaults from invisible infrastructure to first-class objects of scrutiny. First, mandatory default documentation must become standard in methods sections. Every publication should list library versions, architectural presets, data-preprocessing flags, and evaluation protocols alongside explicit hyperparameters, thereby making the compound prior visible. Second, default sensitivity analysis should be adopted as a routine diagnostic: researchers should systematically vary key defaults—such as neighbor cutoffs or pooling operations—and report how predictive behavior changes, revealing the extent to which results depend on hidden parameters rather than explicit design.

Lastly, community standards for sensible defaults in materials-specific tasks are urgently needed. Unlike general machine-learning libraries, materials AI requires domain-tuned presets that reflect periodic boundary conditions, symmetry requirements, and extensive properties; establishing these through consensus rather than historical accident would reduce path dependence. Fourth, default auditing should become part of model reporting checklists, parallel to existing reproducibility guidelines. Journals and conferences could require authors to certify that defaults were examined and justified, thereby institutionalizing the practice of treating defaults as epistemic commitments.

These reforms would elevate materials AI beyond black-box empiricism toward genuine scientific transparency. By making defaults explicit and subject to justification and sensitivity analysis, practitioners can isolate the true contribution of novel methods and ensure that claimed advances are not artifacts of inherited assumptions. The framework, therefore, offers not only a theoretical diagnosis but a practical pathway for aligning modeling practice with the epistemic demands of materials science.

Conclusion

This paper has advanced the theoretical claim that default assumptions in materials AI function as implicit priors that shape posterior predictions as much as explicit parameters. By defining defaults as hidden parameters, distinguishing them from both explicit parameters and tunable hyperparameters, surveying their prevalence in materials pipelines, and developing a taxonomy of architectural, algorithmic, data-related, and evaluation defaults, the analysis has revealed an overlooked dimension of model epistemology. The derived corollaries on model comparison, reproducibility, and transfer learning, together with the explanation of persistence mechanisms and the clarification of defaults' relation to formal Bayesian priors, collectively demonstrate that unexamined defaults introduce systematic epistemic distortion that cannot be ignored.

The central call is therefore to treat defaults as first-class parameters subject to documentation, justification, and sensitivity analysis. Only by making the invisible visible—by surfacing the compound priors that libraries, codebases, and conventions quietly impose—can materials AI achieve the epistemic maturity demanded by the complexity of matter. Future work should extend this framework to other scientific domains and develop automated tools for default auditing, thereby transforming a hidden source of bias into a deliberate object of scientific control. In doing so, the field moves closer to a practice in which every assumption, default, or explicit, is held to the same rigorous standard of justification.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 10 Dec 2021 Revised: 20 Feb 2022 Accepted: 28 Mar 2022

References

Poggio T, Banburski A, Liao Q. Theoretical issues in deep networks. Proc Natl Acad Sci U S A. 2020;117(48):30039-45.

Gondal MW, Wuthrich M, Miladinovic D, Locatello F, Breidt M, Volchkov V, et al. On the transfer of inductive bias from simulation to the real world: A new disentanglement dataset. NeurIPS. 2019;32.

Zhang C, Bengio S, Hardt M, Recht B, Vinyals O. Understanding deep learning (still) requires rethinking generalization. Commun ACM. 2021;64(3):107-15.

Dai M, Demirel MF, Liang Y, Hu JM. Graph neural networks for an accurate and interpretable prediction of the properties of polycrystalline materials. npj Comput Mater. 2021;7(1):103.

Butler KT, Davies DW, Cartwright H, Isayev O, Walsh A. Machine learning for molecular and materials science. Nature. 2018;559(7715):547-55.

Schmidt J, Marques MR, Botti S, Marques MA. Recent advances and applications of machine learning in solid-state materials science. npj Comput Mater. 2019;5(1):83.

Chen C, Ye W, Zuo Y, Zheng C, Ong SP. Graph networks as a universal machine learning framework for molecules and crystals. Chem Mater. 2019;31(9):3564-72.

Wang F, Bian Y, Wang H, Lyu M, Pedrini G, Osten W, et al. Phase imaging with an untrained neural network. Light Sci Appl. 2020;9(1):77.

Ma C, Li Y, Hernández-Lobato JM. Variational implicit processes. In: International Conference on Machine Learning; 2019. p. 4222-33.

De Breuck PP, Hautier G, Rignanese GM. Materials property prediction for limited datasets enabled by feature selection and joint learning with modnet. npj Comput Mater. 2021;7(1):83.

Qayyum A, Ilahi I, Shamshad F, Boussaid F, Bennamoun M, Qadir J. Untrained neural network priors for inverse imaging problems: A survey. IEEE Trans Pattern Anal Mach Intell. 2022;45(5):6511-36.

Borisov V, Leemann T, Seßler K, Haug J, Pawelczyk M, Kasneci G. Deep neural networks and tabular data: A survey. IEEE Trans Neural Netw Learn Syst. 2022;35(6):7499-519.

Neyshabur B. Implicit regularization in deep learning. arXiv preprint arXiv:1709.01953; 2017.

Neyshabur B, Bhojanapalli S, McAllester D, Srebro N. Exploring generalization in deep learning. NeurIPS. 2017;30.

Liu S, Okatani T. Symmetry-aware neural architecture for embodied visual exploration. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR); 2022. p. 17221-30.

Mathew A, Amudha P, Sivakumari S. Deep learning techniques: An overview. In: International Conference on Advanced Machine Learning Technologies and Applications. Singapore: Springer; 2020. p. 599-608.

Canay IA, Romano JP, Shaikh AM. Randomization tests under an approximate symmetry assumption. Econometrica. 2017;85(3):1013-30.

Mantovani RG, Rossi AL, Alcobaça E, Gertrudes JC, Junior SB, de Carvalho AC. Rethinking default values: A low cost and efficient strategy to define hyperparameters. arXiv:2008.00025 [cs.LG]; 2020.

Polkowski LT. Epistemic, default and dynamic logics. In: Logics for computer and data sciences, and artificial intelligence. Cham: Springer; 2022. p. 105-40.

Phan HC, Dhar AS. Predicting pipeline burst pressures with machine learning models. Int J Press Vessel Pip. 2021;191:104384.

Battaglia PW, Hamrick JB, Bapst V, Sanchez-Gonzalez A, Zambaldi V, Malinowski M, et al. Relational inductive biases, deep learning, and graph networks. arXiv preprint arXiv:1806.01261; 2018.

Masters P, Vered M. What’s the context? Implicit and explicit assumptions in model-based goal recognition. In: Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI). Montreal, Canada: AAAI Press; 2021. p. 4516-23.

Van Zwet E. A default prior for regression coefficients. Stat Methods Med Res. 2019;28(12):3799-807.

Chollet F, Kalinowski T, Allaire JJ. Deep learning with R. New York: Simon and Schuster; 2022.

Feinman R, Lake BM. Learning inductive biases with simple neural networks. arXiv preprint arXiv:1802.02745; 2018.

Okereke M, Keates S. Boundary conditions. In: Finite element applications: A practical guide to the FEM process. Cham: Springer International Publishing; 2018. p. 243-97.

Atanasiu R. Hidden assumptions. In: Critical thinking for managers: Structured decision-making and persuasion in business. Cham: Springer; 2021. p. 13-26.

Min H, Tarmoun S, Vidal R, Mallada E. On the explicit role of initialization on the convergence and implicit bias of overparametrized linear networks. In: Proceedings of the 38th International Conference on Machine Learning (ICML). PMLR; 2021. p. 7760-8.

Boon M. How scientists are brought back into science-The error of empiricism. In: A critical reflection on automated science: Will science remain human? Cham: Springer; 2020. p. 43-65.