

Attesting Model Lineage by Consisted Knowledge Evolution with Fine-Tuning Trajectory

Zhuoyi Shang^{1,2,3,*}, Jiasen Li^{1,2,3,*}, Pengzhen Chen^{1,2,3,*},
Yanwei Liu^{1,3,†}, Xiaoyan Gu^{1,2,3,†}, Weiping Wang¹

¹*Institute of Information Engineering, Chinese Academy of Sciences, Beijing, China*

²*School of Cyber Security, University of Chinese Academy of Sciences, Beijing, China*

³*State Key Laboratory of Cyberspace Security Defense, Beijing, China*

{shangzhuoyi, lijiasen, chenpengzhen, liuyanwei, guxiaoyan, wangweiping}@iie.ac.cn

Abstract

The fine-tuning technique in deep learning gives rise to an emerging lineage relationship among models. This lineage provides a promising perspective for addressing security concerns such as unauthorized model redistribution and false claim of model provenance, which are particularly pressing in open-weight model libraries where robust lineage verification mechanisms are often lacking. Existing approaches to model lineage detection primarily rely on static architectural similarities, which are insufficient to capture the dynamic evolution of knowledge that underlies true lineage relationships. Drawing inspiration from the genetic mechanism of human evolution, we tackle the problem of model lineage attestation by verifying the joint trajectory of knowledge evolution and parameter modification. To this end, we propose a novel model lineage attestation framework. In our framework, model editing is first leveraged to quantify parameter-level changes introduced by fine-tuning. Subsequently, we introduce a novel knowledge vectorization mechanism that refines the evolved knowledge within the edited models into compact representations by the assistance of probe samples. The probing strategies are adapted to different types of model families. These embeddings serve as the foundation for verifying the arithmetic consistency of knowledge relationships across models, thereby enabling robust attestation of model lineage. Extensive experimental evaluations demonstrate the effectiveness and resilience of our approach in a variety of adversarial scenarios in the real world. Our method consistently achieves reliable lineage verification across a broad spectrum of model types, including classifiers, diffusion models, and large language models.

1 Introduction

The rapid advancement of deep learning (DL) has been largely driven by a paradigm shift from traditional *training-from-*

* These three authors contributed equally to this work.

† Corresponding author.

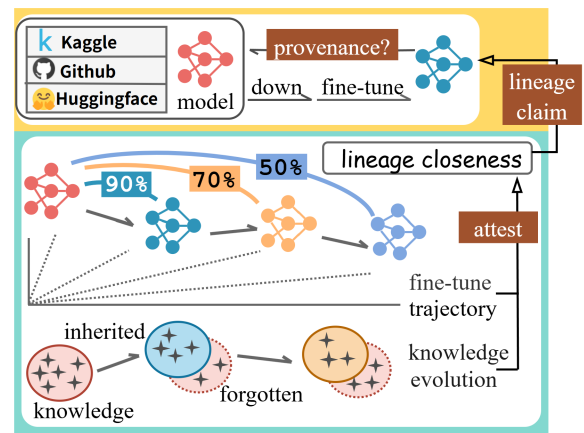


Figure 1: The knowledge evolution mechanism in model fine-tuning. The consistency between the knowledge of parent model and the total knowledge inherited and discarded during the evolution process is used to attest the lineage relationship.

scratch approaches to the *pre-training–fine-tuning* paradigm. This shift enables efficient domain adaptation, particularly for large-scale models, by allowing pre-trained models to be minimally fine-tuned for alignment with specific tasks or human preferences. As a result, fine-tuning introduces a new form of dependency among models, known as *model lineage* [51], where multiple rounds of successive fine-tuning create inter-generational relationships between models.

However, alongside the benefits of fine-tuning, new security challenges have emerged: (i) unauthorized redistribution via derivative models [10], (ii) adversarial fine-tuning for harmful or unethical purposes [30], and (iii) false claim of model provenance [26], especially in commercial settings. The rise of open-weight model repositories (e.g., GitHub and Hugging Face [48]) [35, 38, 40, 44], has further amplified these risks, as current platforms often lack effective lineage attestation mechanisms. This creates opportunities for both inadvertent misuse and deliberate obfuscation of true model ancestry.

Regarding these security concerns, it is critically impor-

tant to develop robust forensic techniques for model lineage attestation. Such tools are essential for intellectual property (IP) protection [4, 9, 23, 46], model authentication [11], and defect tracing [22]. They can provide conclusive evidence of inheritance between the base and derived models, serving as a foundation for legal and ethical accountability. For example, if a model owner suspects that a publicly released model is derived from his proprietary model, he may first conduct behavioral comparisons through black-box access [31]. *Upon gathering sufficient suspicion, the owner could request white-box access through legal channels to extract unique model-specific identifiers and assert ownership.* However, existing authentication methods struggle with model lineage verification due to the disruptive nature of fine-tuning on both parameters and behavioral traits.

Furthermore, recent advancements in model lineage detection [51, 52] have yet to yield unambiguous, verifiable evidence of lineage. These approaches remain vulnerable to dynamic adversarial spoofing, as they primarily rely on static architectural similarities (feature-similarities or parameter similarities) between models. Such structural representations fail to capture the deep, dynamic transformations in a model’s internal knowledge representation during fine-tuning. As a result, these approaches often yield suboptimal attribution performance. Indeed, fine-tuning is not merely a process of parameter adjustment—it reflects a continuous evolution of model knowledge, guided by the semantics of the fine-tuning dataset. *This highlights two key challenges for lineage attestation: (1) how to effectively characterize the intrinsic variations in model knowledge throughout the fine-tuning process, and (2) how to leverage the dynamic process of knowledge transformation for reliable and interpretable lineage attribution.*

To address the above two challenges, we propose to attest model lineage by exploring the knowledge evolution in a fine-tuning process. Our approach is inspired by biological evolution, which serves as a powerful metaphor for model lineage. In human genetics, genes—carriers of hereditary information—are inherited from parents, yet evolution entails a dynamic process where some genetic traits are preserved across generations while others are gradually forgotten. As illustrated in Fig. 1, in the context of model evolution, knowledge functions like genes: the fine-tuned model inherits knowledge from its parent, while the fine-tuning process inevitably alters or forgets certain portions of the original knowledge. Therefore, lineage similarity can be measured by examining the consistency between the knowledge of parent model and the totality of inherited and discarded knowledge during the evolution process.

A model’s knowledge is structurally encoded in its parameters, enabling us to capture the knowledge variations in fine-tuning process at the parameter level. Fortunately, model editing techniques [16] allow the parameters of a fine-tuned (child) model to be decoupled into those of its parent model

and a task vector derived by fine-tuning. The additional task vector can be instantiated as a neural network, referred to an evolution model, delineating knowledge variations caused by fine-tuning. Then we can encode the intermediate or final responses of both the evolution model and the original/fine-tuned models into knowledge vectors using probing datasets. The probing strategies are tailored to different model structures, including decision-boundary probing for classification models, distributional sampling for diffusion models, and response-based probing for LLMs. Our intuition is that a model’s behavior on specific inputs reflects its internal knowledge, enabling an arithmetic way to measure the consistency of knowledge evolution with fine-tuning trajectory.

In summary, we first explicitly parameterize the knowledge modifications introduced during fine-tuning. Building on this, we propose an effective knowledge vectorization mechanism to embed the edited knowledge into a shared latent space, that offers a principled foundation for model lineage attestation, where the consistency between the fine-tuning trajectory and the knowledge evolution path forms the necessary condition for validating model lineage relationships.

Our main contributions are summarized as follows:

- Inspired by the genetic mechanism of human evolution, we uncover for the first time the knowledge evolution mechanism in the model fine-tuning process, grounding the effective model lineage verification in internal representational evidence.
- We propose a novel model lineage attestation methodology by exploring consisted knowledge evolution with fine-tuning trajectory. To verify the knowledge consistency, we devise a novel vectorization mechanism to embed the models’ knowledge in a shared latent space, embracing a variety of model types. Our methodology can capture the intrinsic genealogical links among models, outperforming those with only comparing superficial resemblance among models.
- We perform comprehensive experimental evaluations for two prominent adversarial scenarios. Our method demonstrates robust lineage verification across a wide spectrum of model types, including classification models, diffusion models, and diverse large language models, showing strong resilience against real-world adversarial manipulations.

2 Related Work

2.1 Neural Network IP Protection

To detect unauthorized copies of protected models, several techniques are commonly employed for DL model IP protection. One kind of methods is model watermarking

[1, 19, 27, 53], which involves embedding a detectable identifier into the neural network through processes such as fine-tuning or retraining. They either embed secret information directly into the model weights or introduce backdoor attacks to manipulate training samples to deliberately alter the decision boundaries of the neural network to protect its intellectual property. However, both methods require modifying the original model’s parameters. The former is vulnerable to changes caused by attacks such as pruning and fine-tuning, while the latter lacks robustness when facing downstream tasks that require altering decision boundaries, such as fine-tuning. Another category of technique is model fingerprinting [4, 6, 33, 34, 49], which aims to identify a set of specific samples that trigger similar responses in both the original and pirated models. The classification boundary of a DNN classifier serves as its unique representation. Leveraging this insight, IPGuard [4] identifies data points close to the model owner’s classification boundary and uses them to create a distinctive fingerprint for the classifier. If another DNN classifier produces matching predictions for the majority of these fingerprint data points, it is deemed to be an unauthorized copy of the original classifier. From the perspective of knowledge transfer in DL models, Tian et al. [46] observed the knowledge transferred from training data to a DL model can be uniquely represented by the model’s decision boundary, and they further proposed that the distance from the central training samples to the boundary decision samples can serve as the model’s fingerprint for its IP audit. Although their scheme shows resistance to some fine-tuning attacks and is more robust than IPGuard [4], it remains vulnerable when fine-tuning fails to restore the model’s decision boundary.

2.2 Fine-tuning Relationship Detection

At present, there is relatively limited research on detecting fine-tuning relationship among models. Typically, Neural Lineage [51] and Neural Phylogeny [52] are used to delineate the fine-tuning relationships. As a seminal work, Yu et al. [51] have defined the following branches of study on neural lineage: neural lineage refers to the relationship between a parent model and its child model, specifically identifying which parent model a given child model has been fine-tuned from. In neural lineage, fine-tuning is defined as a standard transfer learning practice: adapting a model from the original dataset to a new task dataset, with the requirement that the fine-tuned model performs well on the new task. The neural lineage detection in [51] relies on access to fine-tuning datasets for validation. This requirement is impractical in real-world model protection scenarios, as malicious publishers are unlikely to disclose their datasets used for fine-tuning. To further identify the direction of fine-tuning relationship, the concept of neural phylogeny detection is then proposed in the succeeding work of Yu et al. [52]. This method relies on static architectural similarities among fine-tuned models, which struggles to grasp

the dynamic knowledge associations among models that truly characterizes the lineage relationships. In this work, for clarity in our discussion, we use the notation of "model lineage" to uniformly describe fine-tuning relationships between models, encompassing directed fine-tuning dependencies.

2.3 Task Arithmetic for Model Editing

Task Arithmetic [12, 16, 32] has been explored as a powerful tool for analyzing and leveraging task relationships in multi-task learning. Its core idea is to treat tasks as vectors and explore how their combinations influence the behavior of the model. One key advantage of task arithmetic lies in its ability to decouple knowledge components about the multi-task models. The performances of pre-trained models on different datasets can be independently modulated by adding or removing task vectors [16], where "addition" improves task-specific performance and "removal" may lead to forgetting. This model editing property provides an arithmetic approach to analyze fine-tuning dependencies between models, directly forming technical basis for model lineage verification in the paper.

3 Problem Statement and Threat Model

3.1 Problem Statement

The purpose of Model Lineage Attestation (MLA) for DL models is to determine whether a given suspect model f_C is a fine-tuned version of a parent model f_P . When the owner P of the parent model f_P discovers a suspected model f_C , he can ask a trusted authentication authority to verify the lineage relationship. The attestation task is therefore to provide a reliable criterion that distinguishes models belonging to the same lineage from those that do not.

Specifically, the goal of MLA is to construct an appropriate *lineage-aware metric space* in which models from the same family exhibit consistently high lineage similarity, while models from different families exhibit low lineage similarity. Under a decision threshold T , the attestation framework aims to induce a similarity metric $\mathcal{S}(\cdot, \cdot)$ space such that f_C and f_P belong to the same family if $\mathcal{S}(f_C, f_P) \geq T$.

3.2 Threat Model

We consider model lineage attestation under two key threat scenarios that focus on adversarial attempts to manipulate model lineage. They are likely often encountered in open-weight model hubs and we illustrate them in Fig. 2. To detail which assets each party may access or manipulate, a fine-grained summary of the capabilities of each role under these threat scenarios is offered in Table 1.

Table 1: Roles and accessible assets under different threat scenarios. In the Illegal Model Derivative setting, the adversary derives and modifies f_C from a stolen parent model without access to the key evidence. In the False Lineage Claim setting, the adversary attempts to fabricate a lineage relationship by supplying forged evidence.

	f_P	f_C	Attest evidence
Defender	Submitted	Submitted	Trusted & Held
Adversary (Illegal Model Derivative)	×	Derived & Modified	×
Adversary (False Lineage Claim)	Forged	✓	Forged

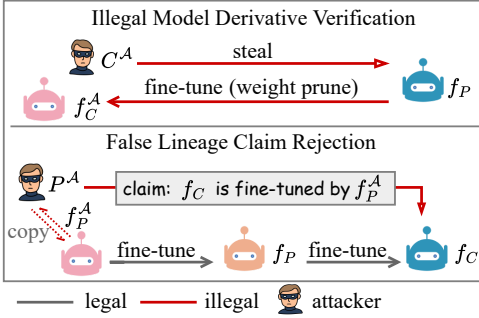


Figure 2: Two key threat scenarios for model lineage attestation. In the illegal derivation scenario, the adversary steals the victim model f_P and derives a new model f_C^A with direct fine-tuning or pruning followed by fine-tuning. In the false lineage claim scenario, the attacker attempts to present a closely related model f_P^A to falsely declare that a victim model f_C was derived from it.

3.2.1 Illegal Model Derivative Verification

A model owner trains and deploys a proprietary model f_P , while an adversary C_A may steal it and release a tampered or fine-tuned variant f_C^A . The model owner therefore requires a reliable attestation mechanism to verify lineage relationships.

Adversary’s Assumptions. The adversary is assumed to have *white-box access* to the victim model f_P , including its architecture and parameters, with sufficient resources to perform fine-tuning. For complex models such as large language models or diffusion models, they may apply advanced fine-tuning techniques such as DreamBooth [36], LoRA [15], or parameter merging to modify the stolen model. Moreover, the adversary may be aware of the lineage attestation technique and further adopt evasion strategies like weight pruning or fine-tuning across multiple generations to obscure lineage. However, C_A does *not* have access to the original training data or configurations of f_P , which constrains their ability to perfectly reproduce the model’s provenance.

Defender’s Assumptions. We assume the defender is a trusted lineage attestation authority whose goal is to verify whether a suspicious model f_C^A is derived from a proprietary model f_P . The authority may receive f_P and any auxiliary ev-

idence (probe dataset and the initialize parameters) provided by the model owner. The defender can obtain the parameters of f_C^A legally by filing a complaint to a judicial body. According to the principle “he who asserts must prove,” it is reasonable for the claimant to provide auxiliary evidence to substantiate ownership. The authority has no knowledge of the adversary’s fine-tuning strategies or modifications. Under these assumptions, the attestation mechanism can determine the lineage relationship between f_C^A and f_P .

3.2.2 False Lineage Claim Rejection

An adversary P_A may attempt to falsely claim that an existing victim model f_C was derived from their model f_P^A , aiming to gain undeserved credit. Consequently, a fair lineage attestation authority would reject any unfounded claims.

Adversary Assumptions. The adversary may attempt to gain undeserved credit by presenting a model f_P^A as the parent of a victim model f_C . To do so, they may steal or copy a close-lineage model of f_C within the same family, such as a grandparent model, along with its corresponding auxiliary evidence, such as a probe dataset or initialization parameters. The adversary has white-box access to f_C , but lacks knowledge of its training configuration. They cannot access the true parent model f_P of f_C and are unaware of the lineage closeness between f_P^A and f_C . Parent-side artifacts are maintained by the model platforms, which prevents the adversary from altering or fabricating these artifacts, including the initialization θ_0 .

Defender’s Assumptions. The defender (lineage attestation authority) is aware of f_C ’s parameters and also its training configuration, and he has also white-box access to f_P^A legally. By comparing the lineage similarity between f_P^A and f_C , the lineage verification can identify assertions where f_C and f_P^A do not share a direct parent-child relationship.

4 Knowledge Evolution in Model Fine-tuning

Fine-tuning a neural network encapsulates knowledge into parameters through a process of selection and consolidation in the training, where the child model selectively inherits knowledge from the parent model, guided by the objective of the fine-tuning task. Critically, model editing theory reveals that

fine-tuned parameters remain task-disentangled—enabling us to systematically analyze the knowledge evolution process through parameter dynamics.

Characterizing Knowledge Evolution through Task Editing. Since neural network parameters can be decomposed into distinct components corresponding to different training tasks [16, 21], the fine-tuning trajectory can be similarly characterized by changes in these parameters, even under nonlinear optimization dynamics [21]. Given a submodel $f_C(\cdot; \theta_C)$ with parameters θ_C , and its parent model $f_P(\cdot; \theta_P)$ with parameters θ_P , the parameter difference from θ_C to θ_P during fine-tuning can be denoted as the task vector $\Delta = \theta_P - \theta_C$. Δ effectively encodes the model evolution process introduced by retraining, which provides a solution to quantify the unknown knowledge evolution during fine-tuning process. Therefore, we can find a surrogate model $f_\Delta(\cdot; \theta_\Delta)$ (called evolution model) with parameters θ_Δ that encodes the evolutionary knowledge offset during the fine-tuning process. This model is constructed within the same search space spanned by the initial parameters θ_0 of f_P as:

$$\theta_\Delta = \theta_0 + \theta_P - \theta_C. \quad (1)$$

Therefore, for a submodel f_C that is fine-tuned from f_P , f_C 's behavior explained from the perspective of the parent model's knowledge is influenced by two parts: one is the direct knowledge transfer from the parent model, the other is the unknown knowledge evolution during the fine-tuning process. Thus, from the perspective of knowledge inheritance, knowledge consistency among these models can be expressed as:

$$\mathcal{K}_P(f_P) = \mathcal{K}_P(f_C) + \mathcal{K}_P(f_\Delta), \quad (2)$$

where \mathcal{K}_P denotes the specific knowledge encoded in f_P , which is tied to its parameters. Specifically, $\mathcal{K}_P(f_P)$ refers to f_P 's understanding of the learned knowledge from f_P 's training dataset D_P , and $\mathcal{K}_P(f_C)$ indicates the child model's observation of the parent model's knowledge. Due to the impact of fine-tuning, the child model inevitably forgets or alters the knowledge of the parent model, which is denoted as $\mathcal{K}_P(f_\Delta)$.

Consistent Knowledge Evolution with Fine-tuning Trajectory. No matter how descendants evolve, a family lineage always traces back to the same ancestral origin, as encoded in their shared genes. Similarly, if a child model f_C is fine-tuned from a parent model f_P , no matter what downstream task is used for fine-tuning, its initialization parameters are necessarily determined by the parent model, even minor changes in the initialization parameters can lead to divergent convergence paths [43]. Therefore, the training trajectory of a model family is unique and difficult to replicate.

Although models from distinct lineages may exhibit similar knowledge due to potential overlaps in their training datasets, Theorem 1 claims that the \mathcal{K}_P knowledge shift that consists with the fine-tuning trajectory described by model parameters modification is an inherent property of model fine-tuning, and

it accordingly provides a intrinsic signature for the model lineage attestation.

Theorem 1. *For a pre-trained model f_P , a fine-tuned model f_C , and the generated model f_Δ calculated through task arithmetic by Eq. (1), that the \mathcal{K}_P knowledge shift from f_P to f_C is consistent with the evolved \mathcal{K}_P knowledge in f_Δ is a necessary condition for confirming that f_C is fine-tuned from f_P .*

Proof. For any true descendant, the parameter update $\Delta = \theta_P - \theta_C$ lies on the task-optimization trajectory and thus satisfies (or approximately satisfies) the task-aligned weight-decoupling structure required for task arithmetic [12, 16, 21]. Assume, for contradiction, that a non-descendant model f'_C with weight θ'_C also satisfies Eq. (2). Since the parameter difference $\Delta' = \theta_P - \theta'_C$ is not generated by fine-tuning the parent on any task, it is not on the task optimization path. Thus, it cannot satisfy the weight-decoupling property, as it violates the fundamental assumption of task arithmetic. Accordingly, f'_Δ fails to act as a valid evolution model explaining the parent knowledge discrepancy among f'_C and f_P , making Eq. (2) impossible. Thus, Eq. (2) holds only if f_C is a fine-tuned descendant of f_P . \square

Importantly, this argument extends beyond the idealized linear fine-tuning case. Realistic fine-tuning may update parameters non-linearly, but it preserves the alignment of Δ with the parent-task knowledge direction, which is the structural property required by Eq. (2). In contrast, Δ' from a non-descendant achieves no such task-induced alignment and cannot satisfy even approximate weight-decoupling. Because the directional alignment is maintained only for true descendants and absent for non-descendants, Eq. (2) continues to hold under realistic fine-tuning dynamics.

5 Model Lineage Attestation

To obtain a more accurate lineage similarity $\mathcal{S}(f_C, f_P)$, we devise a model lineage attestation framework by examining the consistency of knowledge evolution along the fine-tuning trajectory, as shown in Fig. 3. We first parameterize the fine-tuning trajectory as a virtual neural network inspired by task arithmetic theory. Then, we develop an effective knowledge vectorization method, embedding neural network knowledge into a shared latent space for exploring knowledge consistency. The resulting vector space enables us to reliably verify the lineage relationship between the original and fine-tuned models by comparing the resulting knowledge consistency score with a specific threshold T .

5.1 Constructing Evolution Model f_Δ

To depict the process of knowledge evolution for model lineage attestation, the prerequisite is to construct an evolution

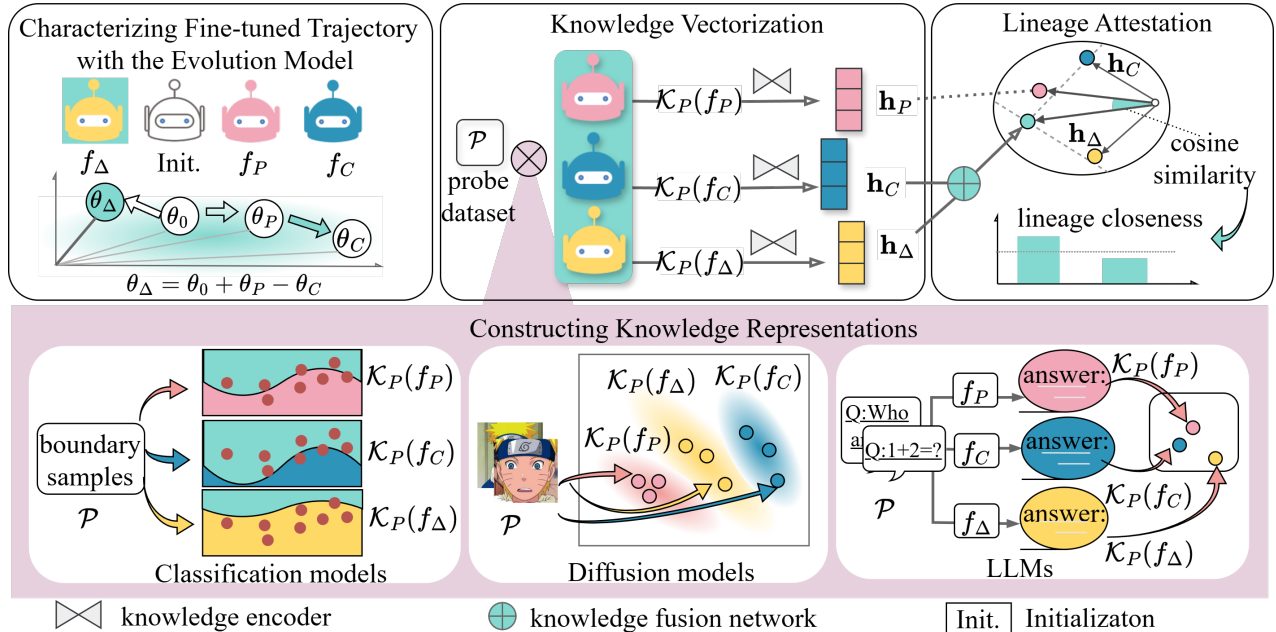


Figure 3: The proposed model lineage attestation (MLA) framework consists of three key components. First, given a parent model f_P and its fine-tuned model f_C , MLA constructs an evolution model f_Δ to characterize knowledge evolution during fine-tuning. Then, all three models are encoded into vector representations (\mathbf{h}_P , \mathbf{h}_Δ , \mathbf{h}_C) through a shared knowledge encoder. Finally, the lineage attestation is performed by evaluating the knowledge consistency condition in Equation (2).

model f_Δ with model editing theory. To this end, we reinterpret the parameter transformation between two models as a neural network $f_\Delta(\cdot; \theta_\Delta)$ in terms of Eq. (1). Thus, f_Δ effectively characterizes the fine-tuned trajectory, and the forward inference response on D_p can be leveraged to probe the changes in its knowledge about D_p induced by fine-tuning.

5.2 Model Knowledge Vectorization

To refine the parametric knowledge in the model, we propose a novel probing-based knowledge vectorization methodology that captures the model’s decision or functional behaviors associated with the probing dataset. This methodology offers an effective mechanism for representing the knowledge in model evolutions that strongly associates with parent model’s training dataset. Given the variations across different model types, we devise specialized knowledge vectorization approaches tailored to each type of model. For classification models, knowledge resides in structured decision boundaries learned from the training dataset. For diffusion models, knowledge is the learned data distribution. We probe by sampling from this distribution to test generative fidelity. Due to tokenizer-induced misalignment, we probe LLM knowledge via behavior-level responses on multi-domain benchmarks. Despite differences in model types, the key steps in vectorizing models remain the same: (1) creating a probing dataset derived from the parent model’s training data, and (2) encoding the model’s

decision/response patterns from this dataset into knowledge vectors.

5.2.1 Knowledge Vectorization of Classification Models

The knowledge learned by a classification model during training is implicitly captured by its decision boundaries, which delineate class-specific regions in the input space.

Constructing Representative Probe Samples. The decision boundary of a neural network encodes its unique learned knowledge, such that samples lying close to this boundary offer a concise representation of the knowledge it has acquired [46]. Since fine-tuning typically alters the parent model’s decision boundary, the child model’s responses to these boundary samples naturally reflect the changes induced by fine-tuning. In particular, by examining the embeddings or predictions on the same probe dataset, we can infer how the original knowledge has been preserved, modified, or lost. On the other hand, centroid samples capture the core semantics of each class. Accordingly, we construct a relatively small probe dataset \mathcal{P} consisting of boundary samples and centroid samples derived from f_P .

Assume the parent model $f_P = g(\varepsilon(\cdot))$ is a k -class DNN classifier, where $\varepsilon(\cdot)$ denotes the feature extractor and $g(\cdot)$ denotes the classification head. Formally, let $\{g_i\}$ denote the decision functions of the target classifier of class i , and a data point x is on the target classifier’s classification boundary if at least two labels have the largest discrimination probability,

i.e., $g_j(x) = g_{j'}(x) \geq \max_{i \neq j, j'} g_i(x)$, where i, j, j' are the category indexes, and $g_i(x)$ is the probability that sample x belongs to category i [4]. Given the training dataset D_P , we construct the probe sample set \mathcal{P} as:

$$\mathcal{P} = \{x_0^0, x_0^1, x_0^2, \dots, x_i^i, x_i^j, \dots, x_k^k\}, \quad (3)$$

where x_i^i is the centroid sample of class i , and x_i^j (with $i \neq j$) denotes a boundary sample that transits from class i to j .

Extracting Feature-Encoded Knowledge via Decision Boundary Probes. To analyze how the parent model’s knowledge evolves through fine-tuning, we begin by examining the feature embeddings that the parent f_P , the evolution model f_Δ and the fine-tuned child f_C extract from the carefully chosen samples. The feature embeddings produced for these samples reflect the model’s ability about the corresponding knowledge, thereby providing an indirect yet effective means to analyze the specific knowledge relationship. Therefore, the parent model’s knowledge $\mathcal{K}_P(f_P)$ with respect to \mathcal{P} is represented by:

$$\mathcal{K}_P(f_P) = \{\varepsilon_P(x) \mid x \in \mathcal{P}\}. \quad (4)$$

Fine-tuning typically alters a model’s decision boundary. We can thus capture the inherited knowledge of the child model by observing its behavior on the parent-derived probe set \mathcal{P} . Specifically, the positions of the parent model’s boundary samples within the child model reflect how fine-tuning has modified the decision regions. Accordingly, the f_C ’s inherited knowledge is represented by:

$$\mathcal{K}_P(f_C) = \{\varepsilon_C(x) \mid x \in \mathcal{P}\}, \quad (5)$$

where $\varepsilon_C(\cdot)$ denote the child model’s feature extractor.

Similarly, the evolution model f_Δ characterizes the transformation from the parent model to the child model, and its behavior on the parent-derived probe set \mathcal{P} captures knowledge that has been lost or modified during fine-tuning. Formally, the knowledge representation $\mathcal{K}_P(f_\Delta)$ is defined as:

$$\mathcal{K}_P(f_\Delta) = \{\varepsilon_\Delta(x) \mid x \in \mathcal{P}\}. \quad (6)$$

Encoding Feature Knowledge into Vectors. To encode semantic dependencies across categories, we construct a knowledge encoder Ψ , in which categories are modeled as sequences and each class occupies a unique position. Feature embeddings are first projected into a shared latent space via a multi-layer perceptron (MLP) block \mathcal{M} . A Transformer encoder \mathcal{T} is then applied k times—once per class—taking the centroid and corresponding boundary samples as input to capture local inter-class relations. The last-token embedding from each sequence is used as the class-specific knowledge vector, and the final representation is obtained by averaging these vectors via a global pooling layer \mathcal{G} . This forms a hierarchical knowledge encoder denoted as $\Psi = \mathcal{G}(\mathcal{T}(\mathcal{M}(\cdot)))$. Formally, for model $f_F \in \{f_P, f_C, f_\Delta\}$, we then obtain the compact knowledge vectors by applying the hierarchical knowledge encoder Ψ :

$$\mathbf{h}_F = \Psi(\mathcal{K}_P(f_F)), \quad F \in \{P, C, \Delta\}. \quad (7)$$

5.2.2 Knowledge Vectorization for Diffusion Models and LLMs

During the forward inference process, the model encodes each input into a semantically rich embedding in its intermediate layers. For multi-domain models, such as diffusion models and large language models, which are trained on massive datasets, it is nearly impossible to access the full training data. Following prior work [24], the model’s knowledge can be defined as the semantic concepts encoded in its feature representations of the probe inputs at a given layer, since these representations reflect what the model has learned to distinguish the query input. However, internal probes (logs or representations) become infeasible when child models alter tokenizers for LLMs, causing internal representation misalignment. We therefore evaluate the consistency of model knowledge by comparing their responses to the same set of probe inputs [7], which serves as a robust proxy for internal knowledge.

Constructing Feature Knowledge Vector for Diffusion Models. For diffusion models, we create a probing dataset, $\mathcal{P} = \{x_1, x_2, \dots, x_n\}$, which consists of n randomly selected images from D_P . The generative capabilities of these models are fundamentally driven by the denoising backbone of the U-Net architecture [41]. Different diffusion models, however, often utilize U-Nets with significantly varied architectures or training objectives, with these differences being substantial across model families. These architectural and objective variations result in distinct intermediate feature representations when given the same input data. Therefore, we leverage the output features of the UpBlock layers within the U-Net as meaningful embeddings to capture this unique knowledge. Let $\varepsilon(\cdot)$ denotes the feature extraction process from the final UpBlock layer of the models, we can instantiate the knowledge, resulting in $\mathcal{K}_P(f_\Delta)$, $\mathcal{K}_P(f_C)$, and $\mathcal{K}_P(f_P)$ as previously established.

To transform the extracted features into compact knowledge representations, we apply a knowledge encoder Ψ that aligns and aggregates feature vectors. Due to the spatial nature of extracted features, we first apply a convolutional layer \mathcal{C} to convert the 3D features into vectors, followed by global average pooling \mathcal{G} , and an MLP block \mathcal{M} for alignment: $\Psi = \mathcal{M}(\mathcal{G}(\mathcal{C}(\cdot)))$. Finally, the knowledge vectors $\mathbf{h}_P, \mathbf{h}_C, \mathbf{h}_\Delta$ for each model is obtained as Eq. (7).

Acquiring Response-based Knowledge Vector for LLMs. We estimate the model knowledge by analyzing the responses of the parent model f_P , child model f_C , and evolution model f_Δ to the same inputs. The probe dataset is defined as $\mathcal{P} = \{x_1, x_2, \dots, x_n\}$, where each query x_i is used to elicit responses from the models f_P, f_C , and f_Δ . To ensure diverse coverage and comprehensively acquiring a model’s knowledge, \mathcal{P} spans multiple domains—such as general Q&A, mathematics, and reasoning. To account for the randomness in LLM generation, 5 responses are collected for each question via in-

dependent inference. Therefore, for a model $f_F \in \{f_P, f_C, f_\Delta\}$, we define its knowledge representation over \mathcal{P} as:

$$\mathcal{K}_P(f_F) = \{f_P(x)^{(r)} \mid x \in \mathcal{P}, r = 1, \dots, 5\}, \quad (8)$$

where $f_P(x)^{(r)}$ denotes the encoded embedding of the r -th response to query x encoded by the parent model’s encoder.

Since LLM progressively refines input representations through stacked layers, and the final hidden states produce a contextualized embedding that captures its understanding of the input prompt, which is subsequently decoded into a response. To obtain a concise representation of a model’s knowledge, we transform the generated responses to each input into continuous embeddings, which act as compact knowledge vectors. In detail, all responses are encoded using the parent model’s encoder ε_P to remove semantic differences across models and map them into a common embedding space, enabling us to better observe how knowledge from the parent model is preserved or transformed in the child and evolution models. For each model f_F , we then aggregate the resulting embeddings using the knowledge encoder Ψ to obtain the final knowledge vector \mathbf{h}_F of the model:

$$\mathbf{h}_F = \Psi(\varepsilon_P(\mathcal{K}_P(f_F))), \quad F \in \{P, C, \Delta\},$$

where $\Psi = \mathcal{G}(\mathcal{M}(\cdot))$ consisting of an MLP layer \mathcal{M} for refining and aligning the initial response-based embeddings and a global average pooling layer \mathcal{G} for aggregating them into the compact knowledge representation.

5.3 Model Lineage Attestation

In the vectorized metric space, the knowledge vectors \mathbf{h}_C , \mathbf{h}_Δ , and \mathbf{h}_P for a fine-tuning process should align with Eq. (2), such that the fusion vector of \mathbf{h}_C and \mathbf{h}_Δ is aligned with \mathbf{h}_P . However, practical limitations arise because of the imperfect knowledge extraction. Firstly, from the perspective of task arithmetic, f_Δ is constructed along the fine-tuning direction guided by the fine-tuning dataset, and is therefore semantically misaligned with the probe dataset, since the latter is semantically aligned with the parent model. Secondly, the evolution model f_Δ generated through model editing only approximates the model’s optimization trajectory and does not perfectly represent the changed knowledge. Consequently, the extracted feature embeddings may contain noise and fail to fully preserve the knowledge of the parent model.

Constructing Knowledge Alignment Space. To mitigate these limitations and refine the approximation, we introduce a trainable knowledge fusion network Φ , which is employed as a linear fully connected network, to further aggregate the knowledge in \mathbf{h}_C and \mathbf{h}_Δ and enable knowledge-consistency alignment. Φ is designed as:

$$\Phi(\mathbf{h}_C, \mathbf{h}_\Delta) = \mathcal{M}(\mathbf{h}_C \parallel \mathbf{h}_\Delta), \quad (9)$$

where \parallel denotes the concatenation operation, and \mathcal{M} is a learnable MLP block that operates on the combined vector.

Overall, we design the attestation framework by jointly optimizing the parameters θ_Ψ of knowledge vectorization network Ψ and θ_Φ of knowledge fusion network Φ to construct a metric space that preserves knowledge consistency between parent and child models. In this space, a higher similarity score indicates greater confidence that the fused knowledge $\Phi(\mathbf{h}_C, \mathbf{h}_\Delta)$ aligns with the parent model’s knowledge \mathbf{h}_P .

Specifically, true parent–child pairs and non-direct pairs are used as positive and negative samples, respectively, to learn the optimal parameters θ_Φ^* and θ_Ψ^* as:

$$(\theta_\Phi^*, \theta_\Psi^*) = \arg \min_{\theta_\Phi, \theta_\Psi} \mathcal{L}, \quad (10)$$

$$\mathcal{L} = \text{sim}(\mathbf{h}_P, \Phi(\mathbf{h}_C, \mathbf{h}_\Delta)) - \text{sim}(\mathbf{h}_P, \Phi(\mathbf{h}'_C, \mathbf{h}'_\Delta)) + m,$$

where \mathbf{h}'_C and \mathbf{h}'_Δ denote the embeddings of a negative child model and its corresponding negative evolutionary direction, respectively, and $\text{sim}(\cdot, \cdot)$ denotes cosine similarity. The margin $m = 0.2$ imposes a minimum distance between positive and negative pairs, thereby encouraging a more discriminative knowledge alignment space.

Finally, given an unknown parent–child model pair (f_P, f_C) , the knowledge consistency score is calculated as:

$$S(f_P, f_C) = \text{sim}(\mathbf{h}_P, \Phi(\mathbf{h}_C, \mathbf{h}_\Delta)). \quad (11)$$

Then the lineage can be claimed if and only if $S(f_P, f_C)$ is above a pre-defined threshold T .

Analysis of Cross-Generational Knowledge Alignment. Models that undergo multiple successive fine-tuning stages often exhibit across-generational lineage (e.g., grandparent–grandchild), which can also be assessed via Eq. (11). However, as fine-tuning progresses, knowledge inherited from the parent model gradually diminishes, making lineage attestation increasingly difficult—especially in skip-generation cases where semantic alignment weakens.

Consider a model f_P fine-tuned to f_{C1} , and then to f_{C2} . Let $f_\Delta = \theta_1 + \theta_2 + \theta_0$ denote the evolution model obtained by task arithmetic, where θ_1 and θ_2 are task vectors of the two stages, respectively, θ_0 is the initial parameters of f_P . As fine-tuning progresses across generations, the child model f_{C2} inherits progressively less knowledge from f_P , and more of the original information is changed, which is captured by the task-vector corresponding to the evolution model f_Δ . Therefore, the information loss caused by the semantic inconsistency between the evolution model and the parent model, as well as the alignment bias introduced by the mismatch between model editing and the true optimization trajectory, accumulates progressively across generations.

This parameter-level accumulation of task vectors implies that the fusion knowledge vector $\Phi(\mathbf{h}_C, \mathbf{h}_\Delta)$ increasingly deviates from the parent model’s original knowledge vector \mathbf{h}_P . The degree of this semantic drift is positively correlated with the number of fine-tuning generations. As a result, the similarity score between \mathbf{h}_P and $\Phi(\mathbf{h}_C, \mathbf{h}_\Delta)$ for skip-generation




Figure 4: Kernel density estimation(KDE) of similarity scores for different lineage relations. Each subplot shows the similarity distribution between a child model and its parent, grandparent, great-grandparent and non-lineage items, respectively. True parent-child pairs exhibit distinctly higher similarity, enabling effective rejection of false lineage claims.

relationships is expected to gradually decrease as the generational distance increases.

How to determine the Decision Threshold T ? To explore an appropriate decision threshold T of lineage closeness, we construct multiple four-generation model families based on the ResNet-18 [45], MobileNet-V2 [37], Diffusion-V2 [36], and LLaMA-3 [13] architectures. Fig. 4 shows the kernel density estimation (KDE) of knowledge alignment similarity scores across different lineage relations of ResNet-18 families, and the other results can be found in Fig. 10 of Appendix B.1. Each subplot originating from parent models trained on Caltech101 [20], Dogs [17], and Flowers [29] dataset, respectively.

In all cases, the parent-child pairs (blue) consistently exhibit the highest similarity scores, with peak values centered 0.83-0.88, indicating strong knowledge consistency between adjacent generations. The grandparent relationships follow, peaking between 0.54-0.57, while the great-grandparent class peaking in the 0.39-0.45 range. In contrast, non-lineage pairs display the lowest similarity, concentrated near 0.00-0.20. This decay across generational distances demonstrates a clear and stable correlation between the embedding vector alignment and the true knowledge evolution, which is consistent with our previous analysis. Additionally, these results demonstrate that our MLA is capable of differentiating at least four generations of models.

Interestingly, the similarity intervals for lineage-nonlinear are well separated without significant overlap. This allows us to define tiered thresholds for lineage inference, enabling reliable differentiation between direct lineage, distant lineage, and non-lineage model pairs. For example, in model families derived from ResNet-18, a lower threshold (e.g., 0.3) allows reliable separation of three-generation lineage from non-lineage models, while a higher threshold (e.g., 0.7) precisely captures direct parent-child relationship. We further note that the exact similarity values slightly vary across different model architectures. In the MobileNet-V2 model family, non-lineage pairs typically have similarity scores below 0.1, while grandparent relations exceed 0.2; in contrast, the corresponding values for the Diffusion-V2 model family are approximately 0.2 and 0.4, respectively. Integrating the experimental findings

from different families reveals that non-lineage similarities are mostly below 0.2, whereas parent-child similarities are generally above 0.7. Thus, a universal T in this paper is used across different model families for coarse-grained attestation—potentially reusing the value from a known family with similar architecture. For a higher accuracy, a family-specific T may be adopted by training in advance with labeled lineage pairs.

6 Experiments

6.1 Experimental Settings

Basic Setup and Implementation Details. We utilized the Adam [18] optimizer with a learning rate of 1×10^{-4} for training the proposed framework, and all experiments were conducted on two NVIDIA A100 GPUs. For the classification-oriented knowledge encoder Ψ , the MLP block \mathcal{M} consists of a fully connected layer followed by a LayerNorm operation, which aims to project the input feature embeddings to a uniform 128-dimensional vector and ensure stable training. Next, the transformer encoder \mathcal{T} with 2 layers and a hidden size of 128 was applied. To handle varying numbers of categories across models, the input sequences were padded to the maximum number of categories in the candidate model using zero vectors. Additionally, the knowledge fusion network Φ was a lightweight projection block consisting of a single linear layer with dimensions (256×128) followed by a ReLU activation.

For the diffusion-oriented knowledge encoder Ψ , the input consists of a 256×256 image and its corresponding textual prompt. The feature extracted from the UpBlock layer has dimensions of $(320, 32, 32)$. The convolution layer \mathcal{C} with a kernel size of 3×3 was then applied to this feature, followed by a global pooling operation \mathcal{G} that transformed it into a 320-dimensional vector. Finally, the MLP \mathcal{M} was used to map this vector into a 160-dimensional knowledge embedding.

For the LLM-oriented knowledge encoder Ψ , input features were first projected to a 512-dimensional vector space through a fully connected layer, followed by a LayerNorm layer. A second position-aware fully connected layer was then applied to capture local contextual patterns, followed by a dropout

layer with a rate of 0.1. Finally, an adaptive average pooling layer aggregated the sequence into a 512-dimensional knowledge vector. Accordingly, the fully connected layer in \mathcal{M} of the knowledge fusion network Φ was configured as 1024×512 dimensions.

Classification Model Library. For the ResNet-18 [45] model families, we developed more than 500 models in total. Specifically, we trained 20 parent models on each of the Flowers [29], Dogs [17], and Caltech-101 [20] datasets. To promote model diversity, we randomly partitioned classes for parent and child models and trained them with varied hyperparameters. The resulting (parent, child) model pairs were then divided, with 80% used for training and 20% for testing. For across-generational relationship, we took the above mentioned parent models as Generation-1 and the child models as Generation-2. We then randomly fine-tuned them down two generations along Caltech101 [20]→ Caltech101 [20]→ CIFAR100 [42]→ Aircraft [28], Dogs [17]→ Dogs [17]→ CIFAR100 [42]→ Oxford-IIIT Pet [39], and Flowers [29]→ Flowers [29]→ CIFAR100 [42]→ Caltech101 [20]. The fourth-generation models were primarily included to measure lineage closeness across generations. We also constructed a model family based on MobileNet-V2 architecture [37], as detailed described in Appendix A.

Diffusion Model Library. We selected Stable-Diffusion-v2 version as the model family structure. From the Hugging Face [48] repository, we collected 30 distinct versions of *Stable-Diffusion-v2* as the set of parent models. Subsequently, we apply two mainstream fine-tuning techniques—DreamBooth [36] and LoRA [15]—to generate 100 fine-tuned variants for each method, which serve as the corresponding child models. To obtain knowledge feature representations for lineage analysis, we randomly sampled 1000 images from COCO2014 [25] dataset as the probe sample set. Further details regarding the implementation of DreamBooth and LoRA are provided in the Appendix A.

LLM Library. We collected 5 derivative models and 22 additional fine-tuned variants of *meta-llama/Llama-3.1-8B*, as well as 2 derivative models and 49 fine-tuned variants of *Qwen/Qwen2.5-1.5B* from Hugging Face¹. For LLaMA-families, we specify the parameters of *meta-llama/Llama-3.1-8B* as the initialization parameters θ_0 and considered the original first-generation fine-tuned models, such as *TwinLlama-3.1-8B*, as parent models, while their fine-tuned descendants (such as *TwinLlama-3.1-8B-DPO*) served as child models. This setup was adopted because the initialization parameters of *meta-llama/Llama-3.1-8B*—which serve as the original parameters for training the parent model—were not available. Based on descriptions in Hugging Face, we constructed probing datasets by sampling 50 prompts from each of six datasets and collecting five responses per prompt to capture model knowledge. Further details regarding the implementation of LLM libraries are provided in the Appendix A.

Baselines. We evaluate our proposed MLA framework

against four typical baselines: the state-of-the-art methods Neural Lineage [51] (Abbreviated as N.L.), Neural Phylogeny [52] (Abbreviated as N.P.), KRM [46], and IPGuard [4]. Since the number of classes in parent–child pairs may differ, and both KRM and IPGuard rely on class-specific boundary samples, we use a small number of classes in each pair for consistency. The N.L. method focuses on the coupling between model weights and features extracted from the fine-tuning dataset. As such access is infeasible under our attack setting, we instead use the parent model’s training data to extract the aligned features for the N.L. method, which preserves the core feature–weight relationship needed for evaluation.

It should be noted that KRM and IPGuard are designed specifically for classification models, and thus they are not applicable in our MLA experiments for diffusion model and LLM, where we compare only with the N.L. and N.P. methods. *Furthermore, the N.L. method requires storing a large number of feature representations and model parameters, making it infeasible to run on the stable diffusion model and LLM experiments.*

Evaluation Metrics. We adopt the following metrics to evaluate the performance of the lineage attestation methods, i.e., the ability to fairly determine the presence or absence of a lineage relationship between attack models and victim models under various threat scenarios. (1) True Positive Rate (TPR): The proportion of positive samples (i.e., model pairs with a true lineage relationship) that are correctly identified. (2) False Positive Rate (FPR): The proportion of negative samples (i.e., model pairs without a lineage relationship) that are incorrectly classified as positive. (3) Receiver Operating Characteristic (ROC) curve: It plots the TPR against the FPR at various decision thresholds.

In the *Illegal Model Derivative Verification* scenario, the lineage attestation mechanism seeks to determine whether a suspect model f_C^A is directly derived from a protected model f_P . The task is to distinguish true lineage pairs from unrelated ones, where a high TPR indicates effective detection and a low FPR reflects strong robustness against false accusations.

In the *False Lineage Claim Reject* scenario, the verifier must distinguish true direct lineage from indirect or non-lineage relations to invalidate false claims. The goal of this task is to discriminate positive pairs (parent, child) from negative pairs (non-direct parent, child), with the ROC curves used to evaluate the identification performance.

6.2 MLA Performance for Illegal Model Derivative Verification

We conduct experiments under two attack scenarios: the Across-Generational Attack (AGA) and Weight Pruning before Fine-tuning Attack (WPA). For AGA, we evaluate MLA method’s capability to detect lineage relationships within three generations. For WPA, we adopt a pruning rate p of 10% following the KRM [46] setup.

Table 2: The performance of model lineage attestation (MLA) against across-generational attack (AGA) and weight pruning attack (WPA). N.L. refers to neural lineage [51], N.P. refers to neural phylogeny [52] and T.I. refers to TinyImageNet dataset.

ResNet-18 Model Family												
Methods	Caltech101->Caltech101->CIFAR100				Dogs->Dogs->CIFAR100				Flowers->Flowers->CIFAR100			
	TPR		FPR		TPR		FPR		TPR		FPR	
	AGA	WPA	AGA	WPA	AGA	WPA	AGA	WPA	AGA	WPA	AGA	WPA
N.L. [51]	0.99	0.85	0.02	0.08	0.94	0.83	0.04	0.06	0.95	0.76	0.05	0.12
N.P. [52]	0.58	0.43	0.21	0.33	0.42	0.27	0.07	0.17	0.60	0.40	0.12	0.15
KRM [46]	0.19	0.15	0.16	0.11	0.15	0.14	0.13	0.13	0.20	0.19	0.10	0.13
IPGuard [4]	0.19	0.16	0.15	0.14	0.16	0.14	0.14	0.12	0.20	0.13	0.11	0.12
Ours	1.00	0.99	0.00	0.01	0.99	0.99	0.02	0.03	0.99	0.99	0.01	0.00

MobileNet-V2 Model Family												
Methods	Caltech101->T.I.->Mixed Dataset				T.I.->Pet->Mixed Dataset				Flowers->CIFAR100->Mixed Dataset			
	TPR		FPR		TPR		FPR		TPR		FPR	
	AGA	WPA	AGA	WPA	AGA	WPA	AGA	WPA	AGA	WPA	AGA	WPA
N.L. [51]	1.00	0.83	0.02	0.08	0.95	0.92	0.13	0.16	0.96	0.94	0.10	0.13
N.P. [52]	0.73	0.66	0.11	0.08	0.65	0.40	0.15	0.20	0.63	0.45	0.17	0.15
KRM [46]	0.18	0.15	0.18	0.11	0.22	0.14	0.18	0.13	0.19	0.18	0.18	0.13
IPGuard [4]	0.18	0.16	0.13	0.14	0.14	0.14	0.12	0.20	0.13	0.11	0.12	0.11
Ours	1.00	0.99	0.00	0.03	0.99	0.99	0.01	0.03	0.99	0.98	0.01	0.01

MLA across Classification Model Families. In our approach, the decision threshold T for lineage detection is set to 0.3. Table 2 shows the lineage attestation results of both ResNet-18 model family and MobileNet-V2 model family.

First of all, it can be observed that our method consistently achieves near-perfect performance against AGA across both ResNet-18 [45] and MobileNet-V2 [37] architectures, with TPR around 0.99 and FPR not exceeding 0.02. This superior performance stems from our design, which captures the consistent trajectory of both model parameters alternation and knowledge evolution during fine-tuning, enabling robust adaptation to diverse fine-tuning strategies. In contrast, the N.L. [51] method achieved suboptimal performance. It attempts to encode the evolution of a model by examining the coupling between its parameters and the extracted data features. However, this method fails to precisely characterize the nuanced knowledge transfer process that occurs during fine-tuning. The N.P. [52] method suffers even more severe degradation, with TPR dropping to 0.42 and FPR rising to 0.21 under the worst-case, due to its reliance on parameter-level differences while overlooking the fine-tuning-induced semantic variations among models. Comparably, both KRM [46] and IPGuard [4] perform poorly, with TPRs ranging from 0.10 to 0.20 and relatively high FPRs. These methods assume stable decision boundaries during fine-tuning, that may hold under lightweight fine-tuning adaptations but fail under substantial task shifts, where models significantly reshape their boundaries to fit new fine-tuning task.

Furthermore, we observe that our method remains effective under WPA scenario, as the pruned models still retain sufficient knowledge to support reliable attestation. In comparison, N.L. [51], which directly encodes parameter changes, shows a noticeable drop in performance under pruning. This

is particularly obvious in the Caltech101-based model family, where it suffers a TPR drop of 0.14 in ResNet-18 model family and 0.17 in MobileNet-V2 model family, highlighting its sensitivity to weight perturbations introduced by pruning. Similarly, these perturbations also impact N.P. [52], reducing its detection capability with TPR dropping as low as 0.27 in the worst case. WPA has limited impact on KRM [46] and IPGuard [4], since it does not significantly alter the model’s decision boundaries—on which these methods heavily rely. However, despite their stability, their overall TPR remains low (typically between 0.10 and 0.20), indicating limited capability in accurately detecting true lineage.

MLA across Diffusion Model Families. For diffusion models, *DreamBooth* and *LoRA* are two widely used fine-tuning techniques, and we evaluate whether our method and N.P. can distinguish models fine-tuned using these techniques from unrelated models. As shown in Table 3, our method shows strong robustness to AGA and WPA across both DreamBooth and LoRA approaches. We attribute this robustness to the fact that these fine-tune approaches enable diffusion models to adapt to specific tasks by embedding new knowledge, while our MLA method effectively captures this knowledge-level alignment. In contrast, the baseline method N.P. relies on parameter-level clustering and lacks a mechanism to capture knowledge evolution. As a result, it exhibits a marginal drop in TPR against AGA (0.09 lower than ours in LoRA), and suffers a more noticeable decline under WPA than under AGA, with a 0.04 drop on both DreamBooth and LoRA. The attestation signal reflecting model lineage becomes increasingly obscured under WPA. Additionally, N.P. consistently shows high FPR (0.20–0.23), likely due to its clustering-based strategy, which fixes cluster boundaries based on training models. This may lead to test models being misclassified into incor-

rect clusters, especially when different parent models exhibit similar parameter patterns.

Table 3: MLA performance on Stable Diffusion models.

	Dreambooth (TPR)		LoRA (TPR)		FPR	
	AGA	WPA	AGA	WPA	AGA	WPA
N.P. [52]	0.94	0.90	0.91	0.87	0.23	0.20
Ours	1.00	0.99	1.00	0.99	0.02	0.02

MLA across LLMs Families. For the LLM library, the models are realistic ones directly collected from the open-source repository, not trained by us. Thus we can’t perform the WPA experiments and have only the results for AGA. As shown in Table 4, our method effectively identifies model lineage relationships, outperforming the baseline (N.P.). Despite potential semantic discrepancies between the probe datasets and the parent models, the experimental results show that our method remains robust. This improvement can be attributed to its ability to capture fine-grained knowledge transfer patterns during fine-tuning, whereas the N.P. method, relying on heuristic and less adaptive features, is more susceptible to dynamically semantic inconsistencies among models.

Table 4: MLA performance on LLMs.

Model	Ours		N.P. [52]	
	TPR	FPR	TPR	FPR
Llama3.1-8B-family	0.97	0.02	0.80	0.33
Qwen2.5-1.5B-family	0.99	0.01	0.79	0.21

6.3 MLA Performance for False Lineage Claim Rejection

Fig. 5 shows the ROC curves comparing our method with two binary-based decision baselines, N.P. [52] and N.L. [51], in the lineage verification setting. Unlike our approach, which provides a continuous confidence score for flexible thresholding, the baselines yield only binary decisions, thus appearing as discrete TPR/FPR points. As a result, their performance is represented as the discrete TPR/FPR points.

As shown in Fig. 5, our method achieves an AUC (area under the curve) of 0.999 on the ResNet-18 family and 0.952 on the MobileNet-V2 family, demonstrating strong discriminative ability across a wide range of thresholds. N.L. achieves a lower FPR (0.10) and a comparable TPR (0.87) on ResNet-18, while suffers on MobileNet-V2 with a TPR of 0.80 and a much higher FPR of 0.18. This degradation may stem from the fact that its decision relies on behavioral similarity in weight or output features, which can also be shared by non-direct lineage models within the same family. N.P. achieves a TPR of 0.65 but incurs a high FPR of 0.22, indicating a

strong tendency to misidentify unrelated models as related. This is likely due to its data-free design, which relies solely on parameter-level similarity. As models from the same family often converge along similar optimization trajectories, N.P. fails to effectively distinguish direct lineage from architectural resemblance.

We further evaluate the generalizability of our method on generative and large language model families, as shown in Figs. 5(c) and (d). Our method achieves an AUC of 0.990 on Stable Diffusion-V2 and 0.977 on the LLaMA3 family, indicating strong and consistent performance across diverse model architectures and modalities. The superiority lies in its knowledge-centric design, especially in large-scale models where richer knowledge enhances signal separability. In contrast, existing baselines such as N.P. and N.L. exhibit a marked inability to differentiate these relationships. Relying solely on feature-level or parameter-level similarity is insufficient for effective alignment of intrinsic knowledge consistency. The inconsistent performance of N.P. across the four model architectures suggests that it is sensitive to both semantic and structural variations, reflecting a lack of robustness.

6.4 Resistance to Adaptive Attacks

In practice, adversaries may become aware of our MLA protocol and develop adaptive attacks to evade verification. They may attempt to disrupt knowledge-consistency verification by altering the model architecture, or to undermine knowledge-vector extraction by perturbing model parameters.

Structure Evasion Attack(via distillation). In this attack, the adversary first steals the protected model f_P and fine-tunes it to obtain a descendant model f_C^A . To evade lineage verification, the adversary then applies architecture transformation strategies such as knowledge distillation, transferring f_C^A into a structurally different student model $f_C^{A'}$. By design, $f_C^{A'}$ compromises parameter or structure-based consistency checks, rendering them ineffective.

Model lineage depicts the fine-tuning relationship between two models, emphasizing their knowledge evolutionary path in fine-tuning, while model distillation (born-again neural networks) resembles knowledge cloning, lacking both parameter inheritance and continuity in the knowledge evolution, and thus it can corrupt model lineage. However, the behavioral traits of the teacher model can still be inherited by the student model, leading to an atypical pattern of knowledge consistency between the two models.

To assess the knowledge consistency between the suspect model $f_C^{A'}$ and the victim model f_P , we perform reverse distillation: starting from the initial parameters of f_P , we distill $f_C^{A'}$ into a proxy model $f_C^{A''}$ that has the same architecture with f_P and then evaluate its lineage similarity with f_P . Thus, the lineage relationship between $f_C^{A'}$ and f_P can be measured by the knowledge consistency between $f_C^{A''}$ and f_P . For evaluating the robustness of our MLA to the structure evasion

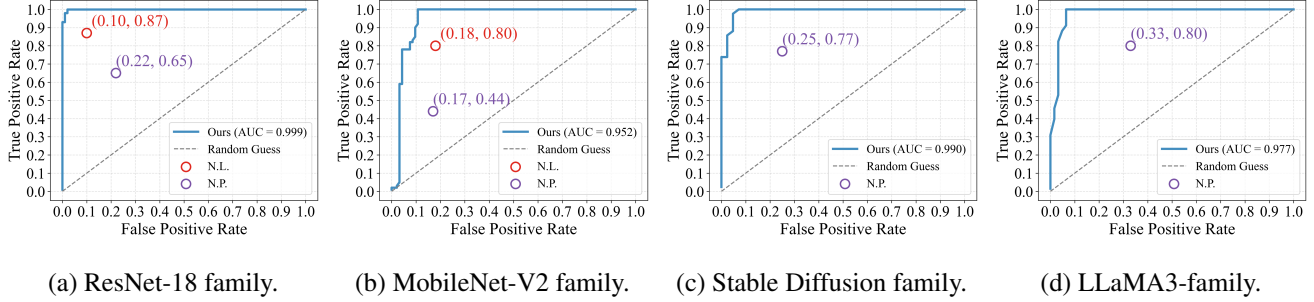


Figure 5: ROC of false lineage claim rejection for various types of model families.

attack (via distillation), we simulate it by distilling the classification model $f_C^{\mathcal{A}}$ into the compact network $f_C^{\mathcal{A}'}$ composed of three convolutional layers in a data-free manner using deep inversion [50]. We then perform reverse distillation from $f_C^{\mathcal{A}'}$ towards f_P to obtain $f_C^{\mathcal{A}''}$ and measure the resulting lineage similarity to verify knowledge inheritance.

Table 5 reports the resistance of our MLA against distillation attacks under varying student model qualities. When the student model achieves a high test accuracy (65.43%), corresponding to a well-distilled model, MLA attains a TPR of 0.76, indicating almost perfect detection of lineage relationships. As the distillation quality decreases (test accuracy drops to 48.62%), the overlap of inherited knowledge diminishes, leading to a lower TPR of 0.36. Usually, the goal of a distillation attack is to preserve a high-quality (i.e. a higher test accuracy more than 65.43%) student model. That said, the attacker is still difficult to bypass our MLA’s verification unless the student model is with very low accuracy.

Table 5: Resistance of MLA to Distillation Attacks. Acc. refers to the classification accuracy.

Acc.(%)	89.90(f_P)	65.43($f_C^{\mathcal{A}'}$)	50.32($f_C^{\mathcal{A}''}$)	48.62($f_C^{\mathcal{A}''}$)
TPR	0.99	0.76	0.54	0.36





Figure 6: Lineage attestation accuracy across different model families against varying parameter perturbation ratio ρ .

Parameter Perturbation Attack. We simulate an adaptive attack that perturbs model parameters to hinder obtaining correct evolution-model and knowledge-vector. Following NPET [47], we add zero-mean uniform noise scaled to each parameter’s mean absolute value. The perturbation ratios $\rho \in \{0.05, 0.1, 0.15\}$ specify the fraction of mean absolute


values used to control the resulting standard deviation of the noise. As shown in Fig. 6, these parameter perturbations do not noticeably affect knowledge alignment, despite slightly altering the model weights without impacting inference. This indicates that MLA is resilient to random parameters noise, and such minor perturbations are insufficient to compromise lineage verification.

Stronger Weight-Pruning Attacks. Weight-pruning serves as an adaptive attack, as removing neurons may disrupt the knowledge patterns used for lineage assessment. We evaluate the resilience of MLA to larger weight-pruning rates. Since weight-pruning attack inevitably degrades model performance that reflects the model’s practical usability, we assess the attack-induced model performance degradation using task-appropriate metrics. Specifically, for LLMs, we measure the relative increase ($\Delta \text{PPL}(\%)$) in Perplexity (PPL) [3] compared to the original model; for diffusion models, we evaluate the relative change ($\Delta \text{FID}(\%)$) in Fréchet Inception Distance (FID) [14]; and for classification models, we report the percentage of accuracy loss ($\Delta \text{Acc}(\%)$). PPL measures a language model’s uncertainty in predicting the next token given the preceding context and an increase in PPL directly reflects a degradation in the model’s ability to generate fluent and coherent text. FID, on the other hand, quantifies the discrepancy between the distributions of generated images and real images in the feature space, and an increase in FID indicates that the generative model approximates the target image distribution poorly.


For classification models, child models are constructed via fine-tuning without introducing new classes, such that the number of target classes remains unchanged. As shown in Fig. 7(a), the TPR (averaged over the ResNet-18 and MobileNet-V2 families) decreases as the pruning ratio increases, accompanied by a corresponding drop in classification accuracy. This behavior can be attributed to the fact that, when the class set is unchanged, the classification knowledge encoded in the decision boundaries largely preserves its fundamental structure even after pruning. As a result, classification models exhibit relatively higher robustness under increasing pruning strength. In contrast, as shown in Fig. 7(b) and Fig. 7(c), LLMs and diffusion models demonstrate signifi-



(a) WPA for classification models.



(b) WPA for diffusion models.



(c) WPA for LLMs.

Figure 7: MLA performance against Weight Pruning Attack.

cantly higher vulnerability under stronger pruning attacks. As the pruning ratio increases, features that encode the model’s sensitivity to training data are progressively removed, leading to a substantial degradation in lineage attestation accuracy. It is also worth noting that, although classification models maintain relatively high lineage identification performance under moderate pruning, their classification accuracy drops sharply once the pruning ratio exceeds 50%, which severely limits their practical usability. A similar trend is observed for diffusion models and LLMs, where aggressive pruning (e.g., 50%) results in substantial degradation of both generation quality and model usability.

Knowledge-Overwriting Attack. In the Illegal Model Derivative Verification setting, an adversary aware of the probe dataset may attempt to corrupt the extracted knowledge by injecting behaviors that conflict with probe-induced signals. We assume that the adversary possesses prior knowledge of the probe dataset (e.g., images for classification models, generative images and captions for diffusion models, and probe prompts for LLMs). We vary the attack strength by perturbing different portions of the probe set for the following mechanisms: (i) *Label-Perturbed Fine-tuning for Classification Model*. The adversary re-labels a subset of samples within each selected class, inducing shifts in the decision regions tied to probe data. (ii) *Caption-Corrupted Fine-tuning for Diffusion Models*. The adversary substitutes probe captions with incorrect descriptions and fine-tunes adapters on the resulting mismatched image–caption pairs, disrupting probe-induced semantic alignment. (iii) *Cross-Family Response Fine-tuning for LLMs*. The adversary replaces the correct probe outputs with responses from an unrelated model (e.g., Llama-3.1-8B), using them as pseudo-targets during fine-tuning to inject cross-family behavioral patterns. As shown in Fig. 8, the TPR values of classification models degrade only slightly because boundary perturbations do not alter their underlying decision structure. For diffusion models, caption corruption is localized to the mismatched image–caption pairs, leaving semantic knowledge within unrelated probe largely unaffected. In contrast, LLMs are more vulnerable since the cross-family pseudo-targets shift their conditional token distribution. We observe that the knowledge overwriting attack can forcibly

inject conflicting behavioral signals into the suspect model. However, such effects can be straightforwardly mitigated by evaluating the model on an independently constructed probe dataset.





Figure 8: MLA performance against Knowledge-Overwriting Attack.

6.5 Ablation Study

In this section, we present ablation studies to examine the influence of critical factors in our method, including the effect of probe data size and the generalization ability for different LLM families. Additional results are provided in Appendix B.2.



(a) Varying data size.

(b) Varying data source.

Figure 9: MLA performance against varying probe dataset.

Effect of Varying Probe Dataset. To evaluate the sensitivity of MLA to probe quality, we evaluate its robustness with respect to both varying probe size and varying probe data

source. The probe data are constructed either by re-sampling an equal number of samples from the same dataset, or by drawing samples from unseen datasets during training (detailed in Section B.2). As shown in Fig. 9, classification models are only mildly affected by different data source. Furthermore, performance on diffusion models and LLMs remains largely stable under both different probe size and varying data source. Even when the probe is drawn from an unseen dataset, diffusion models and LLMs still exhibit non-trivial semantic alignment on these data, owing to their strong priors and extensive pretraining on large-scale datasets.

LLM Families Generalization. To further evaluate the generalizability of MLA beyond the model families used in the main experiments, we test its performance on 20 additional HuggingFace variants spanning four major ecosystems: Microsoft Phi-4, Alibaba Qwen-3, Mistral, and Google Gemma. The complete list of evaluated model variants is provided in our OpenScience repository.¹ As shown in Table 6, MLA achieves consistently high TPR (0.98–0.99) and low FPR (0.01–0.03) across all four families. These results demonstrate that the proposed lineage attestation framework generalizes robustly across heterogeneous architectures and training pipelines, indicating its applicability to a broad range of open-weight LLM ecosystems.

Table 6: MLA performance across additional LLM families.

	Phi-4	Qwen-3	Mistral	Gemma
TPR	0.99	0.98	0.99	0.98
FPR	0.01	0.01	0.03	0.01

7 Conclusion

In this work, we unravel the knowledge mechanism of fine-tuning-induced model evolution. Based on this rationale, we address the model lineage attestation problem by exploring the knowledge evolution in neural networks and propose a model lineage attestation framework. In the framework, knowledge modifications by fine-tuning are explicitly captured through model editing, and further a knowledge vectorization technique is proposed to embed these edited changes into a common latent space. This enables robust alignment in terms of the consistency among the resulting knowledge embeddings, providing the fundamental evidence for confirming model lineage relationship. Extensive experimental evaluations under two key adversarial scenarios validate the effectiveness of our method. Our framework offers a robust solution for verifying the lineage of model derivatives, especially in the context of open-weight model reuse, where unauthorized derivations and false lineage claims are more prevalent.

¹<https://zenodo.org/records/17958570>.

Acknowledgments

We thank the anonymous reviewers and our shepherd for their constructive comments. This work was supported by the Strategic Priority Research Program of the Chinese Academy of Sciences (NO.XDB0690302), NSFC under grant No.62371450 and Project of Key Laboratory of Cyberspace Security Defense under grant No.2025-A03.

Ethical Considerations

In this work, we study the problem of model lineage attestation using both publicly released models and those we trained on open-source datasets. All experiments are carried out entirely within local premises on controlled servers, ensuring that no external systems or third-party services were involved. Nevertheless, we emphasize that any future application of lineage attestation techniques to proprietary or commercial models should be conducted only with explicit authorization and in accordance with relevant privacy policies and regulatory standards.

Lineage attestation introduces considerations for several stakeholders: model developers and model owners may be affected when derivation relationships are inferred, as such findings can influence intellectual-property claims or assertions of model provenance; organizations deploying ML systems may rely on lineage conclusions to support auditing or compliance decisions, and thus have an interest in ensuring that such analyses are conducted appropriately; end-users and the wider public are indirectly connected, since the trustworthiness of deployed AI systems depends in part on the authenticity of their underlying models.

As with any analysis technique that examines model parameters, it is plausible that lineage attestation could be misused to infer relationships between models without authorization. By determining whether one model is derived from another, an actor could potentially employ such information to support unauthorized provenance claims or to dispute ownership inappropriately. However, these risks are limited in practice: lineage attestation requires explicit access to both models and does not reveal training data, internal parameters, or other sensitive content. Under white-box access conditions, the risks are further mitigated because the evaluation environment and model access paths are strictly controlled. Within these constraints, the defensive value of lineage attestation outweighs its potential risks, as it enables the reliable identification of illicit derivatives and prevents false lineage assertions in ML supply chains. The process involves comparing model parameters or derived representations, which underscores the importance of performing such evaluations only with appropriate authorization, particularly when proprietary systems are involved.

Open Science

Implementations of our method as well as the model zoo derived from our study, will be made publicly available. Code available at: <https://zenodo.org/records/17958570>.

References

- [1] Yossi Adi, Carsten Baum, Moustapha Cisse, Benny Pinkas, and Joseph Keshet. Turning your weakness into a strength: watermarking deep neural networks by backdooring. In *Proceedings of the 27th USENIX Conference on Security Symposium*, page 1615–1631, 2018.
- [2] Hayet Benbrahim and Ali Behloul. Fine-tuned xception for image classification on tiny imagenet. In *2021 International Conference on Artificial Intelligence for Cyber Security Systems and Privacy (AI-CSP)*, pages 1–4, 2021.
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [4] Xiaoyu Cao, Jinyuan Jia, and Neil Zhenqiang Gong. Ipguard: Protecting intellectual property of deep neural networks via fingerprinting the classification boundary. In *ASIA CCS '21*, pages 14–25. ACM, 2021.
- [5] Eole Cervenka. Naruto blip captions. <https://huggingface.co/datasets/lambdalabs/naruto-blip-captions/>, 2022.
- [6] Pengzhen Chen, Yanwei Liu, Xiaoyan Gu, Enci Liu, Zhuoyi Shang, Xiangyang Ji, and Wu Liu. Plugmark: A plug-in zero-watermarking framework for diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 17335–17345, 2025.
- [7] Zhuo Chen, Xinyu Wang, Yong Jiang, Zhen Zhang, Xinyu Geng, Pengjun Xie, Fei Huang, and Kewei Tu. Detecting knowledge boundary of vision large language models by sampling-based inference. *CoRR*, abs/2502.18023, 2025.
- [8] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3606–3613, 2014.
- [9] Tianshuo Cong, Xinlei He, and Yang Zhang. Sslguard: A watermarking scheme for self-supervised learning pre-trained encoders. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, CCS '22*, page 579–593, 2022.
- [10] Yingqian Cui, Jie Ren, Yuping Lin, Han Xu, Pengfei He, Yue Xing, Lingjuan Lyu, Wenqi Fan, Hui Liu, and Jiliang Tang. Ft-shield: A watermark against unauthorized fine-tuning in text-to-image diffusion models. *SIGKDD Explor. Newsl.*, 26(2):76–88, 2025.
- [11] Aditya Desu, Xuanli He, Qiongkai Xu, and Wei Lu. Generative models are self-watermarked: Declaring model authentication through re-generation. *Transactions on Machine Learning Research (TMLR)*, 2024.
- [12] Govind Krishnan Gangadhar and Karl Stratos. Model editing by standard fine-tuning. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 5907–5913, Bangkok, Thailand, August 2024.
- [13] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv e-prints:2407.21783*, 2024.
- [14] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- [15] Edward Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *The Tenth International Conference on Learning Representations*, 2022.
- [16] Gabriel Ilharco, Marco Túlio Ribeiro, Mitchell Wortsman, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. In *The Eleventh International Conference on Learning Representations (ICLR)*, 2023.
- [17] Aditya Khosla, Nityananda Jayadevaprakash, Bangpeng Yao, and Fei-Fei Li. Novel dataset for fine-grained image categorization: Stanford dogs. In *Proc. CVPR workshop on fine-grained visual categorization (FGVC)*, volume 2, 2011.
- [18] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations (ICLR)*, 2015.
- [19] Isabell Lederer, Rudolf Mayer, and Andreas Rauber. Identifying appropriate intellectual property protection mechanisms for machine learning models: A systematization of watermarking, fingerprinting, model access,

- and attacks. *IEEE Transactions on Neural Networks and Learning Systems*, 35(10):13082–13100, 2024.
- [20] Fei-Fei Li, Fergus Rob, and Perona Pietro. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *2004 conference on computer vision and pattern recognition workshop*, pages 178–178. IEEE, 2004.
- [21] Hongkang Li, Yihua Zhang, Shuai Zhang, Pin-Yu Chen, Sijia Liu, and Meng Wang. When is task vector provably effective for model editing? A generalization analysis of nonlinear transformers. In *The Thirteenth International Conference on Learning Representations (ICLR)*, 2025.
- [22] Linyang Li, Pengyu Wang, Ke Ren, Tianxiang Sun, and Xipeng Qiu. Origin tracing and detecting of llms. *arXiv preprint arXiv:2304.14072*, 2023.
- [23] Zongjie Li, Chaozheng Wang, Shuai Wang, and Cuiyun Gao. Protecting intellectual property of large language model-based code generation apis via watermarks. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security, CCS ’23*, page 2336–2350, 2023.
- [24] Ruofan Liang, Tianlin Li, Longfei Li, Jing Wang, and Quanshi Zhang. Knowledge consistency between neural networks and beyond. In *8th International Conference on Learning Representations (ICLR)*, 2020.
- [25] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755, 2014.
- [26] Jian Liu, Rui Zhang, Sebastian Szyller, Kui Ren, and N. Asokan. False claims against model ownership resolution. In *Proceedings of the 33rd USENIX Conference on Security Symposium, SEC ’24*, 2024.
- [27] Sofiane Lounici, Mohamed Njeh, Orhan Ermis, Melek Önen, and Slim Trabelsi. Yes we can: Watermarking machine learning models beyond classification. In *2021 IEEE 34th Computer Security Foundations Symposium (CSF)*, pages 1–14, 2021.
- [28] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013.
- [29] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian conference on computer vision, graphics & image processing*, pages 722–729. IEEE, 2008.
- [30] Charles O’Neill, Jack Miller, Ioana Ciuca, Yuan-Sen Ting, and Thang Bui. Adversarial fine-tuning of language models: An iterative optimisation approach for the generation and detection of problematic content. *arXiv e-prints:2308.13768*, 2023.
- [31] Ding Sheng Ong, Chee Seng Chan, Kam Woh Ng, Lixin Fan, and Qiang Yang. Protecting intellectual property of generative adversarial networks from ambiguity attacks. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3629–3638, 2021.
- [32] Guillermo Ortiz-Jiménez, Alessandro Favero, and Pascal Frossard. Task arithmetic in the tangent space: Improved editing of pre-trained models. In *Annual Conference on Neural Information Processing Systems, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- [33] Dario Pasquini, Evgenios M. Kornaropoulos, and Giuseppe Ateniese. Llmmap: Fingerprinting for large language models. *CoRR*, abs/2407.15847, 2024.
- [34] Zirui Peng, Shaofeng Li, Guoxing Chen, Cheng Zhang, Haojin Zhu, and Minhui Xue. Fingerprinting deep neural networks globally via universal adversarial perturbations. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13420–13429, 2022.
- [35] Rahul Ramesh and Pratik Chaudhari. Model zoo: A growing brain that learns continually. In *The Tenth International Conference on Learning Representations (ICLR)*, 2022.
- [36] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22500–22510, 2023.
- [37] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4510–4520, 2018.
- [38] Konstantin Schürholt, Diyar Taskiran, Boris Knyazev, Xavier Giró-i Nieto, and Damian Borth. Model zoos: a dataset of diverse populations of neural network models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, 2022.
- [39] Zhuoyi Shang, Yanwei Liu, Jinxia Liu, Xiaoyan Gu, Ying Ding, and Xiangyang Ji. Know2vec: A black-box proxy for neural network retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 2025.

- [40] Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. Hugginggpt: Solving ai tasks with chatgpt and its friends in hugging face. In *Advances in Neural Information Processing Systems*, volume 36, pages 38154–38180, 2023.
- [41] Chenyang Si, Ziqi Huang, Yuming Jiang, and Ziwei Liu. Freeu: Free lunch in diffusion u-net. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4733–4743, 2024.
- [42] Sahil Singla, Surbhi Singla, and Soheil Feizi. Improved deterministic l2 robustness on CIFAR-10 and CIFAR-100. In *The Tenth International Conference on Learning Representations*, 2022.
- [43] Cecilia Summers and Michael J. Dinneen. Nondeterminism and instability in neural network optimization. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, volume 139, pages 9913–9922, 2021.
- [44] Mina Taraghi, Gianolli Dorcelus, Armstrong Foundjem, Florian Tambon, and Foutse Khomh. Deep learning model reuse in the huggingface community: Challenges, benefit and trends. In *2024 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*, pages 512–523, 2024.
- [45] Sasha Targ, Diogo Almeida, and Kevin Lyman. Resnet in resnet: Generalizing residual architectures. *arXiv preprint arXiv:1603.08029*, 2016.
- [46] Zehao Tian, Zixiong Wang, Ahmed M. Abdelmoniem, Gaoyang Liu, and Chen Wang. Knowledge representation of training data with adversarial examples supporting decision boundary. *IEEE Trans. Inf. Forensics Secur.*, 18:4116–4127, 2023.
- [47] Dongwei Wang and Huanrui Yang. Taming sensitive weights: Noise perturbation fine-tuning for robust llm quantization. *arXiv preprint arXiv:2412.06858*, 2024.
- [48] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, and et al. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, October 2020.
- [49] Kang Yang, Run Wang, and Lina Wang. Metafinger: Fingerprinting the deep neural networks with meta-training. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 776–782, 7 2022.
- [50] Hongxu Yin, Pavlo Molchanov, Jose M Alvarez, Zhizhong Li, Arun Mallya, Derek Hoiem, Niraj K Jha, and Jan Kautz. Dreaming to distill: Data-free knowledge transfer via deepinversion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8715–8724, 2020.
- [51] Runpeng Yu and Xinchao Wang. Neural lineage. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4797–4807. IEEE, 2024.
- [52] Runpeng Yu and Xinchao Wang. Neural phylogeny: Fine-tuning relationship detection among neural networks. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [53] Jie Zhang, Dongdong Chen, Jing Liao, Weiming Zhang, Huamin Feng, Gang Hua, and Nenghai Yu. Deep model intellectual property protection via deep watermarking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(8):4005–4020, 2022.

A Implementation Details for Building Diverse Model Families

Implementation details for building MobileNet-V2 model families. We also constructed a model family based on the MobileNet-V2 architecture [37]. In this family, the parent and child models were trained on different datasets to simulate cross-dataset transfer experiments. We designed multi-generation fine-tuning paths such as Caltech-101 [20] → Tiny ImageNet [2] → MixedDataset, Flowers [29] → CIFAR100 [42] → MixedDataset, and Tiny ImageNet → Oxford-IIIT Pet [39] → MixedDataset. Here, MixedDataset refers to a set of diverse fine-tuning tasks randomly selected from Aircraft [28], DTD [8], and Dogs [17], in order to increase domain diversity. All models in this family were developed using the same randomization strategies as in the ResNet-18 model library.

Implementation Details for Building Diffusion Model Families. For fine-tuning via DreamBooth, we adopt the official dataset provided by the authors, utilizing all 30 sets of image sets. The fine-tuning is performed using the officially recommended hyperparameters to ensure reproducibility and consistency with standard usage. More specifically, the learning rate was set to 1×10^{-5} . During the fine-tuning process, we save model checkpoints at both 10,000 and 20,000 training steps to serve as representative fine-tuned versions for subsequent evaluation and analysis. For LoRA-based fine-tuning, we employ the publicly available *Naruto* [5] dataset and explore two LoRA rank settings: 128 and 256, following common practices in diffusion model adaptation tasks. The learning rate is set to 1×10^{-5} , and we save the model checkpoint after 10,000 training steps as the fine-tuned version for subsequent evaluation.




Figure 10: Kernel density estimation(KDE) of similarity scores across different model architectures. Each subplot shows the similarity distribution between a child model and its parent, grandparent, great-grandparent and non-lineage items, respectively.

Implementation Details for Building LLM Model Families. We randomly selected a subset of models under *meta-llama/Llama-3.1-8B* that serve as parent models, each having at least one identifiable descendant model—most of which are derived via fine-tuning. For *vidia/OpenMath2-Llama3.1-8B*, since only the merged model is publicly available on Hugging Face, we treat the merged version as its derivation in our analysis. We randomly select a subset of models under *Qwen/Qwen2.5-1.5B* that serve as parent models, covering a diverse set of fine-tuned variants, including those derived via instruction tuning, model merging, and adapter-based fine-tuning. This setup ensures that our evaluation spans a broad spectrum of real-world adaptations commonly applied to Qwen-based models. These descendant model variants, covering both LLaMA and Qwen families, are available in our OpenScience repository.² Based on their Hugging Face model descriptions, we selected the datasets “MMLU”, “LLMTwin”, “OpenMath”, “Guard”, “Tulu”, and “Dolphin” to construct probing datasets, from which we randomly sampled 50 prompts per dataset and collected 5 responses per prompt to capture the models’ knowledge. Due to experimental constraints, we limited our evaluation to LLMs within Hugging Face model repository. Since both grandparent–child and sibling relationships involve models that are separated by two generations of fine-tuning datasets, we approximate the similarity of grandparent–child pairs using that of sibling models. This substitution is reasonable because sibling models, though derived from different fine-tuning paths, reflect a comparable degree of lineage distance.

B Additional Experiments

B.1 Fine-Grained Analysis of the Knowledge Similarity

Effect of Model Architecture. Fig. 10 shows the KDE curves of model families based on different network architectures. To demonstrate the effectiveness of our method in realistic model repository scenarios, we report validation results only on parent–child, non-lineage, and grandparent–child relationships in

LLaMA3 family. As the lineage relationship becomes weaker, the similarity between knowledge vectors gradually decreases. Although in Diffusion experiments, there is no obvious margin between grandparent and great-grandparent relationship, we believe this is mainly due to two reasons. First, the randomly sampled probe dataset does not fully cover the semantic scope of the parent model, and therefore yields highly similar features for those two generations. Second, large foundation models trained on diverse data sources usually retain most of their original semantic representations after fine-tuning, which reduces the semantic gap between successive generations.

Effect of Fine-tuning Rounds. We define one generation as fine-tuning a model once with a downstream dataset. As discussed in the main text, the lineage similarity score predicted by our framework decreases as the number of fine-tuning generations increases. To systematically evaluate this effect, we conduct experiments on both MobileNet-V2 and ResNet-18 families, where models are derived from a single source and thus more sensitive to fine-tuning transformations.

We consider two scenarios: (i) **same-domain fine-tuning**, where the original model is trained on a subset of Caltech101 classes and the downstream fine-tuning tasks are sampled from other Caltech101 categories; and (ii) **cross-domain fine-tuning**, where the downstream tasks are drawn from the Dogs dataset. As shown in Fig. 11, the similarity score consistently decreases with the number of fine-tuning rounds, confirming that iterative fine-tuning progressively obfuscates the lineage attestation trace. Importantly, the scores remain distinguishable from those of non-lineage models: for MobileNet-V2, the score stays above 0.2, and for ResNet-18, it remains above 0.3. This indicates that while fine-tuning reduces knowledge consistency, the lineage signal persists across multiple generations. Moreover, the difference between same-domain and cross-domain settings is marginal, suggesting that the depth of fine-tuning has a stronger impact on lineage verification than task distribution shifts.

B.2 Additional Ablation Studies and Details

Detailed implementation under Trusted Alternative Probe Datasets. For in-domain probing, we resample a new set of

²<https://zenodo.org/records/17958570>.




Figure 11: Knowledge consistency scores with increasing rounds of fine-tuning.

data points from the original training distribution to form an alternative probe set, ensuring that the semantic domain remains consistent while the specific samples differ. For cross-domain probing, we replace the original probe dataset with a trusted dataset drawn from a related but distinct distribution—for example, using ImageNet as a substitute for Caltech-101 in classification models, using CIFAR-100-derived captions as a replacement for the COCO captions in diffusion models, or using *mgs*, *gsm8k*, *hellaswag*, *humaneval*, *arc_challenge* as substitutes for the original probe prompts in LLMs.

Effect of the Key Component in Our Framework. We ablate each component by measuring the TPR drop upon removal, evaluating its impact on lineage discrimination. Results of MobileNet-V2 family are shown in Table 7. Firstly, removing the evolution knowledge vector (w/o \mathbf{h}_Δ) reduces the score to 0.66, confirming that the evolution model effectively captures parent-to-child knowledge transfer. Secondly, substituting the knowledge encoder with mean pooling (w/o Ψ) of feature embeddings yields a TPR of 0.91 but incurs a 0.08 accuracy drop, indicating that Ψ compactly encodes lineage-relevant knowledge beyond raw features. Finally, replacing the knowledge fusion network (w/o Φ) with a simple sum $\mathbf{h}_C + \mathbf{h}_\Delta$ decreases performance from 0.99 to 0.43, demonstrating Φ 's role in integrating inherited and evolved knowledge.

Table 7: Ablation study for the proposed framework, where w/o refers to without.

	our MLA	w/o \mathbf{h}_Δ	w/o Ψ	w/o Φ
TPR	0.99	0.66	0.91	0.43

B.3 Extended Adaptive Attack

Knowledge Infusion Attack. In the False Lineage Claim Rejection scenario, the adversary may attempt to construct a forged parent model f_P^A that exhibits behavior similar to the child model f_C . To achieve this, the adversary queries f_C on the probe dataset and uses its responses to fine-tune f_P^A , thereby distilling probe-induced behavior into the

forged model to fabricate a plausible parent–child relationship. As shown in Table 8, its similarity to f_C remains low (0.31/0.31/0.32 for 6%, 30%, and 60% probe usage), well below the 0.7 threshold. These results indicate that iterative fine-tuning progressively distorts the alignment between parameter evolution and underlying knowledge, leading to decreasing lineage similarity instead of successful lineage reconstruction.

Table 8: Resistance to Knowledge Infusion Attacks.

Probe Usage	Similarity to f_C	Verified($T=0.7$)
6%	0.31	×
30%	0.31	×
60%	0.32	×

C Discussion and Limitations

Practical considerations for third-party evidence collection. MLA operates under a model management platform in which a trusted attestation authority evaluates artifacts submitted by the claimant. This requirement aligns with existing model-governance practices: enterprise registries and public model platforms already maintain model parameters, version histories, and provenance metadata, making the forwarding of platform-certified artifacts—such as the initialization θ_0 —a natural extension of current workflows. Under this setting, authenticity of the submitted model parameters can be ensured, as the malicious claimant cannot modify platform-certified checkpoints. However, verifying the authenticity of auxiliary evidence such as probe datasets is more challenging. MLA focuses on validating claimed lineage relationships rather than reconstructing the primordial training origin of a model, and a trusted authority may not be able to guarantee that a submitted probe dataset is the exact one used during fine-tuning. In such cases, the authority may instead construct alternative probe sets drawn from semantically similar sources. As demonstrated in Fig. 9, semantically aligned probe datasets achieve comparable attestation performance; however, developing stronger mechanisms for verifying dataset provenance remains an open direction for future work.