

Abstraction Alignment: Comparing Model-Learned and Human-Encoded Conceptual Relationships

Angie Boggust
CSAIL

Massachusetts Institute of Technology
Cambridge, Massachusetts, USA
aboggust@mit.edu

Hendrik Strobel
IBM Research AI

Cambridge, Massachusetts, USA
hendrik@strobel.com

Hyemin Bang
CSAIL

Massachusetts Institute of Technology
Cambridge, Massachusetts, USA
hbang@mit.edu

Arvind Satyanarayan
CSAIL

Massachusetts Institute of Technology
Cambridge, Massachusetts, USA
arvindsatya@mit.edu

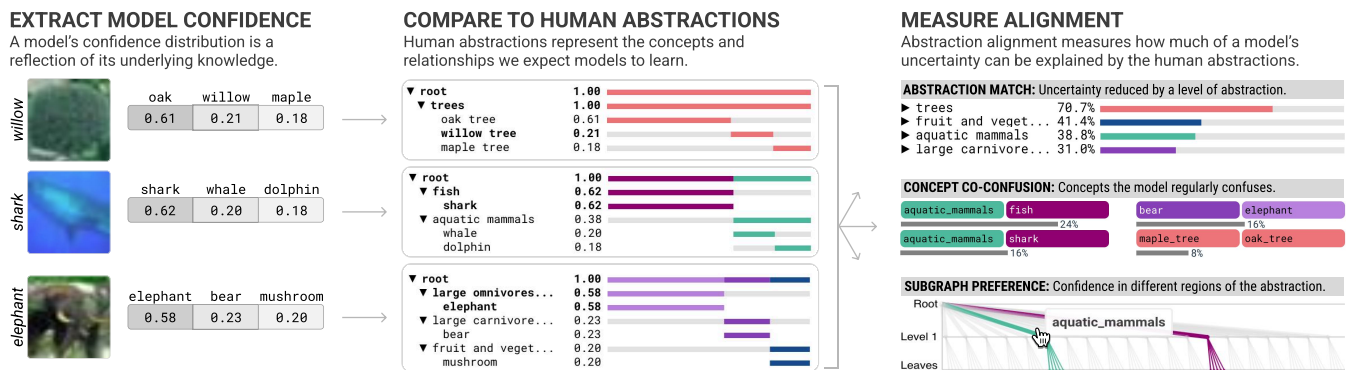


Figure 1: Abstraction alignment measures human-AI alignment by comparing model behavior to known human abstractions.

Abstract

While interpretability methods identify a model's learned concepts, they overlook the relationships between concepts that make up its abstractions and inform its ability to generalize to new data. To assess whether models' have learned human-aligned abstractions, we introduce abstraction alignment, a methodology to compare model behavior against formal human knowledge. Abstraction alignment externalizes domain-specific human knowledge as an abstraction graph, a set of pertinent concepts spanning levels of abstraction. Using the abstraction graph as a ground truth, abstraction alignment measures the alignment of a model's behavior by determining how much of its uncertainty is accounted for by the human abstractions. By aggregating abstraction alignment across entire datasets, users can test alignment hypotheses, such as which human concepts the model has learned and where misalignments recur. In evaluations with experts, abstraction alignment differentiates seemingly similar errors, improves the verbosity of existing model-quality metrics, and uncovers improvements to current human abstractions.



This work is licensed under a Creative Commons Attribution 4.0 International License. *CHI '25, Yokohama, Japan*

© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1394-1/25/04
<https://doi.org/10.1145/3706598.3713406>

CCS Concepts

• Computing methodologies → Natural language processing; Computer vision; Machine learning; • Human-centered computing → Visualization systems and tools.

Keywords

interpretability, human-AI alignment, visualization, abstraction

ACM Reference Format:

Angie Boggust, Hyemin Bang, Hendrik Strobel, and Arvind Satyanarayan. 2025. Abstraction Alignment: Comparing Model-Learned and Human-Encoded Conceptual Relationships. In *CHI Conference on Human Factors in Computing Systems (CHI '25)*, April 26–May 01, 2025, Yokohama, Japan. ACM, New York, NY, USA, 20 pages. <https://doi.org/10.1145/3706598.3713406>

1 Introduction

AI alignment is increasingly critical to meet growing societal and regulatory demands for AI systems that make human-like decisions [136, 143]. To meet these demands, the research community has developed interpretability methods to uncover the concepts models use to reason about their inputs and generate outputs – for instance, using WHEELS to classify car images [50, 104] or TOURIST ATTRIBUTIONS to respond to travel text [142]. By analyzing these concepts, users can better understand model behavior and identify potential misalignments, such as an overreliance on one medical

concept when making complex diagnoses [15, 72] or a propensity for sycophantic responses at the expense of truthfulness [142].

However, existing methods analyze a model’s learned concepts in isolation, quantifying the model’s sensitivity to each concept independently [73, 142]. While such testing procedures identify the importance of each concept to the model’s decision, they overlook the model’s learned relationships between concepts. Yet, conceptual relationships are core to the ability to extrapolate learned concepts to new data, contextualize knowledge across tasks, and flexibly reason at a task-appropriate level of specificity [3, 42, 87, 160]. Interpreting these characteristics of a model requires analyzing its *abstractions*, validating that it has not only learned granular concepts (e.g., SCHNAUZER), but that it also organizes them into progressively more general notions (e.g., DOG and then ANIMAL).

Although existing interpretability techniques have been used to discover that models learn concepts at varying levels of detail [7, 50, 57, 72, 102], analyzing abstractions requires significant manual effort. Users must survey interpretability results piece-by-piece to confirm that the model’s abstractions are human-aligned, checking that all concepts activate as expected [72, 142], the set of concepts is comprehensive for the task [142], and similar concepts are represented similarly by the model [19, 142]. In each of these cases, the process of estimating alignment occurs largely inside the user’s head, requiring significant cognitive effort to compare model concepts against their domain abstractions. Moreover, by relying on individuals’ mental abstractions, existing approaches limit alignment assessment to users with deep domain-specific knowledge, such as medical specialists [15, 72] or chess Grandmasters [127].

To scaffold the process of assessing model alignment, we introduce *abstraction alignment*, a methodology to measure the agreement between a model’s learned abstractions and an explicit human representation of the modeled domain. Abstraction alignment is a form of representational alignment [136] that compares model outputs (a proxy for its internal representations) against codified human abstractions (a proxy for our internal representations). Abstraction alignment begins with a *human abstraction graph*—an agreed upon representation of formal human knowledge containing a set of pertinent concepts spanning levels of abstraction, such as a lexical graph [90] or medical hierarchy [157]. It then measures alignment by evaluating how well the abstraction graph accounts for a model’s decision uncertainty. Through this process, model output probabilities (i.e., confidence in each class or token) are mapped to concepts in the human abstraction graph, and the probabilities of sibling concepts (e.g., GUITARIST and SINGER) are summed and propagated to shared ancestors (e.g., MUSICIAN and ARTIST). The result is the model’s *fitted abstraction graph*, representing its decision making process across concepts at many levels of abstraction.

To aggregate abstraction alignment over many model decisions, we define three metrics. `Abstraction match` measures how much of the model’s confusion is mitigated by moving up a level of abstraction, `concept co-confusion` tests how often the model confuses concepts, and `subgraph preference` quantifies which abstractions the model prefers. By integrating the abstraction alignment metrics into an interactive interface, we enable users, ranging from computer scientists to medical domain experts, to ask and answer alignment hypotheses, such as which human concepts the model has learned and what recurring misalignments the model makes.

We demonstrate how abstraction alignment facilitates alignment analysis through case studies with expert users interpreting an image classification model, benchmarking generative language models, and auditing a clinical ML dataset. In an image classification task, abstraction alignment helps interpret model behavior by distinguishing problematic misalignments from benign low-level errors. Analyzing abstraction alignment across the entire test dataset reveals that most model mistakes are not abstraction-aligned, as it learns to differentiate images using visual abstractions like COLOR instead of the desired human abstractions based on BIOLOGICAL and USAGE differences. These misalignments suggest instances where the model may fail, ways to improve data collection, and possible recategorizations of the human abstraction graph.

Then, we collaborate with three language model researchers, applying abstraction alignment to their alignment task: investigating the specificity of generative language models. While users currently test model specificity by comparing the model’s generated text against a few synonymous words, abstraction alignment expands the verbosity of these benchmarks by testing thousands of words across numerous levels of abstraction. As such, we find that abstraction alignment allows researchers to test more complex hypotheses, such as the model’s preferred level of specificity and which unrelated words the model commonly confuses.

Finally, in a participatory AI setting, four medical experts analyze abstraction alignment to audit the MIMIC-III clinical dataset [67, 68], testing whether its labels reflect appropriate medical abstractions. With abstraction alignment, experts uncover discrepancies between the medical abstractions we expect models to learn and those codified in the dataset. For instance, experts find that how diseases are classified in the dataset does not always align with the World Health Organization’s (WHO) standards [157]. These abstraction misalignments suggest data processing strategies that better reflect human expectations and exposes known issues in the disease hierarchy, some of which have been recently addressed by the WHO [28]. These results signal that, beyond improving the alignment of model representations, abstraction alignment may also inspire improvements to existing human representations.

Abstraction alignment is publicly available with open-source code at <https://github.com/mitvis/abstraction-alignment> and an interactive interface at <https://vis.mit.edu/abstraction-alignment/>.

2 Related Work

2.1 Abstraction and Human Knowledge

Abstraction is the process of distilling many individual data instances into a set of fundamental concepts and relationships that capture essential characteristics of the data [3, 87, 160]. It is a key feature of human cognition as it allows us to flexibly reason at the level of specificity appropriate for our task and generalize our knowledge by fitting abstracted patterns to new data [42, 150, 160]. As a result, abstractions form the basis for human information encodings across domains like linguistics [35, 90], biology [59, 86], and medicine [157, 158]. In machine learning, abstractions are built into many tasks, including image classification [32, 77] and medical coding [67, 68]. Even datasets without built-in abstractions are often linked to existing abstractions by matching their outputs to

corresponding concepts [121]. Encouragingly, researchers have recently integrated human abstractions into model training pipelines, resulting in increased model generalization [97], and advocated for using conceptual relationships, like abstractions, to advance our understanding of foundation models [153]. Building on this rich history, abstraction alignment leverages abstractions to better understand human-AI alignment.

Related research has studied formal representations of human knowledge, known as knowledge graphs [61, 66]. Knowledge graphs reflect the relationships between entities, like distance (EIFFEL TOWER *is near* → ARC DE TRIOMPHE) or connectivity (BOS *direct flight* → MEX) [61]. In abstraction alignment, we represent human abstractions as an abstraction graph, a type of knowledge graph where nodes are concepts and edges encode abstractions from specific to general concepts (e.g., CARDIOLOGIST *type of* → DOCTOR or MONTREAL *located in* → QUEBEC). We make this distinction because we are interested in understanding whether a model has learned to reason with human-like abstractions. These human abstraction graphs provide an explicit representation of formal human knowledge (Section 3.1) that allows us to quantify alignment.

2.2 AI Alignment and Interpretability

Aligned with our goal of understanding machine learning models, interpretability research measures model reliance on known human concepts [37, 118]. For instance, saliency methods highlight relevant input features [15, 25, 114, 128, 133]; example-based methods derive influential inputs [76, 93, 159, 164]; feature visualizations identify concepts that activate model neurons [7, 41, 57, 83, 105]; and concept-based [50, 73] and mechanistic methods [19, 39, 56, 83, 84, 89, 104, 142] identify human concepts encoded in a model’s latent space. Together these methods have identified problematic model correlations [24], made sense of complex neuron activations [26, 84, 98, 106], and discovered novel concepts that advance human understanding [127]. Building on their success, abstraction alignment expands these methods from identifying independent concepts to understanding the relationships between them, ensuring that models learn human-aligned concepts and abstractions.

At the same time, visualization research has explored how to communicate interpretability results to users. Visualizing interpretability results — e.g., saliency heatmaps [17, 71, 126, 133], embedding scatterplots [14, 132, 154], and feature dictionaries [27, 85, 142] — has enabled greater meaning-making from experts and engagement from lay users [9]. Interactive interfaces [8, 20, 122, 135, 155] have allowed users to perform error analysis [140, 159], track model provenance [1, 16], inspect decision boundaries [134], and intervene on models [62]. We instantiate abstraction alignment in an interactive interface (Section 4), allowing users, from ML researchers to domain experts, to actively participate in alignment tasks.

Abstraction alignment also follows a rich history of human-AI alignment research [136, 143], studying methods for measuring [6, 82, 96, 103, 120], bridging [52, 119, 127], and increasing [60, 97, 115, 145] the alignment of model and human representations. Within the alignment framework developed by Sucholutsky et al. [136], abstraction alignment is a form of behavioral alignment [49, 108] where model outputs serve as proxies for its internal representations. Our experiments examine representational alignment by

comparing model representations (proxied by model outputs) and human representations (proxied by the human abstraction graph) across a set of evaluation data. In Section 3.4, we characterize abstraction alignment using the alignment framework developed by Sucholutsky et al. [136]. Unlike prior studies that focus on a single metric to quantify and optimize alignment, abstraction alignment offers a methodology for evaluating how closely model behaviors reflect formal human knowledge. As a result, there are many possible metrics that capture specific aspects of abstraction alignment (we define three in Section 3.3) and abstraction alignment facilitates qualitative, interactive human analysis (Section 5).

2.3 Uncertainty Estimation

Abstraction alignment relies on model uncertainty to compute alignment, working under the assumption that the model’s uncertainty reflects its learned abstractions. Model uncertainty can be broadly categorized as aleatoric [100, 123], arising from irreducible observational noise like noisy data and labeling errors, or epistemic [151], stemming from limited knowledge like insufficient training data or out-of-distribution inputs [46, 64, 74, 165]. Relatedly, uncertainty quantification research focuses on accurately extracting these uncertainties from models such that the model’s purported confidence is an interpretable measure of correctness [2, 29, 47, 51, 78, 79, 163]. Instead of classifying or adjusting model uncertainty, abstraction alignment uses it as a proxy for the model’s internal representations. As a result, abstraction alignment is agnostic to the type of uncertainty, because both types reflect the model’s internal conceptual boundaries and future behavior. Whether the uncertainty arises because the model lacks training data to distinguish a concept (epistemic) or because humans also confuse the concept (aleatoric), it nevertheless represents the model’s understanding of that concept.

3 The Abstraction Alignment Methodology

The goal of abstraction alignment is to measure how well a model’s behavior aligns with human abstractions. Our method is based on the assumption that a model’s uncertainty is a reflection of its learned abstractions. That is, concepts the model commonly confuses are more similar in its abstractions than concepts it perfectly separates. For instance, if a model has learned human abstractions, then, in aggregate, it should be more likely to confuse APPLES with other FRUITS than with unrelated concepts, like MOTORCYCLES.

While there are likely many methods for measuring abstraction alignment, we take a post hoc and model-agnostic approach that compares model outputs against existing human abstractions. To compute abstraction alignment, we represent human abstractions as an *abstraction graph* (Section 3.1). We compare the model’s behavior to human abstractions by mapping the model’s output options (e.g., its classes or tokens) to concept nodes in the human abstraction graph (Section 3.2). Given a dataset instance, like an image or a sentence, we compute the model’s *fitted abstraction graph*, a weighted version of the abstraction graph representing the model’s confidence in a range of concepts across multiple levels of abstraction. We use the model’s fitted abstraction graphs to define abstraction alignment metrics that quantify how well the human abstractions account for the model’s behavior (Section 3.3).

```

1 # Pseudocode to compute the model's fitted abstraction
2 # graph for one dataset instance.
3 def fit(abstractionGraph, model, outputs, instance):
4     # Initialize the values in the human abstraction graph.
5     fittedAbstractionGraph = abstractionGraph.copy()
6     for node in fittedAbstractionGraph:
7         node.value = 0
8
9     # Set node values based on the model's confidence.
10    probabilities = model(instance)
11    for i, probability in enumerate(probabilities):
12        node = fittedAbstractionGraph.getNode(outputs[i])
13        node.value = probability
14
15    # Propagate the values from leaf to root.
16    for level in reverse(fittedAbstractionGraph.levels):
17        for node in level:
18            for child in node.children:
19                node.value += child.value
20
21    return fittedAbstractionGraph

```

Figure 2: To compare model behavior with human abstractions, abstraction alignment computes a fitted abstraction graph for each model decision. First, we map the model’s output space to concepts in the human abstraction graph. Then, we assign each concept a value corresponding to the model’s confidence in that concept or any of its descendants. The resulting fitted abstraction graph represents the model’s confidence in a range of concepts across levels of abstraction.

3.1 Representing Human Abstractions

Abstraction alignment shifts the alignment process from mentally estimating alignment using a human’s internal knowledge to externally inspecting precomputed alignment results. To do so, we externalize formal human knowledge as a directed acyclic graph (DAG) called the *human abstraction graph*. We define an abstraction graph as a type of knowledge graph where nodes represent concepts and edges represent abstraction relationships between precise and broad concepts. For example, in the medical abstraction graph in Section 5.3, nodes represent medical diagnoses and edges represent the abstractions between specific diagnoses, like `FRONTAL SINUSITIS`, and broader diagnostic categories, like `RESPIRATORY INFECTIONS` [157]. Formally, the human abstraction graph is a DAG G containing a set of nodes $N := \{n_k\}$. Nodes are distributed across levels $L := \{l_h\}$ where each $l_h \subseteq N$, and a node’s level is defined as the length of the longest path (h) from the node to a root.

DAGs are well suited to representing human abstractions because they efficiently encode both human concepts and abstraction relationships. We can easily access a concept’s level of abstraction by measuring its height and move up and down levels of abstraction by accessing its ancestors or descendants. Since the graph is acyclic, it guarantees the hierarchical structure that underpins abstraction. Further, DAGs are commonly used to represent human abstractions [90, 157] and are built into ML datasets [32, 67, 68, 77], allowing abstraction alignment to apply to various domains.

We use human abstraction graphs as agreed-upon, external representations of formal human knowledge. While they may not perfectly match any individual’s internal abstractions, they are useful proxies as they reflect collective human meaning-making. For example, while we may not individually know every word and relation in

the WordNet lexical graph, it nevertheless represents collective English language abstractions that we use to communicate [90]. These graphs are often shared between individuals [95, 101, 137, 147], used to educate newcomers to a domain [157], and employed when building additional knowledge representations [32, 77].

3.2 Integrating Model Outputs with Human Abstractions

To compare the model’s behavior against human abstractions, we map the model’s output space (e.g., its classifiable classes or generatable tokens) to nodes in the human abstraction graph. This defines a mapping from each of the model’s outputs $O := \{o_j\}$ to a node $n_k \in G$ that corresponds to the same concept. Often this mapping is straightforward because the human abstraction graph is built into the modeling task—e.g., the CIFAR-100 classes are mapped to higher-level concepts [77] (Section 5.1). However, even when the human abstraction graph is separate from the modeling task, the model’s outputs can often easily be mapped to the graph’s nodes. For instance, in Section 5.2, we map language model tokens to words in the WordNet lexical graph [90] by searching WordNet for the token’s most similar definition.

With a mapping from model output to concept node, we can now compare the model’s behavior against the human abstractions. To do so, we compute a *fitted abstraction graph*, a weighted version of the human abstraction graph representing the model’s confidence in each concept for a given decision. Following the algorithm in Figure 2, we compute a fitted abstraction graph for every instance in an evaluation dataset $D := \{d_i\}$. Given an instance d_i , like an image or sentence, we extract the model’s probability for each possible output concept o_j . Then, we assign each node in the human abstraction graph n_k a value v_{ik} equal to the model’s probability for the node’s concept or any of its descendants. For example, given a CIFAR-100 image classification model as in Section 5.1, the value of `FLOWER` is the sum of the model’s confidence that the given image is a `POPPY`, `ROSE`, `TULIP`, `ORCHID`, or `SUNFLOWER`. By propagating the model’s probabilities through the abstraction graph, the fitted abstraction graph provides a measure of the model’s confidence in a range of concepts across levels of abstraction.

So far, we have defined fitted abstraction graphs using an ML model, but we can also use them to represent other types of encoded abstractions. We can compute a fitted abstraction graph for any function that maps dataset instances d_i to a distribution over human concepts O . This function, $f : D \mapsto \mathbb{R}^{|O|}$, is often the forward function of a machine learning model, such as an image classifier (Section 5.1) or language generation model (Section 5.2). However, as we demonstrate in Section 5.3, this function can also represent the information encoded in a dataset, where f maps clinical notes d_i to clinical codes o_j , assigning each o_j a value based on whether a human labeled the note with that code. In this case, the fitted abstraction graphs represent the alignment between human labeling patterns f and expected medical abstractions G .

3.3 Measuring Abstraction Alignment

The model’s fitted abstraction graphs support various alignment hypotheses, such as identifying concepts prone to misalignment and determining the model’s preferred level of abstraction. While there

are many alignment metrics one could define across the fitted abstraction graphs, we propose three metrics that have proven useful in our alignment analysis of computer vision classification models, generative language models, and medical datasets (Section 5).

Abstraction Match. One way to measure abstraction alignment is to measure how well the human abstractions account for the model’s uncertainty. If the model’s confusion is substantially reduced by moving up a level of abstraction, then the model’s behavior is more abstraction-aligned than if it continues to be confused at higher-levels of abstraction. While there are cases when the model’s confusion may acceptably not fit the human abstractions, such as confusion on an image containing multiple objects, in aggregate we expect the model’s uncertainty to reflect its abstractions – i.e., it will confuse concepts that it considers similar.

We measure abstraction match as the amount of the model’s decision entropy that is reduced by moving up a level of the abstraction graph. Given two levels (l_g and l_h), we compute the difference in entropy, $H(V) = -\sum_{v \in V} v \log(v)$ [129], between the node values $V := \{v_{ik}\}$ at each level. The larger the entropy, the more confused the model is across concepts at that level of abstraction. If the entropy decreases substantially then the model’s behavior aligns with the abstraction mapping the low-level nodes to the higher-level nodes. We aggregate abstraction match across a set of data instances D , which can be a single instance, the entire dataset, or an informative data subset.

$$M(l_g, l_h) = \frac{1}{|D|} \sum_{i=1}^{|D|} H([v_{ik} \forall n_k \in l_h]) - H([v_{ik} \forall n_k \in l_g]) \quad (1)$$

Subgraph Preference. Another valuable metric is to compare the values of different fitted abstraction subgraphs. For instance, in Section 5.2, we compare subgraphs that represent different concepts (e.g., any location vs. CANADIAN locations) and different levels of abstraction (e.g., concepts more specific than JOURNALIST to concepts more general than JOURNALIST). In aggregate, these comparisons help us quantify and compare abstractions the model prefers.

We compute subgraph preference by measuring how often the aggregate value of a node in one subgraph, s_a , is larger than the aggregate value of a node in another subgraph s_b . This is an extension of the specificity testing metric proposed by Huang et al. [63], where s_a represents the specific concept and s_b represents the general concept. However, while the prior metric was designed to test two concepts, abstraction alignment allows us to test a breadth of concepts, including different levels of abstraction, multiple similar concepts, and concepts related to different abstractions. If the model’s outputs span multiple levels and many concepts in the abstraction graph, we can either compute subgraph preference using the nodes’ values (as in Section 5.2.1) or use the unpropagated model probabilities as the nodes’ values (as in Section 5.2.2).

$$P(s_a, s_b) = \frac{1}{|D|} \sum_{i=1}^{|D|} 1[\max([v_{ik} \forall n_k \in s_a]) > \max([v_{ik} \forall n_k \in s_b])] \quad (2)$$

Concept Co-confusion. Finally, the concept co-confusion metric allows us to measure how often a model assigns probability to pairs of concepts. While concepts that are ancestors or descendants of each other will definitionally have high concept co-confusion, unrelated concepts with high concept co-confusion are unrelated human concepts that the model deems similar.

To compute concept co-confusion for a pair of nodes, we compute the entropy (H) [129] of their values divided by the maximum possible entropy for a pair of nodes. By computing the entropy, we weight the concept co-confusion by how confused the two nodes are – e.g., concept co-confusion for nodes with values 0.4 and 0.6 will be higher than nodes with values 0.9 and 0.1 because the model is more confused between the first pair of concepts. We compute concept co-confusion over the data D to identify repeated confusion.

$$C(n_k, n_l) = \frac{\sum_{i=1}^{|D|} H([v_{ik}, v_{il}])}{\sum_{i=1}^{|D|} H([0.5, 0.5])} \quad (3)$$

3.4 Formalizing Abstraction Alignment as Representational Alignment

Representational alignment is a paradigm for measuring the similarity of two systems’ internal representations [136]. Here, we define abstraction alignment using the representational alignment formalism from Sucholutsky et al. [136]. Abstraction alignment compares machine learning model or dataset abstractions (system A) against formal human knowledge (system B) across a set of dataset instances (D). We use the model’s decisions (Q) as a proxy for its internal representations and a human abstraction graph (W) as a proxy for internal human representations. We compute abstraction alignment by comparing Q and W using the fitted abstraction graphs (Section 3.2) and abstraction alignment metrics (Section 3.3).

Data D : We measure abstraction alignment across a set of evaluation data $D := \{d_i\}$. In our experiments, D consists of images or text, but abstraction alignment applies to any data modality.

System A : System A refers to either the machine learning model or dataset under investigation. When A is a model, the focus is on measuring the alignment of its behavior. When A is a dataset, the goal is to evaluate the alignment of its labels. Accordingly, $f : D \mapsto \mathbb{R}^{|O|}$ represents either the model’s function mapping inputs to its probability distribution over outputs or the function mapping dataset instances to their labels.

- **Measurements X :** System A ’s measurements, $X \in \mathbb{R}^{|D| \times |O|}$, is a matrix of model outputs or dataset labels obtained by applying f to each data instance, $X := [f(d_1), \dots, f(d_{|D|})]$.
- **Embeddings Q :** Since abstraction alignment directly studies the model outputs or dataset labels, we let $Q := X$.

System B : System B represents formal human knowledge. Accordingly, $g : D \mapsto \mathbb{R}^{|G|}$ maps the entire dataset D to a relevant human abstraction graph G containing nodes N . It synthesizes domain-specific human knowledge (represented by the data) into core concepts and their relationships (represented by the graph).

- **Measurements Y :** System B ’s measurements, $Y \in \mathbb{R}^{|G|}$, are a human abstraction graph relevant to the data, $Y := g(D) = G$. For comparison with system A , we assume that the output concepts are a subset of the human concepts $O \subseteq N$ in G .
- **Embeddings W :** Since Y already represents human abstractions, we do not apply additional transformations ($W := Y$).

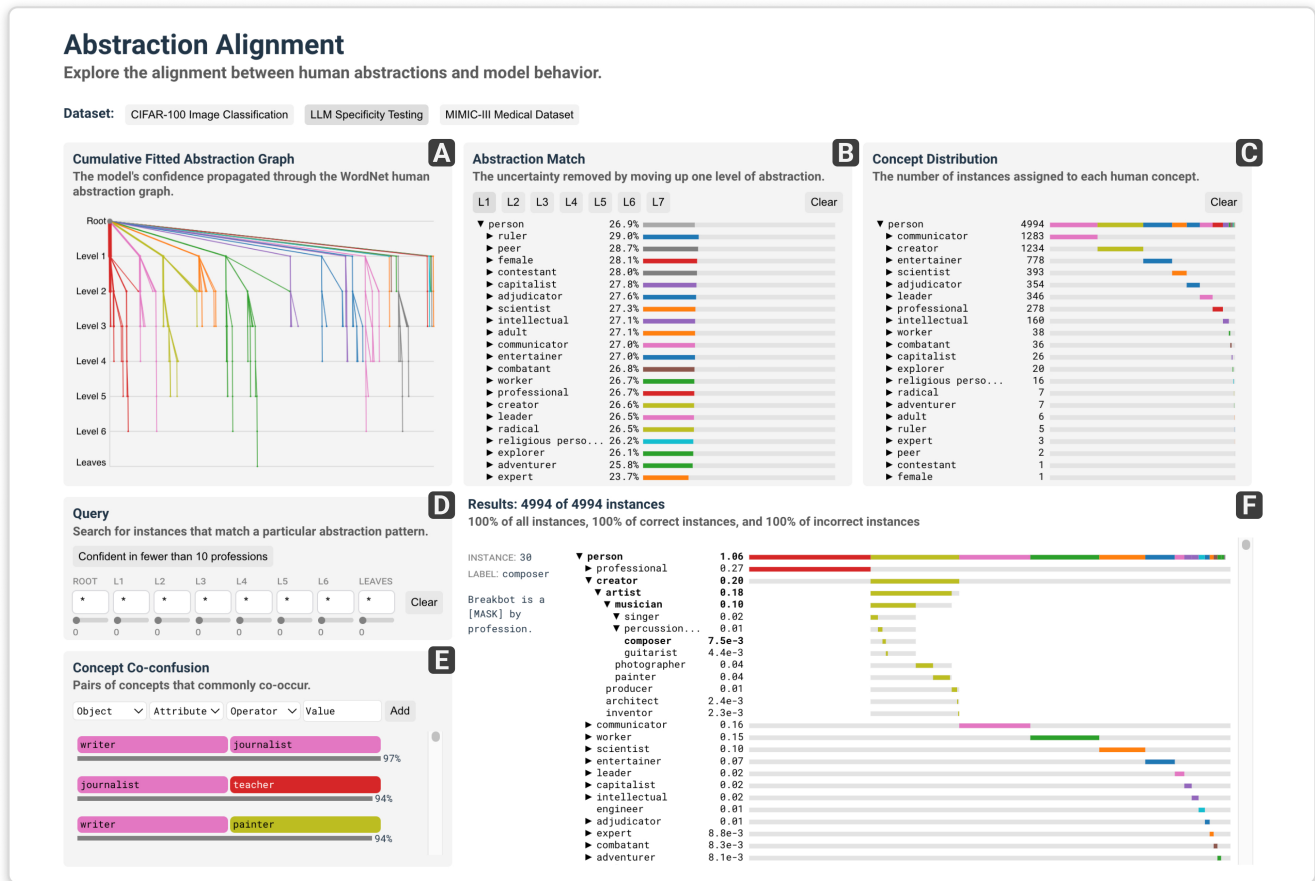


Figure 3: The abstraction alignment interface visualizes a model’s alignment with human abstractions. It displays the cumulative fitted abstraction graph (A), aggregated abstraction match (B), concept distribution (C), and concept co-confusion (E). Interacting with these panels or the query bar (D) updates the instance list (F) to show the fitted abstraction graphs of relevant inputs.

Alignment Function $\delta(Q, W)$: To compare the model outputs or dataset labels (Q) against the human abstraction graph (W), we first create a fitted abstraction graph by projecting the concepts and probabilities in Q onto the graph of concepts in W (Section 3.2). Instead of providing a single quantifier of alignment, this graph-based comparison allows us to define a family of metrics, measuring multiple facets of alignment (Section 3.3).

4 The Abstraction Alignment Interface

We instantiate abstraction alignment in an interactive visual interface (Figure 3) consisting of six interconnected components that express the abstraction alignment metrics (Section 3.3) and fitted abstraction graphs (Section 3.2). Users can interactively test alignment hypotheses by filtering to individual concepts, selecting entire graph regions, and querying for alignment patterns (Figure 4).

Cumulative Fitted Abstraction Graph. The cumulative fitted abstraction graph (Figure 3A) serves as an overview of the model’s fitted abstraction graphs and a visualization of the human abstraction graph. We compute it by summing the fitted abstraction graph

(Section 3.2) for every dataset instance such that a node’s value in the cumulative fitted abstraction graph is the sum of that node’s value across every model decision. We visualize it as a vertical graph across the levels of abstraction, from the most abstract (root) to the most specific concepts (leaves). To display meaningful visual groupings, we color nodes and edges based on their level-1 ancestor. Node radius and edge width reflect the cumulative value assigned to each concept. For instance, in Figure 3A, the thick red edge for PROFESSIONAL indicates that is assigned high confidence in nearly every model prediction. To minimize overlapping edges, we use a recursive depth-first layout and sort children based on their value. Hovering over a node reveals its name, value, and contributing instances, clicking selects the concept and its relatives, and double-clicking selects just that concept (Figure 4A). Upon selection, the interface updates to display the relevant dataset instances.

Abstraction Match. The abstraction match component (Figure 3B) instantiates the abstraction match metric (Eq. (1)), highlighting abstractions the model has learned most effectively. For every concept and level of abstraction, we compute the proportional

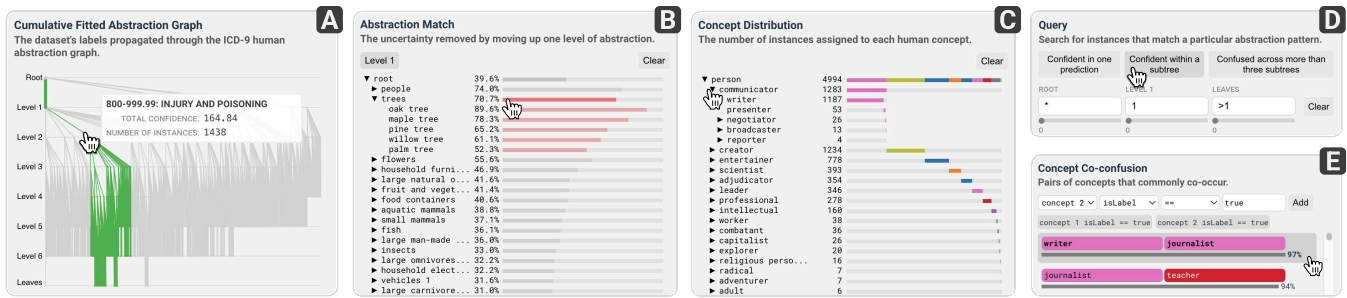


Figure 4: Interacting with the abstraction alignment interface allows users to explore alignment hypotheses. Users can select a concept (A–C), a concept pair (E), or define an alignment query (D) to update the interface with relevant dataset instances.

abstraction match between that level and the next using every dataset instance whose label is a descendant of the concept. For example, for occupation prediction, the level-1 abstraction match for SCIENTIST is shown as the percent decrease in entropy by moving from level-2 to level-1 over all instances whose true occupation is a descendant of SCIENTIST (e.g., PHYSICIST or ASTRONOMER). The abstraction match component is displayed as a nested horizontal bar chart, with bars corresponding to concepts in the human abstraction graph and their lengths representing their percent abstraction match. Selecting a bar (Figure 4B) updates the interface, displaying results for the set of instances that contributed to that concept’s abstraction match score.

Concept Distribution. The concept distribution component visualizes how dataset labels are distributed across levels of abstraction. For each concept in the human abstraction graph, its concept distribution value represents the number of dataset instances whose label is a descendant of that concept. For example, in Figure 3C, the concept distribution component shows that 1,238 instances in the occupation prediction dataset are labeled as a type of COMMUNICATOR. Like abstraction match, the concept distribution is displayed as a nested horizontal bar chart. Selecting a bar updates the interface to display dataset instances labeled under that concept (Figure 4C).

Query. The query component allows users to search for types of model behavior defined over the abstraction graph (Figure 3D). We define a query as an ordered list of layer-wise subqueries that measure the distribution of values across nodes in that layer. A layer’s subquery can be a wildcard (*) that matches any distribution, an integer that defines the number of nodes in that layer with a non-zero value, or a probability distribution (list) of numbers in the range [0, 1] that defines the relative node scores in a layer. We incorporate the modifiers not (!), greater than (>), and less than (<) to expand query expressivity. A query matches an instance if every layer in its fitted abstraction graph matches its corresponding subquery. We can use the results of a query to understand how frequently an alignment pattern occurs and what its common outcomes are. For instance, given a human abstraction graph with three levels, we can query for instances where the model distributes its confidence over multiple leaf nodes in the same subgraph as [* , [1] , >1] (Figure 4D). In the interface, we provide informative pre-defined queries in natural language to help users get started with querying.

Concept Co-confusion. The concept co-confusion component instantiates the concept co-confusion metric (Eq. (3)) to reveal pairs of concepts that the model commonly confuses (Figure 3E). It is visualized as a list of concept pairs, sorted by their concept co-confusion. Each concept in a pair is colored based on its level-1 concept and shown above a sparkline [148] representing the pair’s concept co-confusion. Users can filter the list of concept pairs based on attributes of an individual concept (i.e., its height, depth, name, and if it is a label) or the concept pair (i.e., if they are connected, share a parent, or share an ancestor) (Figure 4E). Selecting any pair updates the rest of the interface to show instances that contributed to that pair’s concept co-confusion.

Instance List. Finally, to allow users to drill down into the model’s alignment on individual decisions, the abstraction alignment interface displays a fitted abstraction graph (Section 3.2) for every dataset instance (Figure 3F). Each fitted abstraction graph is displayed as a nested horizontal bar chart, where a bar represents a node in the fitted abstraction graph and its length represents that node’s value. Bars are colored based on the level-1 concept and the root bar shows a summary of the bars below it. The fitted abstraction graphs are displayed next to the instance text or image and its true labels. To visually indicate salient areas of the fitted abstraction graph, we bold the text corresponding to the instance’s labels and any of its direct relatives. When selections occur in other interface elements, the instance list updates to display the relevant instances. We display summary statistics above the list, showing the number of instances selected and, in classification settings, the proportion of them that are correctly and incorrectly classified.

5 Evaluative Case Studies with Domain Experts

By externalizing formal human knowledge as an abstraction graph, abstraction alignment expands current alignment workflows from internalized comparison to iterative hypothesis testing. To evaluate this shift in perspective, we emulate three real-world alignment tasks across computer vision, natural language, and medicine. First, in Section 5.1, we apply abstraction alignment to interpret an image classification model, finding that it expands interpretation from narrow questions about why a model made a specific decision to broad explorations of the human concepts it has learned. Next, in Section 5.2, we collaborate with researchers to benchmark the specificity of language model responses, revealing that abstraction

Table 1: We evaluate abstraction alignment through case studies with seven experts in language model specificity (§5.2) and medical dataset analysis (§5.3).

Language Model Specificity Case Study (§5.2)			
ID	Title	Affiliation	Role
P1	Professor	University	Tests model specificity
P2	Research Scientist	Tech Company	Tests model specificity
P3	Project Manager	Tech Company	Builds LLM benchmarks

Medical Dataset Analysis Case Study (§5.3)			
ID	Title	Affiliation	Role
P4	Medical Coder	Medical Center	Codes clinical notes
P5	Medical Coder	Medical Center	Codes clinical notes
P6	ICD Manager	Health Org.	Oversees ICD usage
P7	ICD Manager	Health Org.	Oversees ICD usage

alignment expands their conventional benchmarks of isolated pairwise comparisons to more comprehensive comparisons across the entire space of potential outcomes. Finally, in Section 5.3, we leverage abstraction alignment earlier in the ML pipeline, using it with healthcare professionals to assess the human alignment of a medical dataset, revealing discrepancies between medical abstractions and their real-world usage.

Study Method. To evaluate how abstraction alignment supports real-world alignment analysis, we collaborate with seven domain experts across two case studies: language model specificity (Section 5.2) and medical dataset analysis (Section 5.3). We conducted in-depth, semi-structured interviews with each expert to assess how abstraction alignment influenced their analysis. We began with questions about their domain expertise, alignment workflows, and desired outcomes, such as “Tell me about your role as a [title]?” and “How do you currently measure alignment in [case study task]?”. Next, we introduced the abstraction alignment interface (Section 4) using tasks and datasets representative of their domain. Finally, we prompted experts to think aloud as they engaged with abstraction alignment to identify ways the model or dataset was aligned or misaligned with their domain knowledge. This approach allowed us to understand experts’ current processes, observe how abstraction alignment functioned in context, and assess its potential to address experts’ alignment goals. We discuss study limitations in Section 6.

We targeted expert participants to understand how abstraction alignment could impact real-world domains. To identify experts, we reached out to authors of relevant literature, attendees of specialized conferences, and LinkedIn professionals with applicable expertise. We purposively sampled [23] seven participants, ensuring they had deep familiarity with their case study – language model participants regularly tested language models and medical dataset participants had extensive experience with medical codes (Table 1).

We conducted six video interviews each lasting 30–60 minutes (P6 and P7 opted to interview together as colleagues). Our institution deemed our study exempt from full IRB approval, and participants received \$50 gift cards. With consent, all interviews were recorded, resulting in 223 minutes of audio/video data and transcripts. We conducted a thematic analysis, reviewing recordings

and transcripts to code key observations, such as cases where participants recognized an alignment/misalignment (e.g., the dataset is missing domain-specific abstractions), expressed ways abstraction alignment facilitated/hindered their analysis (e.g., it replicated their existing experiment design), or identified an insight that led them to hypothesize about the downstream impact (e.g., the model is overly specific at the expense of correctness). After analyzing recordings individually, we grouped codes into higher-level themes, and used them to structure the results in Section 5.2.1 and Section 5.3.

5.1 Interpreting Image Model Behavior

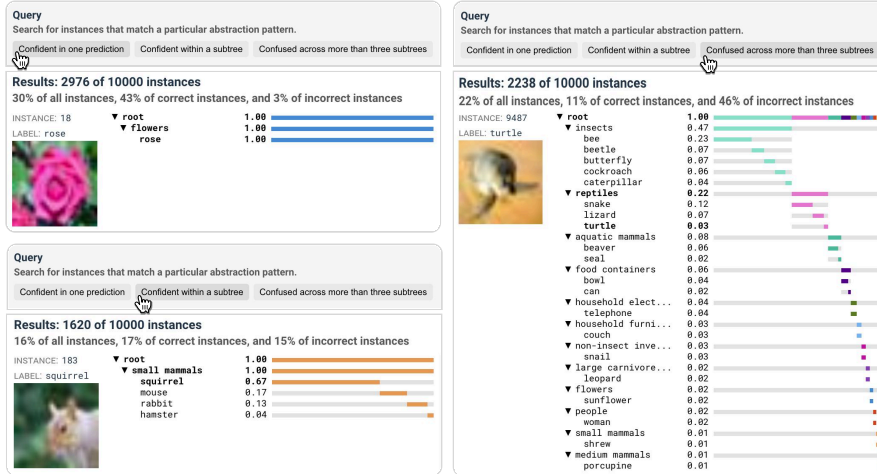
A common interpretability task is understanding a model’s mistakes; however not all mistakes are equally problematic. For instance, in an autonomous driving task, mistaking a TRUCK for a BUS might be harmless, whereas models that mistake a TRUCK for the SKY have, unfortunately, caused real-world accidents [81, 144]. We are more likely to forgive the first mistake because it more closely aligns with our human abstractions – TRUCKS and BUSES are both VEHICLES and we treat them similarly while driving. However, the latter mistake could indicate more severe model generalizability issues where the model’s abstractions do not align with accepted human reasoning (i.e., TRUCKS and the SKY are vastly different concepts that require different interactions).

Applying abstraction alignment to this setting, we find that it helps interpret model behavior and differentiate the severity of a model’s mistakes by expanding the number and complexity of concepts we use to characterize model decisions. In particular, we use abstraction alignment to interpret a ResNet20 [54] computer vision model trained on CIFAR-100 [77]. We use the CIFAR-100 class and superclass structure as the human abstraction graph, as it maps low-level classes, like TRUCK, into higher-level concepts, like VEHICLES [77]. The result is a human abstraction graph with 121 nodes across three levels of abstraction. We compute each test image’s fitted abstraction graph by applying a softmax to the model’s outputs and propagating them through the human abstraction graph.

5.1.1 Interpreting Model Decisions. Analyzing the abstraction alignment of an individual instance can indicate why the model made a particular decision. An instance’s fitted abstraction graph represents how the model made its decision on that instance and the conceptual similarity of other options it considered. For example, in Figure 1, we show three CIFAR-100 test images, their fitted abstraction graphs, and the model’s output probability. The probability distribution for each image follows an approximately 60/20/20 split, so we might assume the model is similarly confused about each image. However, the fitted abstraction graphs reveal that the model’s abstraction alignment differs significantly across images. In the top example, the model’s probability is split between three classes within the same concept, indicating the model is confident the image is a TREE and simply unsure of the species. Whereas, in the bottom instance, the model assigns probability to three distinct high-level concepts, including FRUIT AND VEGETABLES and LARGE OMNIVORES AND HERBIVORES, indicating that it is very confused about this image or has learned a relationship that does not align with human expectations (e.g., a color relation where ELEPHANTS and MUSHROOMS are both GRAY). In a real-world setting, we may be willing to overlook a model’s abstraction-aligned errors (like in

A Interpreting Model Decisions

When the model is correct, it is typically confident in a single decision. But, when it is incorrect, it is often confused across multiple high-level concepts. This suggests the model's confusion stems from misalignments with human concepts.



B Identifying Misaligned Concepts

The model is more aligned (i.e., has higher abstraction match) with concepts like PEOPLE than VEHICLES 2.

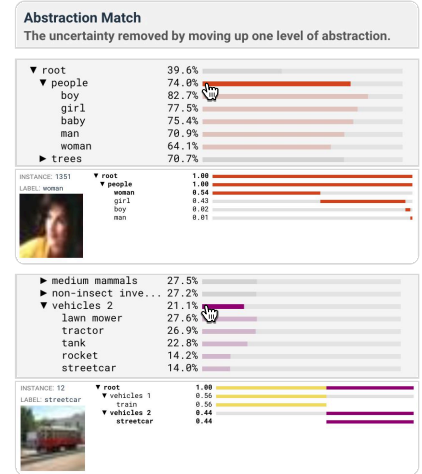


Figure 5: Abstraction alignment offers insights into the behavior of an image classification model. Querying alignment patterns distinguishes benign lack-of-specificity errors from more problematic misalignments (A), and analyzing abstraction match identifies which human concepts the model aligns with, highlighting potential failure cases (B).

the top image) and may want to penalize the model for unaligned confusion even if its answer is correct (as in the bottom image).

Analyzing the abstraction alignment of an instance can also reveal places where human abstractions could be updated to better fit the modeling task. For instance, analyzing the fitted abstraction graph for the middle image in Figure 1 shows that the model splits its output probability between three classes: SHARK, WHALE, and DOLPHIN. Confusion across these three classes may seem abstraction aligned because they are LARGE OCEAN ANIMALS. However, our human abstraction graph splits them into separate high-level concepts because it is based on biological properties, where SHARKS are FISH and WHALES and DOLPHINS are AQUATIC MAMMALS. While the biological principles that separate FISH and AQUATIC MAMMALS (e.g., gills or blow holes) are important to zoology, they are visually subtle and unlikely to come across in low-resolution CIFAR-100 images. Given the model only has access to images, it makes sense that it could learn a visual abstraction where SHARKS and WHALES are closely related. If our use case requires the model to learn biological abstractions, we might consider training on a different dataset with images that visually distinguish biological properties or contain additional metadata about the animal. However, if we are only interested in the model's visual alignment, we may define a new human abstraction graph based only on visual similarity.

5.1.2 Uncovering Global Patterns in Model Behavior. Analyzing abstraction alignment across many model decisions can identify recurring patterns of misalignment that can impact the model's generalizability. To analyze abstraction alignment across an entire dataset, we measure how often a model exhibits a particular type of abstraction alignment/misalignment. We define types of abstraction alignment as queries describing the number of nodes the model considers at each level and how it distributes its confidence across nodes. To measure how often the model's decisions align with

human abstractions we query for instances where the model is confident in a single class or becomes confident in a single high-level concept. In Figure 5A, we see these instances represent nearly half of the model's decisions but only 18% of model mistakes. This means most of the model's mistakes are not harmless low-level errors, but the result of confusion at higher levels of abstraction. Digging into this further, we query for instances where the model considers at least four disjoint concepts and find that almost a quarter of all instances and half of mistakes fall into this category.

5.1.3 Identifying Conceptual Alignment. Given the model and human abstractions seemingly conflict, we can use abstraction alignment to identify which human abstractions the model has learned. To do so, we analyze the model's abstraction match view (Figure 5B) and find that the PEOPLE and TREE abstractions resolve a large proportion of the model's uncertainty, indicating that the model has learned those abstractions. For instance, 74% of the model's uncertainty on images of BABIES, BOYS, GIRLS, MEN, and WOMEN is resolved at the parent concept PEOPLE. While our model only achieves 67.7% test accuracy, seeing that it has learned some human abstractions may increase our trust that its confusion in these categories is harmless, particularly if our setting does not require fine-grained classifications.

Our abstraction alignment metrics also reveal areas where the model is misaligned with human abstractions. For instance, the model's uncertainty is not accounted for by abstractions like VEHICLES 2 nor does it appear to learn animal categorizations like NON-INSECT INVERTEBRATES and MEDIUM MAMMALS. In both cases, we might consider these results to be acceptable model performance in light of ill-fitting human abstractions. In particular, the CIFAR-100 hierarchy artificially restricts each high-level concept to contain exactly 5 children — a constraint that produces two nodes for VEHICLES that arbitrarily distinguish their children rather than

meaningfully capture abstracted patterns. In contrast, although the animal categories are semantically meaningful, they reflect biological concepts like size (MEDIUM) and reproduction (MAMMALS) that are seemingly hard for a model to learn from 32x32 images. If learning accurate biological abstractions are important for our task, then we may prefer to train on an alternate dataset that more precisely expresses these characteristics; on the other hand, if learning visual abstractions is acceptable, we may update our human abstraction graph to better reflect what can be learned from the data (i.e., categorizing animals based on visual similarity). With abstraction alignment we have revealed the model’s learned abstractions, cases of misalignment with human expectations, and mitigation strategies to improve model alignment and existing human abstractions.

5.2 Benchmarking Language Model Specificity with ML Researchers

An essential alignment task for generative language model researchers is ensuring models produce outputs at an appropriate level of abstraction [70, 162, 166, 167]. For instance, given the input “*What is Claude Monet’s profession?*”, we would prefer a model that gave a specific answer, like PAINTER, instead of an overly general answer, like WORKER. On the other hand, if the model is unsure of Monet’s exact profession, then we’d prefer that it outputs a more general answer it is confident in, like ARTIST, than a specific but possibly incorrect guess, like PHOTOGRAPHER. Currently, researchers assess language model specificity using benchmark datasets containing input prompts paired with multiple correct outputs at different levels of abstraction [63, 91, 130, 161]. However, these benchmarks limit researchers to testing a small number of possible answers across a few levels of abstraction (typically 2–4). This can result in an incomplete understanding of model accuracy since the dataset may inadvertently over penalize answers that humans consider synonymous but are not included in the labels. Moreover, since they dichotomize between a set of correct answers and all other incorrect answers, they do not provide researchers with insight into how wrong a particular mistake is (e.g., PHOTOGRAPHER is a better guess for Monet’s profession than COWBOY).

In this case study, we evaluate abstraction alignment’s ability to improve language model specificity testing through interactive analysis with experts (Section 5.2.1) and quantitative benchmarks (Section 5.2.2). We apply abstraction alignment to measure the specificity of five BERT [34], RoBERTa [88], and GPT-2 [117] language models. We evaluate each model on the S-TEST [63] specificity benchmark dataset, containing sentence prompts for masked token prediction of the subject’s occupation, location, and birthplace. Each of the prompts is labeled with a corresponding specific answer and general answer. For instance the prompt “*Lake Louise Ski Resort is located in [MASK]*” is paired with the specific answer “*Alberta*” and the general answer “*Canada*” [63]. To create fitted abstraction graphs, we map the model’s output distribution over every possible answer to words in a lexical graph. We create the human abstraction graph by mapping the S-TEST specific answers to nodes in the WordNet DAG [44, 90]. We compute edges between nodes using WordNet’s hypernym/hyponym and holonym/meronym functions, creating an abstraction graph of precise and general answers related to the task. Since WordNet is an extensive lexical graph, it contains

many concepts relevant to occupations and locations, making it a valuable proxy for human lexical knowledge on these tasks. For example, it expands occupation specificity analysis from two concepts at two levels of abstraction to over 1,500 concepts across 9 levels of abstraction.


5.2.1 Interactively Analyzing Model Specificity with NLP Experts. To study how abstraction alignment impacts researchers’ perspectives on language model specificity, we collaborate with three language model experts (Table 1). All three experts have substantial experience benchmarking and evaluating language models, with P1 and P2 specializing in specificity analysis. For these experts, specificity testing is critical for developing a comprehensive understanding of model behavior. By expanding from a single ground-truth answer to a range of acceptable answers, specificity testing provides experts with a more nuanced assessment of model accuracy. It helps them develop a mental model of the model’s behavior and estimate how it would behave on future inputs. Despite these benefits, experts acknowledge that specificity testing is currently limited to a small number of acceptable answers across a few predefined levels of abstraction, restricting their ability to thoroughly evaluate model behavior against the full range of correct answers and abstraction levels observed in real-world language use.

A key specificity alignment task for experts was analyzing the model’s preferred level of abstraction, so they could flag models whose outputs were uselessly general or misleadingly specific. Experts’ current specificity benchmarks contain a set of correct answers at various levels of abstraction, allowing them to test how often the model prefers a specific answer over a more general one. However, by propagating model confidence through the abstraction graph, abstraction alignment broadened experts’ perspectives on specificity testing. Analyzing the confidence distributions across all nine abstraction levels in the fitted abstraction graphs, users observed that small probabilities on specific answers often summed to substantial confidence in more general responses and there were many cases where they could “*elicit a different prediction by aggregating probabilities on all these very specific [answers]*” (P1). For example, in Figure 6A, experts’ traditional specificity benchmark would have penalized the model for incorrectly predicting Push Button is a PHOTOGRAPHER since it is not a synonym for the correct answer, COMPOSER. But, by propagating the probabilities, abstraction alignment revealed that the model was conceptually correct just non-specific — i.e., all of its probability was assigned to types of ARTISTS. This insight demonstrated that the model’s understanding was more nuanced than a traditional specificity benchmark could capture. For P3, viewing the results through the abstraction alignment lens was “*a much stronger claim for both specificity and categorization than just [comparing] to a [few] words.*”

Experts also found that abstraction alignment allowed them test a range of hypotheses about model specificity that are not possible with current benchmarks. During their analysis, experts often generated questions about the model’s alignment, such as whether there are dissimilar professions that the model thinks are highly related (P2). While current specificity benchmarks only consider a set of synonyms against all other incorrect options, abstraction alignment supported users in testing these alignment hypotheses by measuring the conceptual distance between model outputs. For example, to

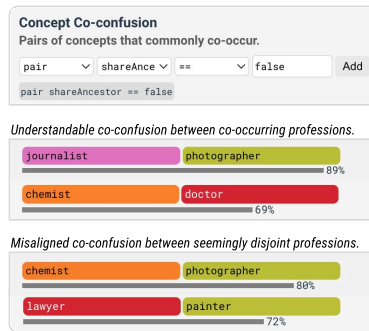
A Hidden Differences in Model Generations

Fitted abstraction graphs show that small probabilities on specific (often incorrect) answers can accumulate into high confidence in broader (often correct) outputs, indicating the model's tendency to be overly-specific.



B Misaligned Concept Relationships

Searching for highly co-confused concepts reveals that the model conflates unrelated professions.



C Concept Overreliance

The cumulative fitted abstraction graph shows that the model assigns some probability to PROFESSIONAL on every generation.

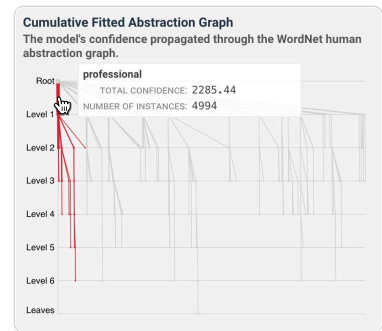


Figure 6: Abstraction alignment helps ML researchers better understand the specificity of generative language models, revealing the model's tendency to generate specific outputs at the expense of correctness (A), confuse seemingly unrelated concepts (B), and overrely on one particular concept (C).

test their hypothesis, P2 examined the concept co-confusion of occupation pairs that did not share an ancestor (Figure 6B). While they found it acceptable that the model assigned probability to co-occurring occupations, like JOURNALIST and PHOTOGRAPHER, they were surprised to see high concept co-confusion between rare pairs of professions, such as LAWYER and PAINTER. Experts worried that “these co-occurrences must [be coming] from the data distribution” (P1), suggesting more serious issues with the underlying training data that could cause “models trained on the same data to have similar co-occurrences” (P1). Relatedly, P3 identified that PROFESSIONAL was assigned some probability across all 4,994 instances (Figure 6C), hypothesizing that it was due to overuse of the word PROFESSIONAL in the dataset outside of an occupation context, such as “[person] is a professional”. Since these findings suggested dataset artifacts were impacting the model's alignment, P2 wanted to confirm that these correlations existed in the dataset, and if so, “inform people who create data [that they] should be careful about these co-occurrences, they cause hallucinations”. As a result, experts found that abstraction alignment expanded the types of specificity tests they perform, from narrow questions about a model's accuracy to broad hypotheses about the model's human-alignment.

Finally, beyond analyzing model specificity, experts hypothesized that abstraction alignment could improve model generation. As P1 described, even “when a model is not confident, it will still produce a very specific answer, but [you know its not confident because] when you sample multiple responses, you will get different answers.” In response to this phenomena, some model generation methods improve accuracy by relaxing the requirement for specificity through repeatedly sampling the model's output and identifying consistent details across the results [161]. However, P1 hypothesized that abstraction alignment could be an alternative method for improving model accuracy. Instead of selecting the model's most probable answer (which is likely overly specific and incorrect), they were interested in using the fitted abstraction graphs to select the most specific concept above a particular confidence threshold. For example, in Figure 6A, instead of generating the model's most likely

answer, MAGICIAN (which is incorrect), with abstraction alignment, they could generate a higher-level and higher-confidence answer, ARTIST (which is correct). This example highlights the versatility of abstraction alignment, showing that by viewing models through this lens, experts not only generated new hypotheses for specificity analysis but also uncovered novel strategies for model generation.

5.2.2 Quantitatively Comparing Model Specificity. While interactively exploring abstraction alignment expanded experts' qualitative analysis, it also improves traditional quantitative specificity benchmarks by generating a more diverse range of testable hypotheses. Existing specificity benchmarks are limited to comparing the model's probability across a small set of correct answers [63, 91, 130, 161]. As a result, existing specificity benchmarks, like S-TEST, only measure the model's preference between one specific and one general answer [63]. Instead, by leveraging human lexical abstractions, abstraction alignment expands the number of possible answers and represents the relationships between answers. Thus, it enables us to test a broader range of specificity questions, such as how often the model prefers any specific answer to any general answer or whether it prefers a correct answer at any level of abstraction over an incorrect but task-specific answer.

To quantify specificity, we use subgraph preference (Eq. (2)) to compare the model's preference for answers in different regions of the lexical abstraction graph (Table 2). As a baseline, we recreate Huang et al. [63]'s specificity metric by comparing the model's probabilities in the dataset-defined specific and general labels ($P(s, g)$). Next, we expand this metric to test specificity across additional words and levels of abstraction. Instead of testing one specific and one general answer, we compare all answers more specific than the specific label (specific label and its descendants) to all answers more general than the specific label (specific label's ancestors) ($P(s_{\downarrow}, s_{\uparrow})$). Finally, we extend these metrics even further, testing whether the model prefers a correct answer at any level of abstraction to an incorrect but task-related answer by comparing all answers related to the specific label to all task-related words ($P(s_{\uparrow}, t)$).

Table 2: Abstraction alignment expands existing language model specificity benchmarks. We compute the subgraph preference and accuracy at 10 (A@10) of BERT [34], RoBERTa [88], and GPT-2 [117] models on the S-TEST dataset’s occupation, location, and birthplace tasks [63]. We compare existing metrics that test model preference between a specific and general answer ($P(s, g)$) [63] to abstraction alignment metrics measuring the model’s preference for any specific answer to any general answer ($P(s_{\downarrow}, s_{\uparrow})$) and a correct answer at any level of abstraction to an incorrect answer on the same task ($P(s_{\uparrow}, t)$).

Model	Occupation				Location				Birthplace			
	A@10	$P(s, g)$	$P(s_{\downarrow}, s_{\uparrow})$	$P(s_{\uparrow}, t)$	A@10	$P(s, g)$	$P(s_{\downarrow}, s_{\uparrow})$	$P(s_{\uparrow}, t)$	A@10	$P(s, g)$	$P(s_{\downarrow}, s_{\uparrow})$	$P(s_{\uparrow}, t)$
bert-base	0.2844	0.7046	0.7901	0.0068	0.4316	0.4909	0.9591	0.2304	0.4142	0.6068	0.9992	0.2450
bert-large	0.2214	0.7176	0.8240	0.0116	0.4564	0.4236	0.9704	0.2744	0.4214	0.5652	0.9967	0.2592
roberta-base	0.2450	0.6180	0.7897	0.0751	0.3659	0.4999	0.9759	0.1790	0.2897	0.5448	1.000	0.2042
roberta-large	0.2244	0.7144	0.8238	0.0797	0.3905	0.4328	0.9785	0.2242	0.2321	0.4216	0.9989	0.2248
gpt-2	0.1610	0.5728	0.5192	0.1684	0.1702	0.4825	0.4489	0.1348	0.3327	0.5972	0.9845	0.1959

Benchmarking models with abstraction alignment reveals aspects of model behavior overlooked by prior metrics. Existing metrics indicate that language models only have a slight preference for specific answers, with most $P(s, g)$ between 40 – 70% [63]. However, by expanding to a larger set of possible answers, abstraction alignment reveals that language models actually have a strong preference for specific answers. For example, bert-large prefers a specific answer on over 80% of instances across all tasks. This result suggests that by only comparing two answers, prior metrics are too strict and do not account for variety of model preferences, whereas abstraction alignment more accurately reflects model specificity by considering a larger set of human-aligned answers.

Beyond making specificity testing more accurate, abstraction alignment also allows us to test other aspects of specificity. Using $P(s_{\uparrow}, t)$, we can test the model’s preference for a correct answer at any level of abstraction to an incorrect answer related to the task. For instance, when predicting “*Enrico Castellani is a [MASK] by profession*” we compare all answers that are direct ancestors or descendants of the specific label PAINTER to all other answers related to any other occupation in the dataset. While previously we found models prefer a specific correct answer to a general correct answer, here, we find that models often prefer a incorrect answer to any correct answer. This is not always correlated with accuracy or other specificity metrics – for instance, gpt-2 has the lowest accuracy and specificity on occupation prediction but the highest preference for correctness. By expanding the set of concepts we can analyze, abstraction alignment expands traditional benchmarks, exposing otherwise hidden aspects of model behavior.

5.3 Analyzing Medical Dataset Encodings with Healthcare Professionals

We also investigate how abstraction alignment can enhance participatory dataset analysis by identifying discrepancies between the abstractions users expect models to learn and those codified in the dataset. Machine learning models build their internal representations by learning correlations between input features and output labels in their training data [111]. However, the correlations encoded in the dataset are not always the correlations we expect the models to learn [13, 22] due to labeling inconsistencies [100], limited data diversity [131], or societal biases [18, 75, 146]. As a


result, a core task for ML developers is to audit their datasets, ensuring they reflect task expectations [69, 139, 149]. Yet, this process is often challenging, especially in applied settings where ML developers lack the domain expertise needed to deeply understand and evaluate the dataset’s content [10, 11, 99]. While participatory AI workflows address this gap by engaging domain experts in dataset curation and auditing [12, 30, 80], eliciting high-quality feedback remains difficult due to the size and complexity of ML datasets [33].

To investigate how abstraction alignment can facilitate participatory data analysis, we collaborated with four healthcare professionals to compare the medical abstractions encoded in a clinical ML dataset against global health standards. Specifically, we focus on a clinical coding task where medical coders label narrative descriptions of patient hospital stays with ICD-9 codes representing diseases the patient had (e.g., 730: BONE INFECTION) and procedures they received (e.g., 78.0: BONE GRAFT) [36]. These codes, derived from the World Health Organization’s (WHO) 9th revision of the International Classification of Disease (ICD-9) hierarchy [157], are critical for justifying costs to insurance companies, tracking epidemiological data, and reporting morbidity statistics [4]. AI-automated clinical coding is an active area of research [36, 38], with many models using the MIMIC-III dataset for training and evaluation [67, 68]. Ideally, models should learn to apply ICD-9 codes based on the WHO’s code guidelines, but discrepancies between real-world code application and prescribed guidance can arise due to coder inexperience, coding system complexity, and intentional misuse to increase insurance payout [107]. Since MIMIC-III contains real-world patient records, its code labels may deviate from the ICD-9 hierarchy’s intended use. This is a critical consideration for ML developers, as models trained on MIMIC-III risk learning and perpetuating these misaligned abstractions during deployment.

We simulate a participatory MIMIC-III dataset audit with four clinical coding experts – two professional medical coders (P4, P5) with decades of experience coding clinical notes at major hospitals and two ICD experts (P6, P7) at national health organizations (Table 1). While these experts have limited ML knowledge, they possess deep knowledge of the medical coding task and ICD-9 guidelines, representing domain experts that participatory AI workflows aim to include. Since the MIMIC-III dataset contains thousands of clinical notes and ICD-9 codes [67, 68], it would be infeasible for experts to manually inspect the dataset’s alignment by reading each clinical

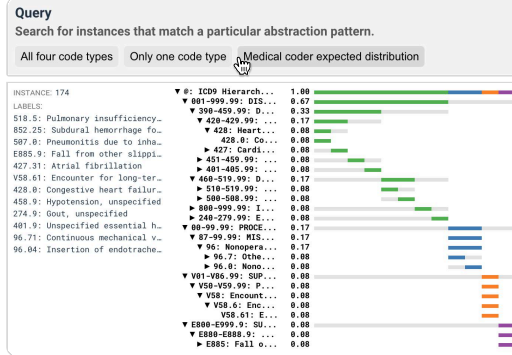
A Comparing Code Distributions

The distribution of codes in the dataset reflect the diseases seen at large hospitals, but may not transfer to specialty clinic settings.



B Querying For Expected Code Distributions

Medical coders have expectations for the distribution of codes that should be applied to each visit — e.g., 2 procedures, 20 diseases that warrant the procedure, and 1-2 codes for supplemental details.



C Identifying Unexpected Concepts

While medical coders must code clinical visits as specifically as possible, the dataset shows an overreliance on unspecified codes.

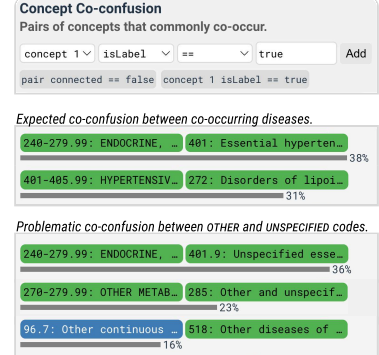


Figure 7: Abstraction alignment can also measure dataset alignment. Here, medical experts analyze a dataset of clinical notes and assigned codes. By comparing the distribution of medical code concepts (A & B) and analyzing commonly co-coded concepts (C), abstraction alignment reveals discrepancies between dataset codes and real-world coding expectations.

note and comparing the labeled codes to their expectations. However, we hypothesize that abstraction alignment may help scaffold dataset alignment by enabling structured comparisons between the dataset’s code application and the ICD-9 hierarchy, which our experts study and use in their professions.

To apply abstraction alignment in this setting, we use the ICD-9 hierarchy [94] as the human abstraction graph, containing 21,116 nodes over 7 levels of abstraction. The nodes represent the ICD-9 codes and edges abstract from low-level codes (e.g., 461.1: ACUTE FRONTAL SINUSITIS) to higher-level code groups (e.g., 460-466: ACUTE RESPIRATORY INFECTIONS). The ICD-9 graph is a relevant proxy for formal human knowledge as it represents the collective health standards our medical coders were trained on during their education. Since in this setting we are comparing the dataset’s encodings to human abstractions, we use the dataset’s labels to compute the fitted abstraction graphs. Here, the fitted abstraction graphs reflect the relationships between the codes assigned to the same medical note, where a node’s value is equal to the number of codes corresponding to that node or any of its descendants. We apply this procedure to the MIMIC-III test set, resulting in 3,372 fitted abstraction graphs.

To ensure the MIMIC-III dataset aligns with medical use, all four experts sought to confirm that the distribution of ICD-9 codes in the dataset reflected the disease and procedure frequencies they experienced in real-world hospital settings. Since clinical notes often contain verbose codes representing highly-specific diseases and procedures, it is challenging to make sense of patterns in the raw label distribution. However, abstraction alignment allowed experts to examine the code distribution at their preferred level of abstraction by aggregating specific code labels (e.g., 427.31: ATRIAL FIBRILLATION) into progressively higher-level code groupings (e.g., 427: CARDIAC DYSRHYTHMIAS and 390–459: DISEASES OF THE CIRCULATORY SYSTEM). Using the instance’s fitted abstraction graphs, experts began by inspecting the dataset’s distribution across the four top-level ICD-9 code groups (00–99.99: PROCEDURES, 001–999.99: DISEASES AND INJURIES, V01–V86.99: SUPPLEMENTARY HEALTH FACTORS, AND E800–E999.9: SUPPLEMENTARY CAUSES OF INJURY AND POISONING)

and iteratively drilled down into subcategories and individual codes. For example, by inspecting the top-level color distribution of individual clinical notes (Figure 7B), P4 confirmed that the dataset’s code distribution aligned with their expectations, “[This distribution] makes sense, because [on a single visit], you could code twenty diagnosis codes, two procedure codes, and maybe a V or E code.”

However, expanding the overall code distribution and moving down a level of abstraction to more closely inspect the disease codes caused P5 to worry that the distribution skewed towards CARDIOVASCULAR and ENDOCRINE diseases, with very few codes related to PREGNANCY AND CHILDBIRTH (Figure 7A). Since the dataset was sourced from a large hospital, it made sense to P5 that the most commonly applied codes matched common diseases, like heart disease and diabetes. However, since P6 and P7 are responsible for nationwide code usage, they were concerned that the dataset might not accurately represent code usage in specialty departments, like stand-alone OB-GYN clinics. Models trained on this dataset would experience a substantial distribution shift between the large hospital codes they saw during training and the specialty codes they would need to assign in practice. For our experts, this discrepancy signaled the need for additional data collection from specialist departments that could be used to fine-tune models for specific use cases (P5) and downstream model evaluations stratified by hospital type to assess performance across varied clinical settings (P6).

Another important alignment task for the medical experts was to assess whether the MIMIC-III dataset captured meaningful relationships between co-occurring medical conditions and procedures. In real-world practice, specific diseases and treatments often co-occur due to established medical correlations. To explore the dataset’s alignment with clinical reality, experts analyzed the filtered concept co-confusion view to test hypotheses about code pairs at different levels of abstraction (Figure 7C), such as how frequently applied codes co-occur with codes from an unrelated subgraph. Experts were not surprised to see that the dataset commonly contains labels for codes like 240–279.99: ENDOCRINE, NUTRITIONAL, AND METABOLIC DISEASES AND IMMUNITY DISORDERS and 401: ESSENTIAL

HYPERTENSION because “*patients with diabetes end up with hypertension*” (P4). However, they were concerned to see frequent co-labeling of OTHER and UNSPECIFIED codes, like 270–279.99: OTHER METABOLIC AND IMMUNITY DISEASES and 285: OTHER AND UNSPECIFIED ANEMIAS. ICD-9 often contains codes representing specific variants of a disease as well as an OTHER or UNSPECIFIED catchall code. Occasionally, applying an UNSPECIFIED code is appropriate when there is no way for the hospital to know the specific disease; however, the medical coders explained that high-quality coding often required them to request additional information from the doctor so they could apply the most accurate codes.

“[This many unspecified] codes is a no, no. Our job is to code to the highest level of specificity because the more information you have, the more likelihood the insurance company understands what’s going on and the claim gets paid.” – P5

Experts worried that the frequent occurrence of UNSPECIFIED code labels in the MIMIC-III dataset could cause models to over apply UNSPECIFIED codes when they are unwarranted or not learn distinctions between the diseases contained within the UNSPECIFIED code, leading to billing issues and statistical discrepancies.

Our medical experts’ abstraction alignment analysis of MIMIC-III reveals discrepancies between how ICD-9 codes are applied in the dataset and human expectations for disease classification. These misalignments suggest that even models that achieve high performance on the dataset may not align with medical standards and could perpetuate code misapplication, leading to inaccurate insurance billing and epidemiological statistics. However, abstraction alignment enables experts to engage in participatory dataset auditing, identifying these issues before models are trained and using their insights to guide corrective actions, like re-coding problematic patient records or reweighting the dataset to balance the frequency of over-applied codes. Notably, abstraction alignment also enabled medical experts to identify limitations in the ICD-9 abstractions themselves. In fact, the overuse of OTHER and UNSPECIFIED codes that our experts found via abstraction alignment corresponds to real-world changes the WHO made during the transition from ICD-9 to ICD-10 to increase code specificity [28, 158]. These findings suggest that beyond supporting participatory dataset analysis, abstraction alignment can also uncover opportunities to refine human-designed abstractions, ensuring they better support both clinical practice and machine learning workflows.

6 Discussion and Future Work

We consider abstraction alignment to be a methodology (i.e., an overarching strategy or conceptual foundation [31]) for comparing the alignment of model-learned and human-encoded abstractions. In this paper, we have instantiated one method for measuring abstraction alignment that compares model outputs to human abstraction graphs. This approach has proven valuable in assessing model and dataset alignment across computer vision, natural language, and medical domains. However, we expect there are many methods for measuring abstraction alignment and believe that developing these methods — tailored to different models, domains, and users — provides exciting opportunities for future work.

Our current method uses the model’s output confidence as a proxy for its internal abstractions. While this approach is model agnostic, allowing us to apply it across a range of models and flexibly extend it to dataset analysis, its reliance on model uncertainty limits our ability to measure abstraction alignment when models are confidently correct. Alternative abstraction alignment methods could overcome this limitation by extracting abstractions directly from the model’s internal representations, drawing inspiration from methods that identify a model’s internal state [58, 83] or its procedure for producing an output [40, 53, 104, 152]. New metrics could measure the alignment between these extracted model representations and existing human abstractions, revealing how a model’s abstractions change across layers and evolve during training.

Another limitation of our abstraction alignment method is its dependence on human abstraction graphs as proxies for human knowledge. Currently, abstraction alignment is limited to domains where abstraction graphs exist, such as linguistics [44, 90] and healthcare [67, 68, 157]. In some domains, abstraction graphs may not perfectly capture task semantics. For instance, while WordNet maps specific areas (e.g. PATAGONIA) to more general locations (e.g., CHILE), it is not a comprehensive location database and omits many towns (e.g., COYHAIQUE) [44, 90]. As a result, in a location prediction task, WordNet may not include every location the model outputs and abstraction alignment may not fully capture the model’s abstractions. However, we are encouraged by the extensive research on knowledge graph generation [61, 66]. Incorporating these methods to generate human abstraction graphs in new domains could enhance abstraction alignment’s applicability and effectiveness.

Further, while human abstraction graphs represent collective domain knowledge, they do not capture the diversity of knowledge that exists across individuals [138], so being abstraction-aligned does not guarantee universal alignment. For example, doctors develop distinct medical abstractions based on their medical training and clinical experiences [21]. Thus, there are many ways to be aligned — e.g., representing a governing body’s medical standards, hospital’s practices, or individual clinician’s perspective. While, currently, abstraction alignment is limited to comparing against collective and formal human knowledge, it offers an opportunity to envision personalized abstraction alignment methods, such those that codify individual knowledge or replace formal abstraction graphs with interactively specified representations. Personalized abstraction alignment could improve human-AI collaboration by ensuring that the human and model are reasoning with the same abstractions. Or, more interestingly, it could identify a user’s optimal collaborator (e.g., a model with complementary abstractions) that expands the user’s expertise and provides alternative perspectives.

Moreover, while the abstraction graph’s DAG structure is computationally valuable, it imposes a rigid view of human knowledge that may not perfectly reflect human cognition. Currently, the abstraction graph is based on Aristotelian concept theory where concepts define precise and discrete membership conditions [125]. Thus, our abstractions are unambiguous — DOGS are animals of the species CANIS LUPUS so SCHNAUZERS and WOLVES are both DOGS. However, cognitive psychology has better represented human reasoning using graded concept theory, where membership is continuous [124]. In this paradigm, a common dog like SCHNAUZER is a strong member of a DOG whereas WOLF is a weak member because, while technically

still a dog, we perceive them differently from domesticated dogs. Adopting graded concept theory may suggest a continuous measurement of abstraction alignment where the abstraction between concepts is weighted based on membership strength.

Likewise, while abstraction graphs provide a structured framework for measuring alignment, they inherently constrain the type of knowledge that can be represented, raising ethical implications about what it means to be aligned and whose knowledge we are aligning to [138]. These graphs are well-suited to domains with explicit and formalized knowledge but struggle to accommodate tacit, subjective, or contextual knowledge that is harder to generalize and abstract into discrete concepts [5, 65, 109, 116]. By privileging knowledge that is easily formalized, abstraction graphs may amplify dominant representations of knowledge, potentially marginalizing alternative ways of knowing [43, 139, 156]. This risks reinforcing existing power dynamics, encouraging the development of models that perpetuate dominant worldviews [45]. Future alignment research should critically examine and document the perspectives we align to [48] and explore more informal abstraction representations that better represent diverse forms of knowledge.

Nevertheless, our case studies demonstrate that the current instantiation of abstraction alignment helps domain experts interpret model and dataset alignment. While there are two limitations with the design of our case studies — they follow a largely exploratory protocol and only engage seven participants — these limitations map to our goal of understanding how grounding alignment analysis in abstractions can aid expert analysts. Thus, we did not conduct comparative evaluations (e.g., against alternative alignment methods), nor do our case studies help us gauge the value of our interface design choices or the usefulness of abstraction alignment in broader contexts (e.g., with less-expert analysts). Nonetheless, our study design allowed us to gather rich, open-ended feedback from experts and evaluate abstraction alignment in real-world tasks. Future work could complement our findings through large-scale comparative studies that directly contrast abstraction alignment to alternative alignment techniques across diverse participants, tasks, and domains. For instance, studies could extend our exploration of abstraction alignment in dataset analysis by applying it to a real-world participatory dataset audit and comparing users' speed and findings against traditional auditing methods. Additionally, a comparative study examining the insights users derive from their models using abstraction alignment versus established interpretability methods (e.g., saliency [92] or probes [58, 83, 84]) could help contextualize the value of abstraction alignment and reveal how it complements existing approaches.

7 Conclusion

In this paper, we study *abstraction alignment*: the agreement between a model's learned behavior and established human abstractions. While current alignment workflows require analysts to mentally compare identified model concepts against their internal representations, abstraction alignment provides structure to alignment analysis by externalizing formal human knowledge as a set of concepts and their abstraction relationships. As a result, across interpretability workflows on computer vision and natural language

tasks, abstraction alignment identifies recurring model misalignments, helping real-world domain experts build a mental model of the model's future behavior. Beyond interactive analysis, abstraction alignment extends the expressiveness of quantitative model benchmarks, revealing aspects of model behavior overlooked by prior metrics, including underestimating models' preference for specific answers. Finally, in a medical dataset analysis task with clinical experts, abstraction alignment reveals differences between the abstractions we expect models to learn and those codified in the dataset, suggesting improvements to existing human abstractions.

Acknowledgments

We thank David Bau, Been Kim, and Martin Wattenberg for their feedback, which significantly improved this work. We also thank our study participants for sharing their experiences and insights. This work was supported by NSF grant #1900991. The first author is supported by the Apple Scholars in AIML PhD Fellowship.

References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems.
- [2] Moloud Abdar, Farhad Pourpanah, Sadiq Hussain, Dana Rezazadegan, Li Liu, Mohammad Ghavamzadeh, Paul W. Fieguth, Xiaochun Cao, Abbas Khosravi, U. Rajendra Acharya, Vladimir Makarenkov, and Saeid Nahavandi. 2021. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Inf. Fusion* 76 (2021), 243–297. <https://doi.org/10.1016/j.inffus.2021.05.008>
- [3] Christopher Alexander, Sara Ishikawa, and Murray Silverstein. 1977. *A Pattern Language: Towns, Buildings, Construction*. Oxford University Press, New York.
- [4] Sherri Alexander, Therese Conner, and Teresa Slaughter. 2003. Overview of inpatient coding. *American journal of health-system pharmacy : AJHP : official journal of the American Society of Health-System Pharmacists* 60 21 Suppl 6 (2003), S11–4.
- [5] J.R. Anderson. 2013. *The Architecture of Cognition*. Taylor & Francis.
- [6] Hyemin Bang, Angie Boggust, and Arvind Satyanarayan. 2024. Explanation Alignment: Quantifying the Correctness of Model Reasoning At Scale. In *European Conference on Computer Vision (ECCV) Explainable Computer Vision (eXCV) Workshop*.
- [7] David Bau, Bolei Zhou, Aditya Khosla, Aude Oliva, and Antonio Torralba. 2017. Network Dissection: Quantifying Interpretability of Deep Visual Representations. In *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, 3319–3327. <https://doi.org/10.1109/CVPR.2017.354>
- [8] Alex Bäuerle, Ángel Alexander Cabrera, Fred Hohman, Megan Maher, David Koski, Xavier Sua, Titus Barik, and Dominik Moritz. 2022. Symphony: Composing Interactive Interfaces for Machine Learning. In *CHI Conference on Human Factors in Computing Systems*. ACM, New York, 210:1–210:14. <https://doi.org/10.1145/3491102.3502102>
- [9] Maalvika Bhat and Duri Long. 2024. Designing Interactive Explainable AI Tools for Algorithmic Literacy and Transparency. In *Designing Interactive Systems Conference (DIS)*. ACM. <https://doi.org/10.1145/3643834.3660722>
- [10] Aditya Bhattacharya, Simone Stumpf, Lucija Gosak, Gregor Stiglic, and Katrien Verbert. 2023. Lessons Learned from EXMOS User Studies: A Technical Report Summarizing Key Takeaways from User Studies Conducted to Evaluate The EXMOS Platform. *CoRR* abs/2310.02063 (2023). <https://doi.org/10.48550/ARXIV.2310.02063> arXiv:2310.02063
- [11] Aditya Bhattacharya, Simone Stumpf, Lucija Gosak, Gregor Stiglic, and Katrien Verbert. 2024. EXMOS: Explanatory Model Steering through Multifaceted Explanations and Data Configurations. In *CHI Conference on Human Factors in Computing Systems*. ACM, 314:1–314:27. <https://doi.org/10.1145/3613904.3642106>
- [12] Abeba Birhane, William Isaac, Vinodkumar Prabhakaran, Mark Diaz, Madeleine Clare Elish, Iason Gabriel, and Shakir Mohamed. 2022. Power to

- the People? Opportunities and Challenges for Participatory AI. In *Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO)*. ACM, 6:1–6:8. <https://doi.org/10.1145/3551624.3555290>
- [13] Abeba Birhane and Vinay Uday Prabhu. 2021. Large image datasets: A pyrrhic win for computer vision?. In *Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 1536–1546. <https://doi.org/10.1109/WACV48630.2021.00158>
- [14] Angie Boggust, Brandon Carter, and Arvind Satyanarayan. 2022. Embedding Comparator: Visualizing Differences in Global Structure and Local Neighborhoods via Small Multiples. In *International Conference on Intelligent User Interfaces (IUI)*. ACM, New York, 746–766. <https://doi.org/10.1145/3490099.3511122>
- [15] Angie Boggust, Benjamin Hoover, Arvind Satyanarayan, and Hendrik Strobel. 2022. Shared Interest: Measuring Human-AI Alignment to Identify Recurring Patterns in Model Behavior. In *CHI Conference on Human Factors in Computing Systems*. ACM, 10:1–10:17. <https://doi.org/10.1145/3491102.3501965>
- [16] Angie Boggust, Venkatesh Sivaraman, Yannick Assogba, Donghao Ren, Dominik Moritz, and Fred Hohman. 2024. Compress and Compare: Interactively Evaluating Efficiency and Behavior Across ML Model Compression Experiments. *IEEE Transactions on Visualization and Computer Graphics* 31, 1 (2024), 809–819. <https://doi.org/10.1109/TVCG.2024.3456371>
- [17] Angie Boggust, Harini Suresh, Hendrik Strobel, John V. Guttag, and Arvind Satyanarayan. 2023. Saliency Cards: A Framework to Characterize and Compare Saliency Methods. In *Conference on Fairness, Accountability, and Transparency (FAccT)*. ACM, 285–296. <https://doi.org/10.1145/3593013.3593997>
- [18] Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Annual Conference on Neural Information Processing Systems (NeurIPS)* 29 (2016).
- [19] Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermy, Tom Conerly, Nick Turner, Cem Anil, Carson Denison, Amanda Askell, Robert Lasenby, Yifan Wu, Shauna Kravec, Nicholas Schiefer, Tim Maxwell, Nicholas Joseph, Zac Hatfield-Dodds, Alex Tamkin, Karina Nguyen, Brayden McLean, Josiah E Burke, Tristan Hume, Shan Carter, Tom Henighan, and Christopher Olah. 2023. Towards Monosemanticity: Decomposing Language Models With Dictionary Learning. *Transformer Circuits Thread* (2023). <https://transformer-circuits.pub/2023/monosemantic-features/index.html>
- [20] Ángel Alexander Cabrera, Erica Fu, Donald Bertucci, Kenneth Holstein, Ameet Talwalkar, Jason I. Hong, and Adam Perer. 2023. Zeno: An Interactive Framework for Behavioral Evaluation of Machine Learning. In *CHI Conference on Human Factors in Computing Systems*. ACM, New York, 419:1–419:14. <https://doi.org/10.1145/3544548.3581268>
- [21] Carrie J. Cai, Samantha Winter, David Steiner, Lauren Wilcox, and Michael Terry. 2019. “Hello AI”: Uncovering the Onboarding Needs of Medical Practitioners for Human-AI Collaborative Decision-Making. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 104:1–104:24. <https://doi.org/10.1145/3359206>
- [22] Aylin Caliskan, Pimparkar Parth Ajay, Tessa Charlesworth, Robert Wolfe, and Mahzarin R. Banaji. 2022. Gender Bias in Word Embeddings: A Comprehensive Analysis of Frequency, Syntax, and Semantics. In *AAAI/ACM Conference on AI, Ethics, and Society (AI/ES)*. ACM, 156–170. <https://doi.org/10.1145/3514094.3534162>
- [23] Steve Campbell, Melanie Greenwood, Sarah Prior, Toniele Shearer, Kerrie Walkem, Sarah Young, Danielle Bywaters, and Kim Walker. 2020. Purposive sampling: complex or simple? Research case examples. *Journal of Research in Nursing* 25, 8 (2020), 652–661. <https://doi.org/10.1177/1744987120927206>
- [24] Brandon Carter, Siddhartha Jain, Jonas W Mueller, and David Gifford. 2021. Overinterpretation reveals image classification model pathologies. *Annual Conference on Neural Information Processing Systems (NeurIPS)* 34 (2021), 15395–15407.
- [25] Brandon Carter, Jonas Mueller, Siddhartha Jain, and David K. Gifford. 2019. What made you do this? Understanding black-box decisions with sufficient input subsets. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*. PMLR, 567–576.
- [26] Shan Carter, Zan Armstrong, Ludwig Schubert, Ian Johnson, and Chris Olah. 2019. Activation Atlas. *Distill* (2019). <https://doi.org/10.23915/distill.00015> <https://distill.pub/2019/activation-atlas>
- [27] Shan Carter, Zan Armstrong, Ludwig Schubert, Ian Johnson, and Chris Olah. 2019. Activation Atlas. *Distill* (2019). <https://doi.org/10.23915/distill.00015> <https://distill.pub/2019/activation-atlas>
- [28] Donna J. Cartwright. 2013. ICD-9-CM to ICD-10-CM Codes: What? Why? How? *Advances in Wound Care* 2, 10 (2013), 588–592. <https://doi.org/10.1089/wound.2013.0478>
- [29] Thomas Y. Chen, Biprateep Dey, Aishik Ghosh, Michael Kagan, Brian Nord, and Nesar Ramachandra. 2022. Interpretable Uncertainty Quantification in AI for HEP. *CoRR abs/2208.03284* (2022). <https://doi.org/10.48550/ARXIV.2208.03284> arXiv:2208.03284
- [30] Ned Cooper and Alexandra Zafiroglu. 2024. From Fitting Participation to Forging Relationships: The Art of Participatory ML. In *CHI Conference on Human Factors in Computing Systems*. ACM, 746:1–746:9. <https://doi.org/10.1145/3613904.3642775>
- [31] Michael Crotty. 1998. *The Foundations of Social Research: Meaning and Perspective in the Research Process*. SAGE Publications.
- [32] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, 248–255. <https://doi.org/10.1109/CVPR.2009.5206848>
- [33] Wesley Hanwen Deng, Bill Boyuan Guo, Alicia DeVrio, Hong Shen, Motahhare Eslami, and Kenneth Holstein. 2023. Understanding Practices, Challenges, and Opportunities for User-Engaged Algorithm Auditing in Industry Practice. In *CHI Conference on Human Factors in Computing Systems*. ACM, 377:1–377:18. <https://doi.org/10.1145/3544548.3581026>
- [34] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*. ACL, 4171–4186.
- [35] Melvil Dewey. 2011. *Dewey Decimal Classification and Relative Index* (23 ed.). OCLC Online Computer Library Center, Inc.
- [36] Hang Dong, Matúš Falis, William Whiteley, Beatrice Alex, Joshua Matterson, Shaoxiang Ji, Jiaoyan Chen, and Honghan Wu. 2022. Automated clinical coding: what, why, and where we are? *NPJ digital medicine* 5, 1 (2022), 159. <https://doi.org/10.1038/S41746-022-00705-7>
- [37] Finale Doshi-Velez and Been Kim. 2017. Towards A Rigorous Science of Interpretable Machine Learning. <https://doi.org/10.48550/arXiv.1702.08608> arXiv:1702.08608 [stat.ML]
- [38] Joakim Edin, Alexander Junge, Jakob D. Havtorn and Lasse Borgholt, Maria Maistro, Tuukka Ruotsalo, and Lars Maaloe. 2023. Automated Medical Coding on MIMIC-III and MIMIC-IV: A Critical Review and Replicability Study. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 2572–2582. <https://doi.org/10.1145/3539618.3591918>
- [39] Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec, Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, Roger Grosse, Sam McCandlish, Jared Kaplan, Dario Amodei, Martin Wattenberg, and Christopher Olah. 2022. Toy Models of Superposition. *Transformer Circuits Thread* (2022). https://transformer-circuits.pub/2022/toy_model/index.html
- [40] Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, Nova Das-Sarma, Dawn Drain, Deep Ganguli, Zac Hatfield-Dodds, Danny Hernandez, Andy Jones, Jackson Kernion, Liane Lovitt, Kamal Ndots, Dario Amodei, Tom Brown, Jack Clark, Jared Kaplan, Sam McCandlish, and Chris Olah. 2021. A Mathematical Framework for Transformer Circuits. *Transformer Circuits Thread* (2021). <https://transformer-circuits.pub/2021/framework/index.html>
- [41] Dumitru Erhan, Yoshua Bengio, Aaron Courville, and Pascal Vincent. 2009. Visualizing higher-layer features of a deep network. *University of Montreal* 1341, 3 (2009), 1.
- [42] Kevin W Eva. 2005. What every teacher needs to know about clinical reasoning. *Medical Education* 39, 1 (2005), 98–106. <https://doi.org/10.1111/j.1365-2929.2004.01972.x>
- [43] Andrew Feenberg. 2005. Critical Theory of Technology: An Overview. *Tailoring Biotechnologies* 1 (01 2005), 47–64.
- [44] Christiane Fellbaum. 1998. *WordNet: An electronic lexical database*. MIT press, Cambridge, MA, USA.
- [45] Michel Foucault. 1972–1977. *Power/Knowledge: Selected Interviews and Other Writings*. Pantheon Books, New York, NY, USA.
- [46] Craig R. Fox and Gulden Ülkümen. 2011. Distinguishing Two Dimensions of Uncertainty. In *Essays in Judgment and Decision Making*. Universitetsforlaget, Oslo, Norway. <https://doi.org/10.2139/ssrn.3695311>
- [47] Yarín Gal and Zoubin Ghahramani. 2016. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In *International Conference on Machine Learning (ICML) (JMLR Workshop and Conference Proceedings, Vol. 48)*. JMLR.org, 1050–1059.
- [48] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna M. Wallach, Hal Daumé III, and Kate Crawford. 2021. Datasheets for datasets. *Commun. ACM* 64, 12 (2021), 86–92. <https://doi.org/10.1145/3458723>
- [49] Robert Geirhos, Kristof Meding, and Felix A. Wichmann. 2020. Beyond accuracy: quantifying trial-by-trial behaviour of CNNs and humans by measuring error consistency. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*.
- [50] Amirata Ghorbani, James Wexler, and Been Kim. 2019. Automating Interpretability: Discovering and Testing Visual Concepts Learned by Neural Networks. *CoRR abs/1902.03129* (2019). <https://doi.org/10.48550/arXiv.1902.03129> arXiv:1902.03129

- [51] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. On Calibration of Modern Neural Networks. In *Proceedings of the International Conference on Machine Learning (ICML) (Proceedings of Machine Learning Research, Vol. 70)*. PMLR, 1321–1330.
- [52] Tammy Gupta, Kevin J. Shih, Saurabh Singh, and Derek Hoiem. 2017. Aligned Image-Word Representations Improve Inductive Transfer Across Vision-Language Tasks. In *International Conference on Computer Vision (ICCV)*. IEEE Computer Society, 4223–4232. <https://doi.org/10.1109/ICCV.2017.452>
- [53] Michael Hanna, Ollie Liu, and Alexandre Variengien. 2023. How does GPT-2 compute greater-than?: Interpreting mathematical abilities in a pre-trained language model. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*.
- [54] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- [55] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Identity Mappings in Deep Residual Networks. In *European Conference on Computer Vision (ECCV) (Lecture Notes in Computer Science, Vol. 9908)*. Springer, 630–645. https://doi.org/10.1007/978-3-319-46493-0_38
- [56] Evan Hernandez, Belinda Z. Li, and Jacob Andreas. 2024. Inspecting and Editing Knowledge Representations in Language Models. In *Conference on Language Modeling (CoLM)*. <https://doi.org/10.48550/ARXIV.2304.00740>
- [57] Evan Hernandez, Sarah Schwettmann, David Bau, Teona Bagashvili, Antonio Torralba, and Jacob Andreas. 2022. Natural Language Descriptions of Deep Visual Features. In *International Conference on Learning Representations (ICLR)*.
- [58] John Hewitt and Christopher D. Manning. 2019. A Structural Probe for Finding Syntax in Word Representations. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, (NAACL-HLT)*. ACL, 4129–4138. <https://doi.org/10.18653/V1/N19-1419>
- [59] Cody E. Hinchliff, Stephen A. Smith, James F. Allman, J. Gordon Burleigh, Ruchi Chaudhary, Lyndon M. Coghill, Keith A. Crandall, Jiabin Deng, Bryan T. Drew, Romina Gazis, Karl Gude, David S. Hibbett, Laura A. Katz, H. Dail Laughinghouse, Emily Jane McTavish, Peter E. Midford, Christopher L. Owen, Richard H. Ree, Jonathan A. Rees, Douglas E. Soltis, Tiffani Williams, and Karen A. Cranston. 2015. Synthesis of phylogeny and taxonomy into a comprehensive tree of life. *Proceedings of the National Academy of Sciences* 112, 41 (2015), 12764–12769. <https://doi.org/10.1073/pnas.1423041112>
- [60] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. Distilling the Knowledge in a Neural Network. *CoRR abs/1503.02531* (2015). <https://doi.org/10.48550/arXiv.1503.02531> arXiv:1503.02531
- [61] Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia d’Amato, Gerard de Melo, Claudio Gutierrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, Axel-Cyrille Ngonga Ngomo, Axel Polleres, Sabbir M. Rashid, Anisa Rula, Lukas Schmelzeisen, Juan Sequeda, Steffen Staab, and Antoine Zimmermann. 2021. *Knowledge Graphs*. Morgan & Claypool Publishers. <https://doi.org/10.2200/S01125ED1V01Y202109DSK022>
- [62] Fred Hohman, Chaogun Wang, Jimmook Lee, Jochen Görtler, Dominik Moritz, Jeffrey P. Bigham, Zhile Ren, Cecile Forest, Qi Shan, and Xiaoyi Zhang. 2024. Talaria: Interactively Optimizing Machine Learning Models for Efficient Inference. In *CHI Conference on Human Factors in Computing Systems*. ACM, New York, 648:1–648:19. <https://doi.org/10.1145/3613904.3642628>
- [63] Jie Huang, Kevin Chen-Chuan Chang, Jinjun Xiong, and Wen-Mei Hwu. 2023. Can Language Models Be Specific? How?. In *Findings of the Association for Computational Linguistics*. ACL, 716–727.
- [64] Eyke Hüllermeier and Willem Waegeman. 2021. Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods. *Machine Learning* 110, 3 (2021), 457–506. <https://doi.org/10.1007/S10994-021-05946-3>
- [65] Ioana Hulpus, Jonathan Kobbe, Heiner Stuckenschmidt, and Graeme Hirst. 2020. Knowledge Graphs meet Moral Values. In *Joint Conference on Lexical and Computational Semantics*. ACL, 71–80.
- [66] Shaoxiong Ji, Shirui Pan, E. Cambria, Pekka Marttinen, and Philip S. Yu. 2022. A Survey on Knowledge Graphs: Representation, Acquisition, and Applications. *IEEE Transactions on Neural Networks and Learning Systems* 33 (2022), 494–514. <https://doi.org/10.1109/TNNLS.2021.3070843>
- [67] Alistair Johnson, Tom Pollard, and Roger Mark. 2016. MIMIC-III Clinical Database (version 1.4). <https://doi.org/10.13026/C2XW26>
- [68] Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. 2016. MIMIC-III, a freely accessible critical care database. *Scientific data* 3, 1 (2016), 1–9. <https://doi.org/10.1038/sdata.2016.35>
- [69] Charles Jones, Daniel C. Castro, Fabio De Sousa Ribeiro, Ozan Oktay, Melissa D. McCradden, and Ben Glocker. 2023. No Fair Lunch: A Causal Perspective on Dataset Bias in Machine Learning for Medical Imaging. *CoRR abs/2307.16526* (2023). <https://doi.org/10.48550/ARXIV.2307.16526> arXiv:2307.16526
- [70] Ehsan Kamaloo, Nouha Dziri, Charles L. A. Clarke, and Davood Rafiei. 2023. Evaluating Open-Domain Question Answering in the Era of Large Language Models. In *Proceedings of the Annual Meeting of the Association for Computational (ACL)*. ACL, 5591–5606. <https://doi.org/10.18653/V1/2023.ACL-LONG.307>
- [71] Andrei Kaphishnikov, Tolga Bolukbasi, Fernanda B. Viégas, and Michael Terry. 2019. XRAI: Better Attributions Through Regions. In *IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, 4947–4956. <https://doi.org/10.1109/ICCV.2019.00505>
- [72] Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viégas, and Rory Sayres. 2018. Interpretability Beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (TCAV). In *International Conference on Machine Learning (ICML)*. PMLR, 2668–2677.
- [73] Been Kim, Martin Wattenberg, Justin Gilmer, Carrie J. Cai, James Wexler, Fernanda B. Viégas, and Rory Sayres. 2018. Interpretability Beyond Feature Attribution: Quantitative Testing with Concept Activation Vectors (TCAV). In *International Conference on Machine Learning (ICML) (Proceedings of Machine Learning Research, Vol. 80)*. PMLR, 2673–2682.
- [74] Armen Der Kiureghian and Ove Ditlevsen. 2009. Aleatory or epistemic? Does it matter? *Structural Safety* 31, 2 (2009), 105–112. <https://doi.org/10.1016/j.strusafe.2008.06.020> Risk Acceptance and Risk Communication.
- [75] Lauren F. Klein and Catherine D’Ignazio. 2024. Data Feminism for AI. In *Conference on Fairness, Accountability, and Transparency (FAccT)*. ACM, 100–112. <https://doi.org/10.1145/3630106.3658543>
- [76] Pang Wei Koh and Percy Liang. 2017. Understanding Black-box Predictions via Influence Functions. In *Proceedings of the International Conference on Machine Learning (ICML) (Proceedings of Machine Learning Research, Vol. 70)*. PMLR, 1885–1894.
- [77] Alex Krizhevsky, Geoffrey Hinton, et al. 2009. Learning multiple layers of features from tiny images. (2009).
- [78] Ananya Kumar, Percy Liang, and Tengyu Ma. 2019. Verified Uncertainty Calibration. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*. 3787–3798.
- [79] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. 2017. Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*. 6402–6413.
- [80] Michelle S Lam, Mitchell L Gordon, Danaë Metaxa, Jeffrey T Hancock, James A Landay, and Michael S Bernstein. 2022. End-user audits: A system empowering communities to lead large-scale investigations of harmful algorithmic behavior. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–34. <https://doi.org/10.1145/3555625>
- [81] Fred Lambert. [n. d.]. Understanding the fatal Tesla accident on Autopilot and the NHTSA probe. *electrek* ([n. d.]). <https://electrek.co/2016/07/01/understanding-fatal-tesla-accident-autopilot-nhtsa-probe/>
- [82] Thomas A. Langlois, H. Charles Zhao, Erin Grant, Ishita Dasgupta, Thomas L. Griffiths, and Nori Jacoby. 2021. Passive attention in artificial neural networks predicts human visual selectivity. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*. 27094–27106.
- [83] Belinda Z. Li, Maxwell I. Nye, and Jacob Andreas. 2021. Implicit Representations of Meaning in Neural Language Models. In *Annual Meeting of the Association for Computational (ACL)*. ACL, 1813–1827. <https://doi.org/10.18653/V1/2021.ACL-LONG.143>
- [84] Kenneth Li, Aspen K. Hopkins, David Bau, Fernanda B. Viégas, Hanspeter Pfister, and Martin Wattenberg. 2023. Emergent World Representations: Exploring a Sequence Model Trained on a Synthetic Task. In *International Conference on Learning Representations (ICLR)*. OpenReview.net.
- [85] Johnny Lin. 2023. Neuronpedia: Interactive Reference and Tooling for Analyzing Neural Networks. <https://www.neuronpedia.org>
- [86] Carl Linnaeus. 1758. *Systema Naturae per Regna Tria Naturae, Secundum Classes, Ordines, Genera, Species, cum Characteribus, Differentiis, Synonymis, Locis*. Laurentius Salvius. <https://doi.org/10.5962/bhl.title.156783>
- [87] Barbara Liskov and John V. Guttag. 1986. *Abstraction and Specification in Program Development* (2 ed.). Vol. 20. MIT press Cambridge.
- [88] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *CoRR abs/1907.11692* (2019). <https://doi.org/10.48550/ARXIV.1907.11692> arXiv:1907.11692
- [89] Jack Merullo, Carsten Eickhoff, and Ellie Pavlick. 2023. Circuit component reuse across tasks in transformer language models. *arXiv preprint arXiv:2310.08744* (2023).
- [90] George A. Miller. 1995. WordNet: A Lexical Database for English. *Commun. ACM* 38, 11 (1995), 39–41. <https://doi.org/10.1145/219717.219748>
- [91] Sewon Min, Julian Michael, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2020. AmbigQA: Answering Ambiguous Open-domain Questions. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*. ACL, 5783–5797. <https://doi.org/10.18653/V1/2020.EMNLP-MAIN.466>
- [92] Christoph Molnar. 2022. *Interpretable Machine Learning* (2 ed.). <https://christophm.github.io/interpretable-ml-book>
- [93] Ramaravind Kommiya Mothilal, Amit Sharma, and Chenhao Tan. 2020. Explaining machine learning classifiers through diverse counterfactual explanations. In *Conference on Fairness, Accountability, and Transparency (FAT*)*. ACM, 607–617. <https://doi.org/10.1145/3351095.3372850>

- [94] James Mullenbach, Sarah Wiegrefe, Jon Duke, Jimeng Sun, and Jacob Eisenstein. 2018. Explainable Prediction of Medical Codes from Clinical Text. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*. ACL, 1101–1111.
- [95] Mark A Musen. 1992. Dimensions of knowledge sharing and reuse. *Computers and biomedical research* 25, 5 (1992), 435–467.
- [96] Lukas Muttenthaler, Jonas Dippel, Lorenz Linhardt, Robert A. Vandermeulen, and Simon Kornblith. 2023. Human alignment of neural network representations. In *International Conference on Learning Representations (ICLR)*. OpenReview.net.
- [97] Lukas Muttenthaler, Klaus Greff, Frieda Born, Bernhard Spitzer, Simon Kornblith, Michael C. Mozer, Klaus-Robert Müller, Thomas Unterthiner, and Andrew K. Lampinen. 2024. Aligning Machine and Human Visual Representations across Abstraction Levels. *CoRR abs/2409.06509* (2024). <https://doi.org/10.48550/ARXIV.2409.06509> arXiv:2409.06509
- [98] Neel Nanda. 2024. Actually, Othello-GPT Has A Linear Emergent World Representation. *Transformer Circuits Thread* (2024).
- [99] Zabir Al Nazi and Wei Peng. 2024. Large Language Models in Healthcare and Medical Domain: A Review. *Informatics* 11, 3 (2024), 57. <https://doi.org/10.3390/INFORMATIC11030057>
- [100] Curtis G. Northcutt, Anish Athalye, and Jonas Mueller. 2021. Pervasive Label Errors in Test Sets Destabilize Machine Learning Benchmarks. In *Neural Information Processing Systems (NeurIPS) Track on Datasets and Benchmarks*.
- [101] N. Noy and Deborah McGuinness. 2001. Ontology Development 101: A Guide to Creating Your First Ontology. *Knowledge Systems Laboratory* 32 (01 2001).
- [102] Tuomas P. Oikarinen and Tsui-Wei Weng. 2023. CLIP-Dissect: Automatic Description of Neuron Representations in Deep Vision Networks. In *International Conference on Learning Representations (ICLR)*.
- [103] Kerem Oktar, Ilija Sucholutsky, Tania Lombrozo, and Thomas L. Griffiths. 2023. Dimensions of Disagreement: Unpacking Divergence and Misalignment in Cognitive Science and Artificial Intelligence. *CoRR abs/2310.12994* (2023). <https://doi.org/10.48550/ARXIV.2310.12994> arXiv:2310.12994
- [104] Chris Olah, Nick Cammarata, Ludwig Schubert, Gabriel Goh, Michael Petrov, and Shan Carter. 2020. Zoom In: An Introduction to Circuits. *Distill* (2020). <https://doi.org/10.23915/distill.00024.001> <https://distill.pub/2020/circuits/zoom-in>.
- [105] Chris Olah, Alexander Mordvintsev, and Ludwig Schubert. 2017. Feature Visualization. *Distill* (2017). <https://doi.org/10.23915/distill.00007> <https://distill.pub/2017/feature-visualization>.
- [106] Chris Olah, Arvind Satyanarayan, Ian Johnson, Shan Carter, Ludwig Schubert, Katherine Ye, and Alexander Mordvintsev. 2018. The Building Blocks of Interpretability. *Distill* (2018). <https://doi.org/10.23915/distill.00010> <https://distill.pub/2018/building-blocks>.
- [107] Kimberly O'Malley, Karon F. Cook, Matt D. Price, Kimberly Raiford Wildes, John F. Hurdle, and Carol M. Ashton. 2005. Measuring diagnoses: ICD code accuracy. *Health services research* 40 5 Pt 2 (2005), 1620–39.
- [108] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*.
- [109] Xiny Pan, Daniel Hernández, Philipp Seifer, Ralf Lämmel, and Steffen Staab. 2024. eSPARQL: Representing and Reconciling Agnostic and Atheistic Beliefs in RDF-star Knowledge Graphs. *CoRR abs/2407.21483* (2024). <https://doi.org/10.48550/ARXIV.2407.21483>
- [110] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Z. Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Annual Conference on Neural Information Processing Systems (NeurIPS)*. 8024–8035.
- [111] Amandalynne Paullada, Inioluwa Deborah Raji, Emily M. Bender, Emily Denton, and Alex Hanna. 2021. Data and its (dis)contents: A survey of dataset development and use in machine learning research. *Patterns* 2, 11 (2021), 100336. <https://doi.org/10.1016/J.PATTERN.2021.100336>
- [112] Fabio Petroni, Patrick S. H. Lewis, Aleksandra Piktus, Tim Rocktäschel, Yuxiang Wu, Alexander H. Miller, and Sebastian Riedel. 2020. How Context Affects Language Models' Factual Predictions. In *Conference on Automated Knowledge Base Construction (AKBC)*. <https://doi.org/10.24432/C5201W>
- [113] Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick S. H. Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander H. Miller. 2019. Language Models as Knowledge Bases?. In *Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. ACL, 2463–2473. <https://doi.org/10.18653/V1/D19-1250>
- [114] Vitali Petsiuk, Abir Das, and Kate Saenko. 2018. RISE: Randomized Input Sampling for Explanation of Black-box Models. In *British Machine Vision Conference (BMVC)*. BMVA Press, 151.
- [115] Mary Phuong and Christoph Lampert. 2019. Towards Understanding Knowledge Distillation. In *International Conference on Machine Learning (ICML) (Proceedings of Machine Learning Research, Vol. 97)*. PMLR, 5142–5151.
- [116] Michael Polanyi. 1966. *The Tacit Dimension*. Routledge & Kegan Paul, London.
- [117] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language Models are Unsupervised Multitask Learners. *OpenAI* (2019).
- [118] Arun Rai. 2020. Explainable AI: From black box to glass box. *Journal of the Academy of Marketing Science* 48, 1 (2020), 137–141.
- [119] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical Text-Conditional Image Generation with CLIP Latents. *CoRR abs/2204.06125* (2022). <https://doi.org/10.48550/ARXIV.2204.06125> arXiv:2204.06125
- [120] Sunayana Rane, Polyphony J. Bruna, Ilija Sucholutsky, Christopher T. Kello, and Thomas L. Griffiths. 2024. Concept Alignment. *CoRR abs/2401.08672* (2024). <https://doi.org/10.48550/ARXIV.2401.08672> arXiv:2401.08672
- [121] Joseph Redmon and Ali Farhadi. 2017. YOLO9000: Better, Faster, Stronger. In *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 6517–6525. <https://doi.org/10.1109/CVPR.2017.690>
- [122] Samantha Robertson, Zijie J Wang, Dominik Moritz, Mary Beth Kery, and Fred Hohman. 2023. Angler: Helping Machine Translation Practitioners Prioritize Model Improvements. In *CHI Conference on Human Factors in Computing Systems*. ACM, New York, 832:1–832:20. <https://doi.org/10.1145/3544548.3580790>
- [123] Natália Vila Nova Rodrigues, Luis Raul Abramo, and Nina Sumiko Tomita Hirata. 2021. The information of attribute uncertainties: what convolutional neural networks can learn about errors in input data. *Machine Learning: Science and Technology* 4 (2021).
- [124] Eleanor Rosch. 2002. *Principles of categorization*. MIT Press, Cambridge, MA, USA. 251–270 pages.
- [125] Eleanor H. Rosch. 2011. "Slow Lettuce": Categories, Concepts, Fuzzy Sets, and Logical Deduction. In *Concepts and Fuzzy Logic*. The MIT Press. <https://doi.org/10.7551/mitpress/8842.003.0006>
- [126] Hendrik Schuff, Alon Jacovi, Heike Adel, Yoav Goldberg, and Ngoc Thang Vu. 2022. Human Interpretation of Saliency-based Explanation Over Text. In *Conference on Fairness, Accountability, and Transparency (FACT)*. ACM, 611–636. <https://doi.org/10.1145/3531146.3533127>
- [127] Lisa Schut, Nenad Tomasev, Tom McGrath, Demis Hassabis, Ulrich Paquet, and Been Kim. 2023. Bridging the Human-AI Knowledge Gap: Concept Discovery and Transfer in AlphaZero. *CoRR abs/2310.16410* (2023). <https://doi.org/10.48550/ARXIV.2310.16410>
- [128] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. In *International Conference on Computer Vision (ICCV)*. IEEE, 618–626.
- [129] Claude Elwood Shannon. 1948. A mathematical theory of communication. *The Bell System Technical Journal* 27, 3 (1948), 379–423. <https://doi.org/10.1002/J.1538-7305.1948.TB01338.X>
- [130] Chenglei Si, Chen Zhao, and Jordan L. Boyd-Graber. 2021. What's in a Name? Answer Equivalence For Open-Domain Question Answering. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*. ACL, 9623–9629. <https://doi.org/10.18653/V1/2021.EMNLP-MAIN.757>
- [131] Zoe De Simone, Angie W. Boggust, Arvind Satyanarayan, and Ashia Wilson. 2023. What is a Fair Diffusion Model? Designing Generative Text-To-Image Models to Incorporate Various Worldviews. *CoRR abs/2309.09944* (2023). <https://doi.org/10.48550/ARXIV.2309.09944> arXiv:2309.09944
- [132] Venkatesh Sivaraman, Yiwei Wu, and Adam Perer. 2022. Emblaze: Illuminating Machine Learning Representations through Interactive Comparison of Embedding Spaces. In *International Conference on Intelligent User Interfaces (IUI)*. ACM, New York, 418–432. <https://doi.org/10.1145/3490099.3511137>
- [133] Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda B. Viégas, and Martin Wattenberg. 2017. SmoothGrad: removing noise by adding noise. <https://doi.org/10.48550/arXiv.1706.03825> arXiv:1706.03825
- [134] Jan-Tobias Sohns, Christoph Garth, and Heike Leitte. 2023. Decision Boundary Visualization for Counterfactual Reasoning. *Computer Graphics Forum* 42, 1 (2023), 7–20. <https://doi.org/10.1111/CGF.14650>
- [135] Hendrik Strobelt, Sebastian Gehrmann, Michael Behrisch, Adam Perer, Hanspeter Pfister, and Alexander M Rush. 2018. Seq2seq-vis: A visual debugging tool for sequence-to-sequence models. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (2018), 353–363. <https://doi.org/10.1109/TVCG.2018.2865044>
- [136] Ilija Sucholutsky, Lukas Muttenthaler, Adrian Weller, Andi Peng, Andreea Bobu, Been Kim, Bradley C. Love, Erin Grant, Jascha Achterberg, Joshua B. Tenenbaum, Katherine M. Collins, Katherine L. Hermann, Kerem Oktar, Klaus Greff, Martin N. Hebart, Nori Jacoby, Qiuyu Zhang, Raja Marjhe, Robert Geirhos, Sherol Chen, Simon Kornblith, Sunayana Rane, Talia Konkle, Thomas P. O'Connell, Thomas Unterthiner, Andrew K. Lampinen, Klaus-Robert Müller, Mariya Toneva, and Thomas L. Griffiths. 2023. Getting aligned on representational alignment. *CoRR abs/2310.13018* (2023). <https://doi.org/10.48550/ARXIV.2310.13018>

- arXiv:2310.13018
- [137] York Sure, Michael Erdmann, Jürgen Angele, Steffen Staab, Rudi Studer, and Dirk Wenke. 2002. OntoEdit: Collaborative Ontology Development for the Semantic Web. In *International Semantic Web Conference (ISWC) (Lecture Notes in Computer Science, Vol. 2342)*. Springer, 221–235. https://doi.org/10.1007/3-540-48005-6_18
- [138] Harini Suresh, Steven R. Gomez, Kevin K. Nam, and Arvind Satyanarayan. 2021. Beyond Expertise and Roles: A Framework to Characterize the Stakeholders of Interpretable Machine Learning and their Needs. In *CHI Conference on Human Factors in Computing Systems*. ACM, 74:1–74:16. <https://doi.org/10.1145/3411764.3445088>
- [139] Harini Suresh and John V. Guttag. 2021. A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle. In *Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO)*. ACM, 17:1–17:9. <https://doi.org/10.1145/3465416.3483305>
- [140] Harini Suresh, Divya Shanmugam, Tiffany Chen, Annie G Bryan, Alexander D'Amour, John Guttag, and Arvind Satyanarayan. 2023. Kaleidoscope: Semantically-grounded, context-specific ML model evaluation. In *CHI Conference on Human Factors in Computing Systems*. ACM, New York, 775:1–775:13. <https://doi.org/10.1145/3544548.3581482>
- [141] Ilya Sutskever, James Martens, George E. Dahl, and Geoffrey E. Hinton. 2013. On the importance of initialization and momentum in deep learning. In *International Conference on Machine Learning (ICML) (JMLR Workshop and Conference Proceedings, Vol. 28)*. JMLR.org, 1139–1147.
- [142] Adly Templeton, Tom Conerly, Jonathan Marcus, Jack Lindsey, Trenton Bricken, Brian Chen, Adam Pearce, Craig Citro, Emmanuel Ameisen, Andy Jones, Hoagy Cunningham, Nicholas L. Turner, Callum McDougall, Monte MacDiarmid, C. Daniel Freeman, Theodore R. Sumers, Edward Rees, Joshua Batson, Adam Jermyn, Shan Carter, Chris Olah, and Tom Henighan. 2024. Scaling Monosemanticity: Extracting Interpretable Features from Claude 3 Sonnet. *Transformer Circuits Thread* (2024). <https://transformer-circuits.pub/2024/scaling-monosemanticity/index.html>
- [143] Michael Terry, Chinmay Kulkarni, Martin Wattenberg, Lucas Dixon, and Meredith Ringel Morris. 2023. AI Alignment in the Design of Interactive AI: Specification Alignment, Process Alignment, and Evaluation Support. *CoRR abs/2311.00710* (2023). <https://doi.org/10.48550/ARXIV.2311.00710> arXiv:2311.00710
- [144] Tesla. [n. d.]. A Tragic Loss. https://www.tesla.com/es_mx/blog/tragic-loss. Accessed: 2024-06-16.
- [145] Yonglong Tian, Dilip Krishnan, and Phillip Isola. 2019. Contrastive Representation Distillation. *CoRR abs/1910.10699* (2019). <https://doi.org/10.48550/arXiv.1910.10699> arXiv:1910.10699
- [146] Antonio Torralba and Alexei A. Efros. 2011. Unbiased look at dataset bias. In *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 1521–1528. <https://doi.org/10.1109/CVPR.2011.5995347>
- [147] Tania Tudorache, Natalya Fridman Noy, Samson W. Tu, and Mark A. Musen. 2008. Supporting Collaborative Ontology Development in Protégé. In *International Conference on The Semantic Web (ISWC) (Lecture Notes in Computer Science, Vol. 5318)*. Springer, 17–32. https://doi.org/10.1007/978-3-540-88564-1_2
- [148] Edward Tufte. 2006. Beautiful Evidence.
- [149] Julie Vaughn, Avital Baral, Mayukha Vadari, and William Boag. 2020. Dataset Bias in Diagnostic AI systems: Guidelines for Dataset Collection and Usage. *ACM Conference in Health, Inference, and Learning, Workshop* (2020).
- [150] Bret Victor. 2011. Up and Down the Ladder of Abstraction; A Systematic Approach to Interactive Visualization. <https://worrydream.com/LadderOfAbstraction/>.
- [151] Hanjing Wang and Qiang Ji. 2024. Epistemic Uncertainty Quantification for Pre-trained Neural Networks. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2024), 11052–11061. <https://doi.org/10.1109/CVPR52733.2024.01051>
- [152] Kevin Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. 2022. Interpretability in the wild: a circuit for indirect object identification in gpt-2 small. *arXiv preprint arXiv:2211.00593* (2022).
- [153] Martin Wattenberg and Fernanda B. Viégas. 2024. Relational Composition in Neural Networks: A Survey and Call to Action. *CoRR abs/2407.14662* (2024). <https://doi.org/10.48550/ARXIV.2407.14662> arXiv:2407.14662
- [154] Martin Wattenberg, Fernanda Viégas, and Ian Johnson. 2016. How to Use t-SNE Effectively. *Distill* (2016). <https://doi.org/10.23915/distill.00002>
- [155] James Wexler, Mahima Pushkarna, Tolga Bolukbasi, Martin Wattenberg, Fernanda B. Viégas, and Jimbo Wilson. 2020. The What-If Tool: Interactive Probing of Machine Learning Models. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 56–65. <https://doi.org/10.1109/TVCG.2019.2934619>
- [156] Langdon Winner. 1980. Do Artifacts Have Politics? *Daedalus* 109, 1 (1980), 121–136.
- [157] World Health Organization. 1978. *International Classification of Diseases, Ninth Revision (ICD-9)*. World Health Organization, Geneva, Switzerland.
- [158] World Health Organization. 2022. *International Classification of Diseases, 10th Revision*. World Health Organization, Geneva, Switzerland.
- [159] Tongshuang Wu, Marco Tulio Ribeiro, Jeffrey Heer, and Daniel S Weld. 2019. Errudite: Scalable, Reproducible, and Testable Error Analysis. In *Annual Meeting of the Association for Computational Linguistics (ACL)*. ACL, Stroudsburg, PA, 747–763. <https://doi.org/10.18653/V1/P19-1073>
- [160] Eiling Yee. 2019. Abstraction and concepts: when, how, where, what and why? *Language, Cognition and Neuroscience* 34, 10 (2019), 1257–1265. <https://doi.org/10.1080/23273798.2019.1660797>
- [161] Gal Yona, Roei Aharoni, and Mor Geva. 2024. Narrowing the Knowledge Evaluation Gap: Open-Domain Question Answering with Multi-Granularity Answers. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*. ACL, 6737–6751. <https://doi.org/10.18653/V1/2024.ACL-LONG.365>
- [162] Hiyori Yoshikawa and Naoaki Okazaki. 2023. Selective-LAMA: Selective Prediction for Confidence-Aware Evaluation of Language Models. In *Findings of the Association for Computational Linguistics*. ACL, 1972–1983. <https://doi.org/10.18653/V1/2023.FINDINGS-EACL.150>
- [163] Bianca Zadrozny and Charles Elkan. 2001. Obtaining calibrated probability estimates from decision trees and naive Bayesian classifiers. In *International Conference on Machine Learning (ICML)*. Morgan Kaufmann, 609–616.
- [164] Xianlong Zeng, Fanghao Song, Zhongen Li, Krerkkiat Chusap, and Chang Liu. 2021. Human-in-the-Loop Model Explanation via Verbatim Boundary Identification in Generated Neighborhoods. In *Machine Learning and Knowledge Extraction International Cross-Domain Conference (CD-MAKE) (Lecture Notes in Computer Science, Vol. 12844)*. Springer, 309–327. https://doi.org/10.1007/978-3-030-84060-0_20
- [165] Chrysoula Zerva, Taisiya Glushkova, Ricardo Rei, and André F. T. Martins. 2022. Disentangling Uncertainty in Machine Translation Evaluation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*. ACL, 8622–8641. <https://doi.org/10.18653/V1/2022.EMNLP-MAIN.591>
- [166] Muru Zhang, Ofir Press, William Merrill, Alisa Liu, and Noah A. Smith. 2024. How Language Model Hallucinations Can Snowball. In *International Conference on Machine Learning (ICML)*.
- [167] Shen Zheng, Jie Huang, and Kevin Chen-Chuan Chang. 2023. Why Does ChatGPT Fall Short in Providing Truthful Answers? *CoRR abs/2304.10513* (2023). <https://doi.org/10.48550/ARXIV.2304.10513> arXiv:2304.10513

A Appendix

A.1 Case Study Setup

Here we describe our case study setup (Section 5). We provide code (<https://github.com/mitvis/abstraction-alignment>) and an interactive interface (<https://vis.mit.edu/abstraction-alignment/>) to explore the results.

A.1.1 CIFAR-100. In Section 5.1, we use abstraction alignment to interpret a CIFAR-100 image classification model. We train a PyTorch [110] ResNet20 model [54] on the CIFAR-100 training set [77] for 200 epochs with a batch size of 128. We apply random crop and horizontal flip augmentations to the images following He et al. [55]. We use cross-entropy loss optimized via stochastic gradient descent and Nesterov momentum [141] (momentum = 0.9; weight decay = $5e - 4$). We use a learning rate of 0.1 and reduce it at epoch 60, 120, and 160 using gamma of 0.2. The trained model achieves 67.7% accuracy on the CIFAR-100 test set.

To apply abstraction alignment we use the CIFAR-100 class/superclass structure [77] to form the human abstraction graph. The graph contains 121 nodes across 3 levels – 100 class nodes, 20 superclass nodes, and a root node. We create a fitted abstraction graph for every dataset instance in the CIFAR-100 test set. To do so, we compute the model’s output probability for each image by applying a softmax to the model’s output logits across all 100 classes. Nodes that correspond to a CIFAR-100 class are assigned a value equal to the model’s output probability for that class. All other nodes’ values are the sum of their reachable leaves’ values. For instance, TULIP’s value is the model’s output probability that the image is a tulip, whereas FLOWER’s value is the sum of the model’s output probability for ORCHID, ROSE, TULIP, SUNFLOWER, and POPPY.

A.1.2 Language Models and WordNet. In Section 5.2, we apply abstraction alignment to benchmark the specificity of language models. Following the benchmarking procedure in Huang et al. [63], we compare pretrained bert-base [34], bert-large [34], roberta-base [88], roberta-large [88], and gpt-2 [117] models from the LAMA benchmark [112, 113]. We test each model on the occupation, location, and birthplace tasks from the S-TEST dataset [63]. Each S-TEST dataset instance is a text input paired with one specific and one general label. For each model, we compute its top-10 accuracy, measured as the proportion of instances where the specific label was in the model’s top 10 predicted tokens.

To measure abstraction alignment, we create a human abstraction graph for each of the occupation, location, and birthplace tasks. For a task, we map each of its specific labels to its corresponding node (i.e., synset) in WordNet [44, 90]. We do this process by searching for the specific answer label in the NLTK WordNet corpus¹. If there are multiple WordNet nodes that hit for a given search, we select the most appropriate node by manually inspecting their WordNet definitions. Then, we expand the graph by including all direct ancestors and descendants of any specific label nodes. The result is a human abstraction graph containing all the vocabulary words related to any of the data instances’ specific labels.

To qualitatively explore abstraction alignment in Section 5.2.1, we create fitted abstraction graphs for every model decision on the

occupation prediction task. First, we compute the model’s output probability across every word in its vocabulary for every data instance. Then, for each data instance, we assign the model’s output probabilities to their corresponding nodes in the human abstraction graph. Finally, we propagate the values, such that each node’s value is the sum of its value and its children’s values.

To quantitatively benchmark the models in Section 5.2.2, we use the fitted abstractions to compute three specificity metrics, using subgraph preference (Eq. (2)). Since, in this case, the model outputs map to concepts at many different levels of abstraction, we do not propagate the values through the fitted abstraction. Instead, we assign each node a value corresponding to the model’s probability of outputting that word. First, we replicate the specificity testing metric from Huang et al. [63] (originally called p_r). We compute it as $P(s, g)$, where s is the single-node graph containing the specific label and g is the single-node graph containing the general label. Next, we compute $P(s_{\downarrow}, s_{\uparrow})$ to compare all words more specific than the specific label s_{\downarrow} (specific label and its descendants) to all words at a higher level of abstraction than the specific label s_{\uparrow} (specific label’s ancestors). Finally, we compute $P(s_{\uparrow}, t)$ to compare ancestors and descendants of the specific label s_{\uparrow} to any other word in the task DAG t .

A.1.3 MIMIC-III Medical Dataset with ICD-9 Codes. In Section 5.3, we apply abstraction alignment to analyze the abstractions in the MIMIC-III dataset [67, 68]. The dataset contains textual medical notes paired with a set of ICD-9 code labels. We use the ICD-9 medical hierarchy as the human abstraction graph [157]. We pair the dataset’s ICD-9 code labels with their corresponding code in the ICD-9 abstraction graph. In this task, non-leaf nodes are codable – e.g., both 282.6: SICKLE-CELL ANEMIA and its direct parent 282: HEREDITARY HEMOLYTIC ANEMIAS can be applied to the same medical note. To compute the fitted abstractions graphs, we set the code node’s value equal to one if the code was labeled on that instance and zero otherwise. Then we propagate scores following Figure 2, setting each node’s value equal to its value plus the summed values of its children. As a result the value of a node is equivalent to the number of times it or one of its children was labeled on the medical note.

A.2 Compute Resources and Efficiency

Time to compute abstraction alignment depends on the model, dataset, and human abstraction graph. Extracting the model outputs/dataset labels, setting up the human abstraction graph, creating the fitted abstraction graphs, and measuring the abstraction alignment metrics take on the order of 10 minutes for the image model case study (Section 5.1), 5 minutes for the language model case study (Section 5.2), and 30 minutes for the medical dataset analysis case study (Section 5.3). We train and evaluate our models on 1 NVIDIA V100 GPU with 1TB of memory.

A.3 Interface Implementation Details

The abstraction alignment interface (<https://vis.mit.edu/abstraction-alignment/>) uses Svelte to build responsive visualizations and the HTML5 Canvas to render performance-intensive charts.

¹<https://www.nltk.org/howto/wordnet.html>