

DIGGING DEEPER: LEARNING MULTI-LEVEL CONCEPT HIERARCHIES

Oscar Hill[†], Mateo Espinosa Zarlenga^{†,*}, Mateja Jamnik[†]

[†]Department of Computer Science and Technology, University of Cambridge, UK

^{*}Department of Computer Science, University of Oxford, UK

ogh22@cam.ac.uk, mateo.espinosazarlenga@trinity.ox.ac.uk, mateja.jamnik@cl.cam.ac.uk

ABSTRACT

Concept-based models explain predictions using human-understandable concepts, but they typically rely on exhaustive annotations and treat concepts as flat and independent. To circumvent this, recent work has introduced *Hierarchical Concept Embedding Models* (HiCEMs) to explicitly model concept relationships, and *Concept Splitting* to discover sub-concepts using only coarse annotations. However, both methods are restricted to shallow hierarchies. We overcome this limitation with *Multi-Level Concept Splitting* (MLCS), which discovers multi-level concept hierarchies from only top-level supervision, and *Deep-HiCEMs*, an architecture that represents these discovered hierarchies and enables interventions at multiple levels of abstraction. Experiments show that MLCS discovers human-interpretable concepts absent during training and that Deep-HiCEMs maintain high accuracy while supporting test-time concept interventions that can improve task performance.

1 INTRODUCTION

Concept-based models (Koh et al., 2020) explain their decisions by predicting human-understandable concepts. By exposing intermediate concepts, these models allow users to inspect, debug, and intervene on model reasoning. However, most concept-based approaches assume concepts are flat and independent, leading to their representations failing to capture known inter-concept relationships (Raman et al., 2024). This is problematic because real-world concepts are often interrelated, and human cognition utilises these relationships (McClelland & Rogers, 2003).

To address this limitation, *Hierarchical Concept Embedding Models* (HiCEMs, Hill et al. (2026)) explicitly model concept relationships and utilise a procedure called *Concept Splitting* to reduce annotation costs by discovering sub-concepts using only coarse top-level concept labels. However, these methods are restricted to shallow hierarchies, supporting only a single layer of sub-concepts in addition to the provided coarse concepts. This prevents models from capturing richer structure and from offering interventions at several levels of abstraction.

To overcome these issues, we propose *Multi-Level Concept Splitting* (MLCS), a method for discovering multi-level concept hierarchies from concept embeddings trained only with top-level supervision, and *Deep-HiCEM*, an architecture designed to support arbitrarily deep concept hierarchies. Together, these contributions enable hierarchical explanations and interventions without exhaustive annotations.

We evaluate our approach across several datasets, including a synthetic dataset designed for hierarchical concept discovery. Our results show that multi-level human-interpretable concept hierarchies can be discovered reliably, and used to explain predictions without sacrificing performance. Our contributions are:

- We introduce **MLCS** for discovering deep concept hierarchies from top-level supervision.
- We propose **Deep-HiCEMs**, which model arbitrarily deep concept hierarchies and allow human interventions at any level in the hierarchy.
- We demonstrate that Deep-HiCEMs trained via MLCS can discover interpretable concept hierarchies. Deep-HiCEMs trained with MLCS achieve competitive task accuracies and are receptive to test-time concept interventions at different levels of granularity.

2 BACKGROUND AND RELATED WORK

Concept learning Concept-based methods (CMs) aim to explain model predictions using human-understandable concepts (e.g., “colour” or “size”) (Bau et al., 2017; Fong & Vedaldi, 2018; Kim et al., 2018). Some CMs, such as Concept Bottleneck Models (CBMs, Koh et al. (2020)) and Concept Embedding Models (CEMs, Espinosa Zarlenga et al. (2022; 2025)), explicitly incorporate concepts into model architectures, leading to inherently interpretable models that provide concept-based explanations. These models, however, typically require large concept-annotated training sets and may suffer from suboptimal predictive performance when concept labels are incomplete (Espinosa Zarlenga et al., 2022) or noisy (Penaloza et al., 2025). Furthermore, most CMs ignore relationships between concepts, instead assuming all concepts are independent (Havasi et al., 2022; Dominici et al., 2024).

Concept Splitting and HiCEMs Concept Splitting (Hill et al., 2026) discovers fine-grained sub-concepts from coarse concept labels. It uses sparse autoencoders (Bricken et al., 2023) to discover sub-concepts in concept-aligned embeddings taken from a pretrained CEM (Espinosa Zarlenga et al., 2022). This reveals interpretable, finer concepts without extra annotations, enabling more granular explanations and interventions. However, Concept Splitting is limited to a single additional level of granularity: it can reveal sub-concepts of a provided concept, but it cannot expose deeper hierarchical structure among those sub-concepts. In this work, we introduce MLCS to overcome this limitation.

HiCEMs (Hill et al., 2026) model hierarchical concept structure, using Concept Splitting to identify parent-child groups. From a latent code \mathbf{h} , a HiCEM learns two embeddings per top-level concept. The positive embedding, $\hat{\mathbf{c}}_i^{+'}$, represents concept c_i ’s active state, and the negative embedding, $\hat{\mathbf{c}}_i^{-'}$, represents its inactive state. These embeddings are then passed through sub-concepts modules, which produce new embeddings ($\hat{\mathbf{c}}_i^{+}$ and $\hat{\mathbf{c}}_i^{-}$) that include information about sub-concepts. The sub-concepts modules also output the most likely sub-concept probabilities, which are used to calculate top-level concept probabilities. These probabilities are used to construct an embedding for each concept via a weighted mixture of positive and negative embeddings.

HiCEMs support concept interventions: a user can correct concept predictions at test time, and the downstream task prediction is recomputed using the corrected concepts, enabling human-in-the-loop control. However, HiCEMs support only two levels of hierarchy. This leaves open the question of how to represent and learn deeper, multi-level hierarchies. In this paper, we address this question with Deep-HiCEM, a generalisation that supports hierarchical concept structure beyond a single parent-child layer.

3 MULTI-LEVEL CONCEPT SPLITTING

To overcome the single-level limitation of Concept Splitting, we introduce MLCS, which replaces the single-level SAE with our Hierarchical Sparse AutoEncoder (HiSAE) that learns structured sparse features at multiple levels. This enables the discovery of both sub-concepts and sub-sub-concepts from the same embedding space, allowing us to construct deeper concept hierarchies from only top-level concept supervision.

The use of a HiSAE is the primary difference between MLCS and Concept Splitting, so we focus here on the HiSAE architecture. The HiSAE is designed to learn sparse structure at two or more levels simultaneously: a top level captures candidate sub-concepts (e.g., “apple”), and a lower level, conditioned on each discovered sub-concept, captures finer-grained refinements (e.g., “red apple”).

3.1 HIERARCHICAL SPARSE AUTOENCODER

The HiSAE, inspired by Muchane et al. (2025), consists of (i) a top-level encoder that maps an input embedding to a dictionary of size K , (ii) a set of sub-encoders, one per top-level latent, each mapping the input to a sub-dictionary of size K_s , and (iii) corresponding linear decoders for both levels.

Given an input embedding \mathbf{e} , the top-level encoder produces activations over the K top-level latents. A top- k sparsification step (Bussmann et al., 2024) retains only the k largest activations.

For each of the k active top-level latents ℓ , the corresponding sub-encoder processes the input embedding \mathbf{e} to produce activations over its sub-dictionary. A second top- k_s operation selects k_s active

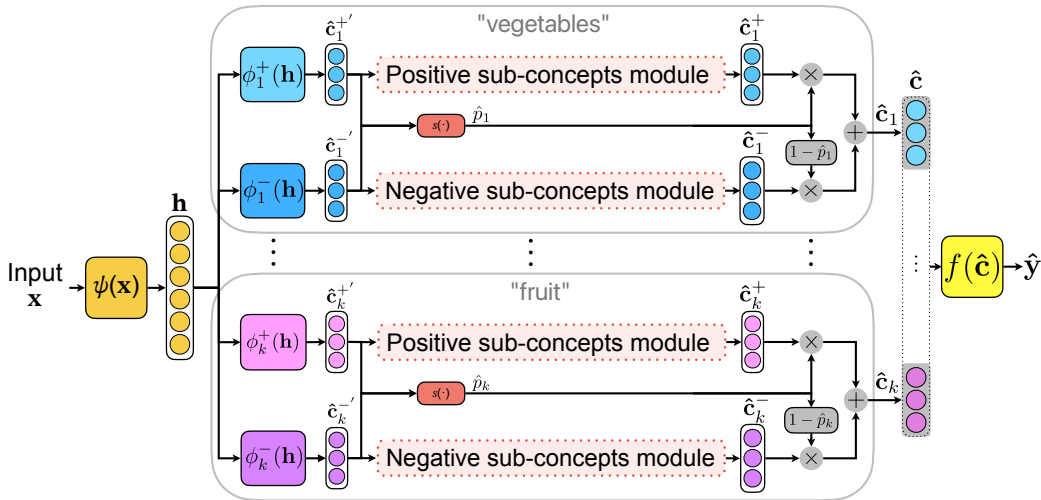


Figure 1: **Deep-HiCEM**: as in a HiCEM, from a latent code \mathbf{h} , we learn two embeddings per concept ($\hat{\mathbf{c}}_i^{+'}$ and $\hat{\mathbf{c}}_i^{-'}$), which are then passed through sub-concepts modules, which produce new embeddings ($\hat{\mathbf{c}}_i^+$ and $\hat{\mathbf{c}}_i^-$) that include information about sub-concepts and their descendants.

sub-latents. The reconstruction is formed by summing the contributions of the active top-level and sub-level decoders, and the model is trained with a standard mean-squared error reconstruction loss.

Crucially, the sub-level is gated by the top-level: sub-latents only contribute when their parent latent is active. This ties sub-sub-concepts to their parent sub-concepts.

3.2 DEPTH OF THE HIERARCHY

In principle, the HiSAE architecture can be extended recursively to yield arbitrarily deep hierarchies. MLCS is therefore a general mechanism for hierarchical concept discovery from embedding spaces.

In this work, however, we restrict ourselves to *two discovered levels*, to make evaluation more straightforward. Even with this restriction, MLCS allows models to represent richer conceptual structures than single-level Concept Splitting while still requiring supervision only at the top level.

As with Concept Splitting (Hill et al., 2026), we can interpret discovered concepts using prototypes, which provide training examples that strongly activate the concept. This approach enables experts to assign potential semantics to discovered concepts.

4 DEEP HIERARCHICAL CEMS

We introduce the Deep-HiCEM architecture (Figure 1), which extends HiCEMs (Hill et al., 2026) to support deeper concept hierarchies, like those produced by MLCS. We focus on the differences from HiCEMs; we do not discuss training Deep-HiCEMs in detail, as the process is the same as in HiCEMs.

4.1 CONCEPT STRUCTURE

Deep-HiCEMs represent concepts that have been organised into a tree. Each node in the tree represents a concept and can have both *positive* and *negative* sub-concepts. Positive sub-concepts can only be present in an example if their parent concept is present, and negative sub-concepts can only be present if their parent concept is not. The concept trees may be arbitrarily deep.

4.2 ARCHITECTURE

As in HiCEMs, for each top-level concept, a Deep-HiCEM learns a mixture of two embeddings with semantics representing the concept’s activity. Each top-level concept c_i is represented with the

embeddings $\hat{\mathbf{c}}_i^+, \hat{\mathbf{c}}_i^- \in \mathbb{R}^m$. Here, $\hat{\mathbf{c}}_i^+$ represents c_i 's active state, and $\hat{\mathbf{c}}_i^-$ represents its inactive state. We also want $\hat{\mathbf{c}}_i^+$ and $\hat{\mathbf{c}}_i^-$ to contain information about c_i 's positive and negative sub-concepts (and their descendants), respectively. To achieve this, a backbone network $\psi(\mathbf{x})$ produces a latent representation $\mathbf{h} \in \mathbb{R}^{m_{\text{hidden}}}$, which is the input to top-level embedding generators ϕ_i^+ and ϕ_i^- . These top-level embedding generators produce intermediate embeddings $\hat{\mathbf{c}}_i^{+'} = \phi_i^+(\mathbf{h}), \hat{\mathbf{c}}_i^{-'} = \phi_i^-(\mathbf{h}) \in \mathbb{R}^m$. As in HiCEMs, we implement the top-level embedding generators as single fully connected layers.

To produce final concept embeddings that contain information about sub-concepts, the embeddings $\hat{\mathbf{c}}_i^{+'}$ and $\hat{\mathbf{c}}_i^{-'}$ are passed through a positive and a negative *sub-concepts module*, respectively. The positive sub-concepts module, which we describe in further detail below, is responsible for learning the positive sub-concepts of c_i and their descendants. Intuitively, this module recursively refines a concept embedding by incorporating information from its sub-concepts. It outputs the positive concept embedding for c_i , $\hat{\mathbf{c}}_i^+$. Similarly, the negative sub-concepts module outputs the negative concept embedding for c_i , $\hat{\mathbf{c}}_i^-$. If concept c_i has no positive sub-concepts, then we take $\hat{\mathbf{c}}_i^+ = \hat{\mathbf{c}}_i^{+'}$. We proceed analogously in the absence of negative sub-concepts. The probability of concept c_i is calculated as $\hat{p}_i = s([\hat{\mathbf{c}}_i^{+'}, \hat{\mathbf{c}}_i^{-'}]^T)$, where s is a shared scoring function that calculates concept probabilities from concept embeddings. As in HiCEMs, the final concept embedding $\hat{\mathbf{c}}_i$ for c_i is calculated as a weighted mixture of $\hat{\mathbf{c}}_i^+$ and $\hat{\mathbf{c}}_i^-$: $\hat{\mathbf{c}}_i = \hat{p}_i \hat{\mathbf{c}}_i^+ + (1 - \hat{p}_i) \hat{\mathbf{c}}_i^-$. Downstream task predictions are calculated as in HiCEMs.

Like HiCEMs, Deep-HiCEMs support concept interventions, such that when a sub-concept is intervened on, its parent may also be updated (e.g., if a human expert informs the model that a positive sub-concept is present, this implies that its parent is also present).

4.3 SUB-CONCEPTS MODULES

We describe a positive sub-concepts module, but negative sub-concepts modules operate similarly. Sub-concept embeddings are produced in the same way as top-level concept embeddings (Figure 1 in Section 4.2), with nested sub-concepts modules if the sub-concept has sub-concepts itself. However, instead of the latent code \mathbf{h} , sub-concept embedding generators take as input the preliminary embedding of the parent concept. The concatenated mixed embeddings for all the sub-concepts are passed through an embedding compressor (implemented as a fully connected layer) to produce a fixed-length positive embedding for the parent concept, $\hat{\mathbf{c}}_i^+$. This recursive structure enables information from deeper descendants to be propagated upwards through the hierarchy.

5 EXPERIMENTS

We evaluate MLCS and Deep-HiCEMs by exploring the following research questions:

- RQ1** Does MLCS discover interpretable concept hierarchies?
- RQ2** How do Deep-HiCEMs' task accuracies compare to those of standard HiCEMs?
- RQ3** Can a Deep-HiCEM's task accuracy be improved by intervening on discovered concepts?

5.1 PSEUDOKITCHENS-2

To evaluate how well MLCS and Deep-HiCEMs can discover and represent two levels of hierarchy (sub-concepts and sub-sub-concepts), we adapt the PseudoKitchens dataset introduced by Hill et al. (2026) so that both sub-concepts and sub-sub-concepts can be discovered. The discoverable sub-concepts are ingredients (e.g., "apple" for the "fruit" top-level concept), and the sub-sub-concepts correspond to different variants of that ingredient (e.g., "red apple"). Full details are given in Appendix A.

5.2 SETUP

Datasets We evaluate our method across five datasets: MNIST-ADD (Hill et al. (2026), based on LeCun et al. (2010)), the SHAPES dataset (Hill et al., 2026), Caltech-UCSD Birds-200-2011 (CUB) (Wah et al., 2011), Animals with Attributes 2 (AwA2) (Xian et al., 2019), and PseudoKitchens-2. Appendix A describes PseudoKitchens-2, and details for the other datasets are given in (Hill et al., 2026).

Table 1: Mean ROC-AUC for discovered concepts. LF-CBMs were unable to discover concepts on some datasets. Concepts discovered with MLCS are predicted accurately.

	MNIST-ADD	SHAPES	CUB	AwA2	PseudoKitchens-2
LF-CBM	–	0.75 \pm 0.00	0.77 \pm 0.00	0.78 \pm 0.00	–
HiCEM + Concept Splitting	0.93 \pm 0.01	0.93 \pm 0.01	0.85 \pm 0.01	0.88 \pm 0.01	0.80 \pm 0.00
Deep-HiCEM + MLCS (ours)	0.94 \pm 0.01	0.93 \pm 0.01	0.84 \pm 0.01	0.85 \pm 0.01	0.79 \pm 0.01

Metrics Following Hill et al. (2026), we first run MLCS on the provided concepts using an initial CEM, and then train a Deep-HiCEM with both the provided top-level concepts and the concepts discovered by MLCS. We evaluate discovered concepts and perform interventions by automatically aligning discovered concepts with a human-interpretable “left-out” concept from a predefined “concept bank”. This bank contains anticipated concepts that are excluded during the initial CEM training (e.g., “the first digit is 5” in MNIST-ADD), and each bank concept is associated in advance with the parent concept whose sub-concept it may represent.

To match discovered concepts to the bank, we compute ROC-AUC scores between the discovered concept labels and their potential parent-concept-associated matches in the bank. As in Hill et al. (2026), each concept in the bank is assigned to the sub-concept with the highest ROC-AUC score, as long as that score is greater than 0.7. Discovered concepts without a bank match above this threshold are excluded from evaluation.

For PseudoKitchens-2, the bank contains both sub-concepts (e.g., “apple”) and sub-sub-concepts (e.g., “red apple”). When matching discovered sub-concepts to the concept bank, we use the average ROC-AUC of the sub-concept and any associated sub-sub-concepts. For all other datasets, we match and evaluate only sub-concepts, as ground truth sub-sub-concept labels are not available.

To answer **RQ1**, we report the average discovered concept ROC-AUC of the Deep-HiCEM. For **RQ2**, we compare task accuracy and provided-concept ROC-AUC between a standard HiCEM and a Deep-HiCEM. For **RQ3**, we measure the change in the task accuracy of Deep-HiCEMs as concepts are intervened on. All metrics are computed on the test sets using three random seeds, and we report means and standard deviations. Concept accuracy is always summarised using mean ROC-AUC to avoid misleading results from majority-class predictors. For further details, see Appendix B.

Baselines We compare our Deep-HiCEMs against standard HiCEMs (Hill et al., 2026). We also report results for the baselines considered by Hill et al. (2026): black-box models, CEMs (Espinosa Zarlenga et al., 2022), CBMs (Koh et al., 2020), label-free CBMs (*LF-CBMs*, (Oikarinen et al., 2023)), Post-hoc CBMs (*PCBMs*, (Yuksekgonul et al., 2023)), and PCBMs with residual connections (*PCBM-hs*). For the PseudoKitchens-2 dataset, we implement the baselines following the protocol of Hill et al. (2026); for the other datasets, we use the results reported in their work.

5.3 RESULTS

Discovered concepts are human-interpretable and can be predicted accurately (RQ1, Table 1).

We report the accuracy of the discovered concept predictions made by our models using the ground truth labels of the corresponding human-interpretable concept bank concepts on the test datasets. Table 1 shows that the mean discovered concept ROC-AUCs of Deep-HiCEMs are consistently high, always within a few percentage points of those of HiCEMs. This indicates that the labels produced by MLCS align Deep-HiCEM concept activations with human-interpretable concepts. Matching HiCEMs within a few percentage points is a strong result, given that our method discovers additional levels of hierarchy. On PseudoKitchens-2, the majority of sub-sub-concepts in the bank were successfully matched to discovered sub-sub-concepts on all runs. The mean sub-concept ROC-AUC was 0.80 \pm 0.00, and the mean sub-sub-concept ROC-AUC was 0.79 \pm 0.01. This demonstrates that MLCS and Deep-HiCEMs can discover and represent deeper concept hierarchies while maintaining discovered-concept ROC-AUCs close to those of HiCEMs after Concept Splitting.

Deep-HiCEMs have high task accuracy (RQ2, Table 2). We measure the task accuracy of Deep-HiCEMs and our baselines. The results are in Table 2. Deep-HiCEMs achieve high task accuracy compared to the baselines. In particular, their performance remains closely aligned with

Table 2: Task accuracies. The task accuracy of Deep-HiCEMs is competitive with all our baselines.

	MNIST-ADD	SHAPES	CUB	AwA2	PseudoKitchens-2
Black box (not interpretable)	0.94 \pm 0.00	0.89 \pm 0.00	0.80 \pm 0.00	0.98 \pm 0.00	0.63 \pm 0.02
LF-CBM	–	0.59 \pm 0.01	0.80\pm0.00	0.94 \pm 0.00	–
PCBM	0.16 \pm 0.03	0.54 \pm 0.01	0.65 \pm 0.01	0.95 \pm 0.00	0.12 \pm 0.00
PCBM-h	0.53 \pm 0.01	0.73 \pm 0.00	0.73 \pm 0.00	0.96 \pm 0.00	0.48 \pm 0.00
CBM	0.23 \pm 0.01	0.78 \pm 0.01	0.65 \pm 0.00	0.97 \pm 0.00	0.60\pm0.01
CEM	0.92\pm0.01	0.89\pm0.00	0.76 \pm 0.01	0.98\pm0.00	0.59 \pm 0.01
HiCEM + Concept Splitting	0.92\pm0.00	0.87 \pm 0.02	0.74 \pm 0.01	0.98\pm0.00	0.57 \pm 0.03
Deep-HiCEM + MLCS (ours)	0.92\pm0.00	0.87 \pm 0.01	0.73 \pm 0.01	0.97 \pm 0.00	0.58 \pm 0.01

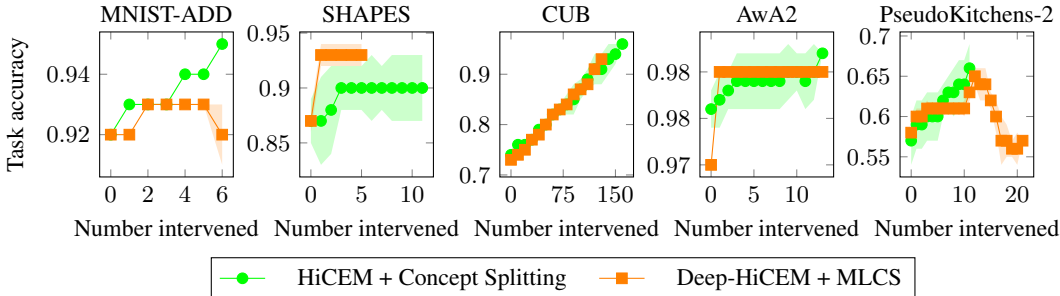


Figure 2: Task accuracy as discovered concepts are intervened on. Intervening on discovered concepts improves task accuracy, with a few exceptions that would benefit from further investigation.

standard HiCEMs, despite modelling deeper concept hierarchies. We investigate the provided concept accuracy of Deep-HiCEMs in Appendix C, and find that Deep-HiCEMs can predict provided concepts as well as HiCEMs.

Intervening on concepts discovered by MLCS can enhance task accuracy, with a few exceptions that warrant further investigation (RQ3, Figure 2) As shown in Figure 2, intervening on discovered concepts can lead to an increase in task accuracy, although interventions on some discovered concepts are not beneficial. The decreases in task accuracy caused by discovered concept interventions, particularly in PseudoKitchens-2, warrant further investigation. We hypothesise that these decreases are due to inaccuracies or biases in the discovered concept labels. We evaluate interventions on the provided top-level concepts in Appendix D and find that provided concept interventions perform equally well in Deep-HiCEMs and HiCEMs. Overall, these results indicate that many of the concepts discovered by MLCS are not only interpretable but also actionable, and that leveraging them through interventions can translate into performance gains.

6 LIMITATIONS AND CONCLUSION

We introduced MLCS and Deep-HiCEMs, a framework for discovering and modelling deep concept hierarchies from only top-level supervision. We showed that our approach reliably uncovers human-interpretable concept hierarchies absent during training while maintaining competitive task accuracy. However, interventions on discovered concepts are not consistently beneficial, motivating further study. Another limitation of our approach is that SAEs are not guaranteed to discover meaningful concepts. Future work includes scaling to larger, more complex datasets, and extending evaluations to investigate even deeper hierarchies. Despite the limitations mentioned, modelling deeper concept structure is a practical step towards more expressive concept-based interpretability.

ACKNOWLEDGMENTS

This work was supported by the Engineering and Physical Sciences Research Council [EP/Y030826/1]. A significant portion of this work was carried out whilst MEZ was at the University of Cambridge, funded by the Gates Cambridge Trust via a Gates Cambridge Scholarship.

REFERENCES

- David Bau, Bolei Zhou, Aditya Khosla, Aude Oliva, and Antonio Torralba. Network dissection: Quantifying interpretability of deep visual representations. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3319–3327, Los Alamitos, CA, USA, 07 2017. IEEE Computer Society. doi: 10.1109/CVPR.2017.354. URL <https://doi.ieeecomputersociety.org/10.1109/CVPR.2017.354>.
- Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermyn, 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. Towards monosemanticity: Decomposing language models with dictionary learning. *Transformer Circuits Thread*, 2023. <https://transformer-circuits.pub/2023/monosemantic-features/index.html>.
- Bart Bussmann, Patrick Leask, and Neel Nanda. Batchtopk sparse autoencoders. In *NeurIPS 2024 Workshop on Scientific Methods for Understanding Deep Learning*, 2024. URL <https://openreview.net/forum?id=d4dpOCqybL>.
- Gabriele Dominici, Pietro Barbiero, Mateo Espinosa Zarlenga, Alberto Termine, Martin Gjoreski, Giuseppe Marra, and Marc Langheinrich. Causal concept graph models: Beyond causal opacity in deep learning. *arXiv preprint arXiv:2405.16507*, 2024.
- Mateo Espinosa Zarlenga, Pietro Barbiero, Gabriele Ciravegna, Giuseppe Marra, Francesco Giannini, Michelangelo Diligenti, Zohreh Shams, Frederic Precioso, Stefano Melacci, Adrian Weller, Pietro Lió, and Mateja Jamnik. Concept embedding models: Beyond the accuracy-explainability trade-off. In *Advances in Neural Information Processing Systems*, volume 35, pp. 21400–21413. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/867c06823281e506e8059f5c13a57f75-Paper-Conference.pdf.
- Mateo Espinosa Zarlenga, Dominici Gabriele, Barbiero Pietro, Zohreh Shams, and Mateja Jamnik. Avoiding leakage poisoning: Concept interventions under distribution shifts. In *International Conference on Machine Learning (ICML)*, 2025.
- Ruth Fong and Andrea Vedaldi. Net2vec: Quantifying and explaining how concepts are encoded by filters in deep neural networks. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 8730–8738, Los Alamitos, CA, USA, 06 2018. IEEE Computer Society. doi: 10.1109/CVPR.2018.00910. URL <https://doi.ieeecomputersociety.org/10.1109/CVPR.2018.00910>.
- Marton Havasi, Sonali Parbhoo, and Finale Doshi-Velez. Addressing Leakage in Concept Bottleneck Models. In *Advances in Neural Information Processing Systems*, 2022.
- Oscar Hill, Mateo Espinosa Zarlenga, and Mateja Jamnik. Hierarchical concept-based interpretable models. In *The Fourteenth International Conference on Learning Representations*, 2026. URL <https://openreview.net/forum?id=h610IERd38>.
- Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viegas, and Rory Sayres. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (TCAV). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2668–2677. PMLR, 07 2018. URL <https://proceedings.mlr.press/v80/kim18d.html>.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. Concept bottleneck models. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 5338–5348. PMLR, 07 2020. URL <https://proceedings.mlr.press/v119/koh20a.html>.

- Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database. Online, 2010. URL <http://yann.lecun.com/exdb/mnist>.
- James L. McClelland and Timothy T. Rogers. The parallel distributed processing approach to semantic cognition. *Nature Reviews Neuroscience*, 4(4):310–322, 04 2003. ISSN 1471-0048. doi: 10.1038/nrn1076. URL <https://doi.org/10.1038/nrn1076>.
- Mark Muchane, Sean Richardson, Kiho Park, and Victor Veitch. Incorporating hierarchical semantics in sparse autoencoder architectures, 2025. URL <https://arxiv.org/abs/2506.01197>.
- Tuomas P. Oikarinen, Subhro Das, Lam M. Nguyen, and Tsui-Wei Weng. Label-free concept bottleneck models. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL <https://openreview.net/pdf?id=FlCg47MNvBA>.
- 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 Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. *PyTorch: an imperative style, high-performance deep learning library*. Curran Associates Inc., Red Hook, NY, USA, 2019.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-learn: Machine learning in python. *J. Mach. Learn. Res.*, 12(null):2825–2830, November 2011. ISSN 1532-4435.
- Emiliano Penalosa, Tianyue H Zhang, Laurent Charlin, and Mateo Espinosa Zarlenga. Addressing concept mislabeling in concept bottleneck models through preference optimization. In *International Conference on Machine Learning*, pp. 48762–48788. PMLR, 2025.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763. PMLR, 2021. URL <http://proceedings.mlr.press/v139/radford21a.html>.
- Naveen Raman, Mateo Espinosa Zarlenga, and Mateja Jamnik. Understanding inter-concept relationships in concept-based models. In *Proceedings of the 41st International Conference on Machine Learning, ICML’24*. JMLR.org, 2024.
- Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. Technical Report CNS-TR-2011-001, August 2011.
- Yongqin Xian, Christoph H. Lampert, Bernt Schiele, and Zeynep Akata. Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(9):2251–2265, 2019. doi: 10.1109/TPAMI.2018.2857768.
- Mert Yuksekgonul, Maggie Wang, and James Zou. Post-hoc concept bottleneck models. In *The Eleventh International Conference on Learning Representations, 2023*. URL <https://openreview.net/forum?id=nA5AZ8CEyow>.

A PSEUDOKITCHENS-2

This appendix describes PseudoKitchens-2, an extension of the PseudoKitchens dataset (Hill et al., 2026) designed to enable the discovery and evaluation of two-level concept hierarchies.

PseudoKitchens-2 is generated in the same way as PseudoKitchens (Hill et al., 2026). The recipes described by Hill et al. (2026) are altered so that different variations of the same ingredient are distinguished (e.g., some recipes require “red apples” while others require “green apples”). The ingredients that have variations are apple (5 variations), potato (5 variations), and pepper (3 variations).

A.1 RECIPES

We adapted the recipes used by Hill et al. (2026) that define valid combinations of ingredients for the classification task. Some ingredients are organised into groups as shown in Table 3 (these groups are the same as the ones used by Hill et al. (2026)). If a recipe contains an ingredient group, a random number of ingredients are selected from that group, unless the group is pasta, in which case only one type of pasta is selected. The recipes used in PseudoKitchens-2 are shown in Table 4. Where a list of numbers is specified in brackets after an ingredient (e.g., “Apple (1)”), this refers to the variants of that ingredient that are acceptable. For each instance, a recipe is chosen uniformly at random.

Table 3: Ingredient groups in PseudoKitchens and PseudoKitchens-2.

Group	Ingredients
Fruit	Banana, Orange, Apple, Pear, Pineapple
Vegetables	Onion, Carrot, Potato, Pepper, Courgette
Pasta	Macaroni, Spaghetti

Table 4: Recipes in PseudoKitchens-2.

Recipe	Ingredients
Fruit Salad	Fruit
Vegetable Pasta	Pasta, Onion, Garlic, Oil, Vegetables, Spice, Tin Tomatoes
Risotto	Cheese, Onion, Garlic, Vegetables, Oil, Spice, Rice
Chips	Potato (2, 3, 4), Oil, Flour, Garlic, Spice
Chilli	Mince, Oil, Onion, Garlic, Chilli, Tin Tomatoes, Spice, Rice
Smoothie	Milk, Yoghurt, Fruit
Hot Chocolate	Chocolate, Milk
Banana Bread	Butter, Sugar, Egg, Flour, Banana
Chocolate Fudge Cake	Egg, Sugar, Oil, Flour, Chocolate, Syrup, Milk
Carbonara	Garlic, Meat, Butter, Cheese, Egg, Spaghetti, Spice
Apple Crumble	Apple (3), Sugar, Flour, Butter
Salad	Pepper (2, 3), Apple (1, 2, 4, 5), Potato (1, 5)

The concepts provided to the initial CEM are the ingredient groups in Table 3 (e.g., “contains fruit”), as well as all the ingredients that are not part of a group. The concept bank sub-concepts correspond to the ingredients in the ingredient groups (e.g., “contains apples”), and the sub-sub-concepts correspond to variations of the same ingredient (e.g., “contains green apples”).

A.2 DATASET COMPOSITION

Like the original PseudoKitchens, the complete PseudoKitchens-2 dataset comprises 10,000 training images, 1,000 validation images, and 1,000 test images. Each image is rendered at 512×512 resolution.

Table 5: Mean ROC-AUCs for provided concepts. Deep-HiCEMs are able to predict provided concepts just as well as HiCEMs.

	MNIST-ADD	SHAPES	CUB	AwA2	PseudoKitchens-2
CBM	0.99 \pm 0.00	1.00 \pm 0.00	0.89 \pm 0.00	1.00 \pm 0.00	0.90 \pm 0.00
CEM	0.99 \pm 0.00	1.00 \pm 0.00	0.95 \pm 0.00	1.00 \pm 0.00	0.91 \pm 0.00
HiCEM + Concept Splitting	0.99 \pm 0.00	1.00 \pm 0.00	0.93 \pm 0.01	1.00 \pm 0.00	0.91 \pm 0.00
Deep-HiCEM + MLCS (ours)	0.99 \pm 0.00	1.00 \pm 0.00	0.94 \pm 0.00	1.00 \pm 0.00	0.91 \pm 0.00

B MODEL ARCHITECTURES, TRAINING AND HYPERPARAMETERS

Model architectures We use the CLIP ViT-L/14 foundation model (Radford et al., 2021) as the backbone for all of our models and baselines. We do not fine-tune the foundation model: we just use the representations it outputs. For all the models, we precompute the representations with standard image preprocessing pipelines and do not use any data augmentations.

When training BatchTopK SAEs, like Hill et al. (2026), we use the default hyperparameters in the code released by Bussmann et al. (2024) for all datasets.¹ The key distinction of this method is its enforcement of sparsity: it selects the top $n \cdot k$ activations across an entire batch of n samples. This allows the number of active features to vary per sample, targeting an average of $k = 32$ active features. The SAEs are trained for 300 epochs with a dictionary size of 12,288 and a learning rate of 3×10^{-4} .

For MLCS, we train Hierarchical Sparse Autoencoders (HiSAEs) to discover concept hierarchies. The HiSAEs use a top-level dictionary of size 4096 with top- k sparsification ($k = 32$), and for each top-level latent we learn a sub-dictionary of size 512 with top- k_s sparsification ($k_s = 16$). HiSAEs are trained with a batch size of 1000 for up to 100 epochs.

All of our CEM and HiCEM concept embeddings have $m = 16$ activations. Across all datasets, we always use a single fully connected layer for label predictor f .

Training hyperparameters Our models are trained using the Adam optimisation algorithm (Kingma & Ba, 2015) with a learning rate of 1×10^{-3} . They are trained for a maximum of 300 epochs, and training is stopped if the validation loss does not improve for 75 epochs. We use a batch size of 256.

In all CEMs, (Deep-)HiCEMs and CBMs, the weight of the concept loss is set to $\lambda = 10$. Following (Koh et al., 2020), in MNIST-ADD, CUB and PseudoKitchens-2 we use a weighted cross-entropy loss for concept prediction to mitigate imbalances in concept labels. In MNIST-ADD, we also use a weighted cross-entropy loss for task prediction to mitigate imbalances in task labels.

When training CEMs and (Deep-)HiCEMs, the RandInt (Espinosa Zarlenga et al., 2022) regularisation strategy is used: at training time, concepts are intervened independently at random, with the probability of an intervention being $p_{\text{int}} = 0.25$. We choose $p_{\text{int}} = 0.25$ because Espinosa Zarlenga et al. (2022) find that it enables effective interventions while giving good performance.

C EVALUATING PROVIDED CONCEPT ACCURACY

As shown in Table 5, Deep-HiCEMs are able to predict provided concepts just as well as HiCEMs.

D EVALUATING PROVIDED CONCEPT INTERVENTIONS

Figure 3 shows that provided concept interventions perform equally well in Deep-HiCEMs and HiCEMs.

¹<https://github.com/bartbussmann/BatchTopK/blob/main/config.py>

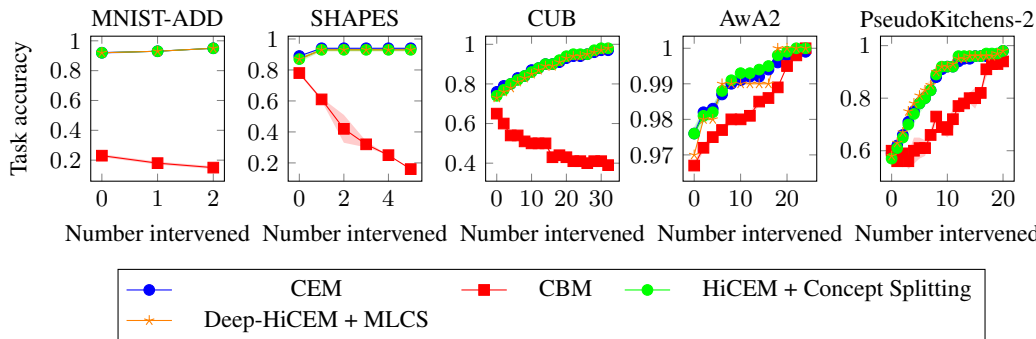


Figure 3: Change in task accuracy as provided concepts are intervened on. Provided concept interventions work just as well in Deep-HiCEMs as they do in HiCEMs.

E CODE, LICENSES, AND RESOURCES

Assets We used the CLIP foundation models (Radford et al., 2021) (<https://github.com/openai/CLIP>), whose code is available under the MIT license. To run our experiments, we made use of the CEM (Espinosa Zarlenga et al., 2022) (<https://github.com/mateoespinosa/cem>, MIT license), Post-hoc CBM (Yuksekgonul et al., 2023) (<https://github.com/mertyg/post-hoc-cbm>, MIT license) and Label-free CBM (Oikarinen et al., 2023) (<https://github.com/Trustworthy-ML-Lab/Label-free-CBM>) repositories. We implemented our experiments in Python 3.12 and used open-source libraries such as PyTorch 2.9 (Paszke et al., 2019) (BSD license) and Scikit-learn (Pedregosa et al., 2011) (BSD license). We have released the code required to recreate our experiments in a MIT-licensed public repository.²

Resources All of our experiments were run on virtual machines with at least 8 CPU cores, 18GB of RAM, and an NVIDIA GPU (Quadro RTX 8000 or GeForce RTX 4090). We estimate that approximately 100 GPU hours were required to complete our work.

Use of AI We used Large Language Models (LLMs) as assistants for drafting and improving the clarity and grammar of this manuscript. LLMs were also used to generate boilerplate code. However, all core research ideas, experimental design, and analysis of the results were conducted by the authors.

²<https://github.com/OscarPi/cem-concept-discovery>