



HAL
open science

CRAFT: Concept Recursive Activation FacTORIZATION for Explainability

Thomas Fel, Agustin Picard, Louis Bethune, Thibaut Boissin, David
Vigouroux, Julien Colin, Rémi Cadène, Thomas Serre

► **To cite this version:**

Thomas Fel, Agustin Picard, Louis Bethune, Thibaut Boissin, David Vigouroux, et al.. CRAFT: Concept Recursive Activation FacTORIZATION for Explainability. Proceedings of the IEEE / CVF Computer Vision and Pattern Recognition Conference (CVPR), 2023., 2023, Vancouver, Canada. hal-04049806

HAL Id: hal-04049806

<https://hal.science/hal-04049806v1>

Submitted on 28 Mar 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

CRAFT: Concept Recursive Activation FacTORization for Explainability

Thomas Fel^{1,3,5*} Agustin Picard^{3,6*} Louis Bethune^{3*} Thibaut Boissin^{3,4*}
David Vigouroux^{3,4} Julien Colin^{1,3} Rémi Cadène^{1,2} Thomas Serre^{1,3}

¹Carney Institute for Brain Science, Brown University, USA ²Sorbonne Université, CNRS, France

³Artificial and Natural Intelligence Toulouse Institute, Université de Toulouse, France

⁴Institut de Recherche Technologique Saint-Exupery, France

⁵Innovation & Research Division, SNCF , ⁶Scalian

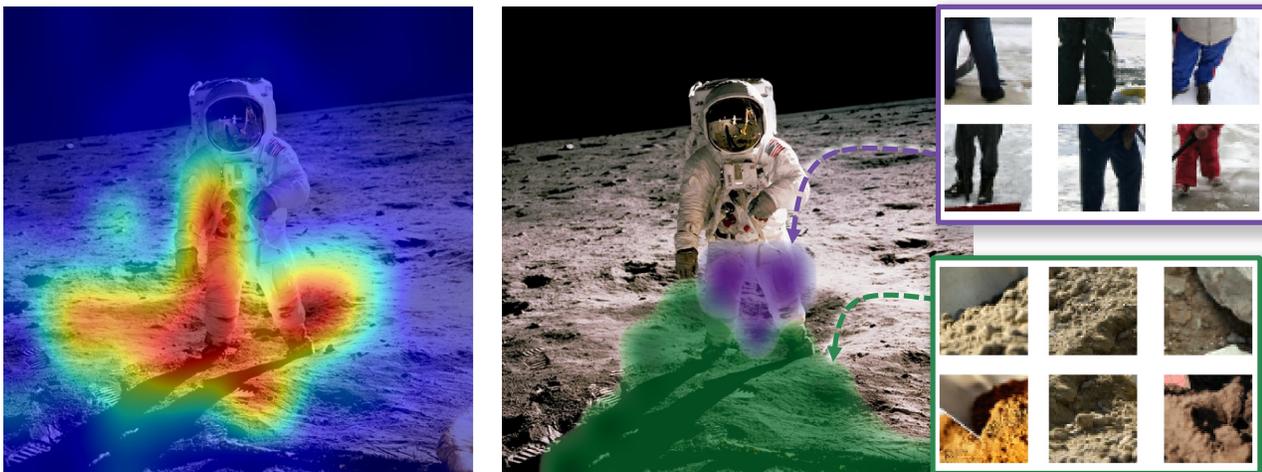


Figure 1. The “Man on the Moon” incorrectly classified as a “shovel” by an ImageNet-trained ResNet50. Heatmap generated by a classic attribution method [55] (left) vs. *concept attribution maps* generated with the proposed CRAFT approach (right) which highlights the two most influential concepts that drove the ResNet50’s decision along with their corresponding locations. CRAFT suggests that the neural net arrived at its decision because it identified the concept of “dirt” ● commonly found in members of the image class “shovel” and the concept of “ski pants” ● typically worn by people clearing snow from their driveway with a shovel instead the correct concept of astronaut’s pants (which was probably never seen during training).

Abstract

Attribution methods, which employ heatmaps to identify the most influential regions of an image that impact model decisions, have gained widespread popularity as a type of explainability method. However, recent research has exposed the limited practical value of these methods, attributed in part to their narrow focus on the most prominent regions of an image – revealing “where” the model looks, but failing to elucidate “what” the model sees in those areas. In this work, we try to fill in this gap with CRAFT – a novel approach to identify both “what” and “where” by generating concept-based explanations. We introduce 3 new ingredients to the automatic concept extraction literature: (i) a recursive strategy to detect and decompose concepts across layers, (ii) a novel method for a more faithful estimation of concept importance using Sobol indices, and (iii) the use of implicit differentiation to unlock Concept Attribution Maps.

We conduct both human and computer vision experiments to demonstrate the benefits of the proposed approach. We show that the proposed concept importance estimation technique is more faithful to the model than previous methods. When evaluating the usefulness of the method for human experimenters on a human-centered utility benchmark, we find that our approach significantly improves on two of the three test scenarios. Our code is freely available: github.com/deel-ai/Craft

1. Introduction

Interpreting the decisions of modern machine learning models such as neural networks remains a major challenge. Given the ever-increasing range of machine learning applications, the need for robust and reliable explainability

* Equal contribution

Proceedings of the IEEE / CVF Computer Vision and Pattern Recognition Conference (CVPR), 2023.



Figure 2. **CRAFT results for the prediction “chain saw”**. First, our method uses Non-Negative Matrix Factorization (NMF) to extract the most relevant concepts used by the network (ResNet50V2) from the train set (ILSVRC2012 [10]). The global influence of these concepts on the predictions is then measured using Sobol indices (right panel). Finally, the method provides local explanations through *concept attribution maps* (heatmaps associated with a concept, and computed using grad-CAM by backpropagating through the NMF concept values with implicit differentiation). Besides, concepts can be interpreted by looking at crops that maximize the NMF coefficients. For the class “chain saw”, the detected concepts seem to be: ● the chainsaw engine, ● the saw blade, ● the human head, ● the vegetation, ● the jeans and ● the tree trunk.

methods continues to grow [11, 35]. Recently enacted European laws (including the General Data Protection Regulation (GDPR) [37] and the European AI act [43]) require the assessment of explainable decisions, especially those made by algorithms.

In order to try to meet this growing need, an array of explainability methods have already been proposed [14, 50, 59, 61, 62, 66, 69, 71, 76]. One of the main class of methods called attribution methods yields heatmaps that indicate the importance of individual pixels for driving a model’s decision. However, these methods exhibit critical limitations [1, 27, 63, 65], as they have been shown to fail – or only marginally help – in recent human-centered benchmarks [8, 26, 41, 49, 60, 64]. It has been suggested that their limitations stem from the fact that they are only capable of explaining *where* in an image are the pixels that are critical to the decision but they cannot tell *what* visual features are actually driving decisions at these locations. In other words, they show where the model looks but not what it sees. For example, in the scenario depicted in Fig. 1, where an ImageNet-trained ResNet mistakenly identifies an image as containing a shovel, the attribution map displayed on the left fails to explain the reasoning behind this misclassification.

A recent approach has sought to move past attribution methods [40] by using so-called “concepts” to communicate information to users on how a model works. The goal is to find human-interpretable concepts in the activation space of a neural network. Although the approach exhibited potential, its practicality is significantly restricted due to the need for prior knowledge of pertinent concepts in its original formulation and, more critically, the requirement for a labeled dataset of such concepts. Several lines of work have focused on trying to automate the concept discovery process based only on the training dataset and without explicit human supervision.

The most prominent of these techniques, ACE [24], uses a combination of segmentation and clustering techniques but requires heuristics to remove outliers. However, ACE provides a proof of concept that it might be possible to discover concepts automatically and at scale – without additional labeling or human supervision. Nevertheless, the approach suffers several limitations: by construction, each image segment can only belong to a single cluster, a layer has to be selected by the user to be used to retrieve the relevant concepts, and the amount of information lost during the outlier rejection phase can be a cause of concern. More recently, Zhang et al. [77] proposes to leverage matrix decompositions on internal feature maps to discover concepts.

Here, we try to fill these gaps with a novel method called CRAFT which uses Non-Negative Matrix Factorization (NMF) [46] for concept discovery. In contrast to other concept-based explanation methods, our approach provides an explicit link between their global and local explanations (Fig. 2) and identifies the relevant layer(s) to use to represent individual concepts (Fig. 3). Our main contributions can be described as follows:

(i) A novel approach for the automated extraction of high-level concepts learned by deep neural networks. We validate its practical utility to users with human psychophysics experiments.

(ii) A recursive procedure to automatically identify concepts and sub-concepts at the right level of granularity – starting with our decomposition at the top of the model and working our way upstream. We validate the benefit of this approach with human psychophysics experiments showing that (i) the decomposition of a concept yields more coherent sub-concepts and (ii) that the groups of points formed by these sub-concepts are more refined and appear meaningful to humans.

(iii) A novel technique to quantify the importance of individual concepts for a model’s prediction using Sobol indices [34, 67, 68] – a technique borrowed from Sensitivity Analysis.

(iv) The first concept-based explainability method which produces *concept attribution maps* by backpropagating concept scores into the pixel space by leveraging the implicit function theorem in order to localize the pixels associated with the concept of a given input image. This effectively opens up the toolbox of both white-box [15, 59, 62, 66, 69, 71, 78] and black-box [14, 47, 55, 57] explainability methods to derive concept-wise attribution maps.

2. Related Work

Attribution methods Attribution methods are widely used as post-hoc explainability techniques to determine the input variables that contribute to a model’s prediction by generating importance maps, such as the ones shown in Fig. 1. The first attribution method, Saliency, introduced in [62], generates a heatmap by utilizing the gradient of a given classification score with respect to the pixels. This method was later improved upon in the context of deep convolutional networks for classification in subsequent studies, such as [66, 69, 71, 76]. However, the image gradient only reflects the model’s operation within an infinitesimal neighborhood around an input, which can yield misleading importance estimates [22] since gradients of large vision models are notoriously noisy [66]. That is why several methods leverage perturbations on the input image to probe the model and create importance maps that indicate the most crucial areas for the decision, such as Rise [55], Sobol [14], or more recently HSIC [50].

Unfortunately, a severe limitation of these approaches – apart from the fact that they only show the “*where*” – is that they are subject to confirmation bias: while they may appear to offer useful explanations to a user, sometimes these explanations are actually incorrect [1, 23, 65]. These limitations raise questions about their usefulness, as recent research has shown by using human-centered experiments to evaluate the utility of attribution [8, 26, 41, 49, 60].

In particular, in [8], a protocol is proposed to measure the usefulness of explanations, corresponding to how much they help users identify rules driving a model’s predictions (correct or incorrect) that transfer to unseen data – using the concept of meta-predictor (also called simulatability) [11, 18, 39]. The main idea is to train users to predict the output of the system using a small set of images along with associated model predictions and corresponding explanations. A method that performs well on this benchmark is said useful, as it help users better predict the output of the model by providing meaningful information about the internal functioning of the model. This framework being agnostic to the type of explainability method, we have chosen to use it in Section 4

in order to compare CRAFT with attribution methods.

Concepts-based methods Kim et al. [40] introduced a method aimed at providing explanations that go beyond attribution-based approaches by measuring the impact of pre-selected concepts on a model’s outputs. Although this method appears more interpretable to human users than standard attribution techniques, it requires a database of images describing the relevant concepts to be manually curated. Ghorbani et al. [24] further extended the approach to extract concepts without the need for human supervision. The approach, called ACE [24], uses a segmentation scheme on images, that belong to an image class of interest. The authors leveraged the intermediate activations of a neural network for specific image segments. These segments were resized to the appropriate input size and filled with a baseline value. The resulting activations were clustered to produce prototypes, which they referred to as “concepts”. However, some concepts contained background segments, leading to the inclusion of uninteresting and outlier concepts. To address this, the authors implemented a postprocessing cleanup step to remove these concepts, including those that were present in only one image of the class and were not representative. While this improved the interpretability of their explanations to human subjects, the use of a baseline value filled around the segments could introduce biases in the explanations [28, 31, 42, 70].

Zhang et al. [77] developed a solution to the unsupervised concept discovery problem by using matrix factorizations in the latent spaces of neural networks. However, one major drawback of this method is that it operates at the level of convolutional kernels, leading to the discovery of localized concepts. For example, the concept of “grass” at the bottom of the image is considered distinct from the concept of “grass” at the top of the image.

3. Overview of the method

In this section, we first describe our concept activations factorization method. Below we highlight the main differences with related work. We then proceed to introduce the three novel ingredients that make up CRAFT: (1) a method to recursively decompose concepts into sub-concepts, (2) a method to better estimate the importance of extracted concepts, and (3) a method to use any attribution method to create *concept attribution maps*, using implicit differentiation [4, 25, 44].

Notations In this work, we consider a general supervised learning setting, where $(\mathbf{x}_1, \dots, \mathbf{x}_n) \in \mathcal{X}^n \subseteq \mathbb{R}^{n \times d}$ are n inputs images and $(y_1, \dots, y_n) \in \mathcal{Y}^n$ their associated labels. We are given a (machine-learned) black-box predictor $\mathbf{f} : \mathcal{X} \rightarrow \mathcal{Y}$, which at some test input \mathbf{x} predicts the output $\mathbf{f}(\mathbf{x})$. Without loss of generality, we establish that \mathbf{f} is a neural network that can be decomposed into two distinct

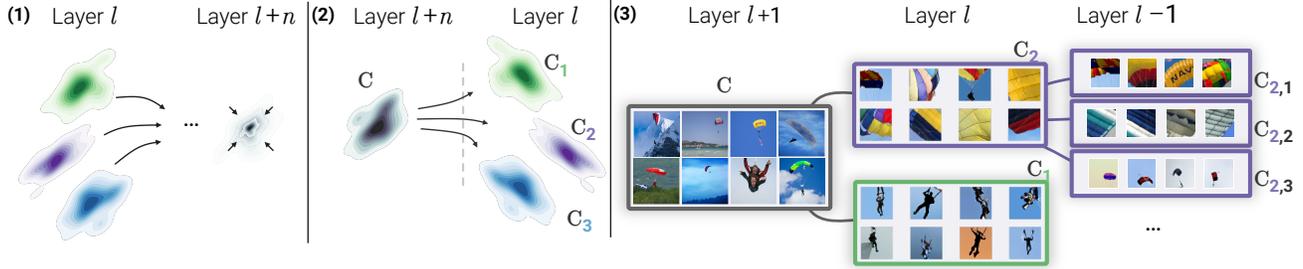


Figure 3. **(1) Neural collapse (amalgamation).** A classifier needs to be able to linearly separate classes by the final layer. It is commonly assumed that in order to achieve this, image activations from the same class get progressively “merged” such that these image activations converge to a one-hot vector associated with the class at the level of the logits layer. In practice, this means that different concepts get ultimately blended together along the way. **(2) Recursive process.** When a concept is not understood (e.g., C), we propose to decompose it into multiple sub-concepts (e.g., C_1, C_2, C_3) using the activations from an earlier layer to overcome the aforementioned neural collapse issue. **(3) Example of recursive concept decomposition** using CRAFT on the ImageNet class “parachute”.

components. The first component is a function g that maps the input to an intermediate state, and the second component is h , which takes this intermediate state to the output, such that $f(x) = (h \circ g)(x)$. In this context, $g(x) \subseteq \mathbb{R}^p$ represents the intermediate activations of x within the network. Further, we will assume non-negative activations: $g(x) \geq 0$. In particular, this assumption is verified by any architecture that utilizes *ReLU*, but any non-negative activation function works.

3.1. Concept activation factorization.

We use Non-negative matrix factorization to identify a basis for concepts based on a network’s activations (Fig.4). Inspired by the approach taken in ACE [24], we will use image sub-regions to try to identify coherent concepts.

The first step involves gathering a set of images that one wishes to explain, such as the dataset, in order to generate associated concepts. In our examples, to explain a specific class $y \in \mathcal{Y}$, we selected the set of points \mathcal{C} from the dataset for which the model’s predictions matched a specific class $\mathcal{C} = \{x_i : f(x_i) = y, 1 \leq i \leq n\}$. It is important to emphasize that this choice is significant. The goal is not to understand how humans labeled the data, but rather to comprehend the model itself. By only selecting correctly classified images, important biases and failure cases may be missed, preventing a complete understanding of our model.

Now that we have defined our set of images, we will proceed with selecting sub-regions of those images to identify specific concepts within a localized context. It has been observed that the implementation of segmentation masks suggested in ACE can lead to the introduction of artifacts due to the associated inpainting with a baseline value. In contrast, our proposed method takes advantage of the prevalent use of modern data augmentation techniques such as randaugment, mixup, and cutmix during the training of current models. These techniques involve the current practice of models be-

ing trained on image crops, which enables us to leverage a straightforward crop and resize function denoted by $\pi(\cdot)$ to create sub-regions (illustrated in Fig.4). By applying π function to each image in the set \mathcal{C} , we obtain an auxiliary dataset $\mathbf{X} \in \mathbb{R}^{n \times d}$ such that each entries $\mathbf{X}_i = \pi(x_i)$ is an image crop.

To discover the concept basis, we start by obtaining the activations for the random crops $\mathbf{A} = g(\mathbf{X}) \in \mathbb{R}^{n \times p}$. In the case where f is a convolutional neural network, a global average pooling is applied to the activations.

We are now ready to apply Non-negative Matrix Factorization (NMF) to decompose positive activations \mathbf{A} into a product of non-negative, low-rank matrices $\mathbf{U} \in \mathbb{R}^{n \times r}$ and $\mathbf{W} \in \mathbb{R}^{p \times r}$ by solving:

$$(\mathbf{U}, \mathbf{W}) = \arg \min_{\mathbf{U} \geq 0, \mathbf{W} \geq 0} \frac{1}{2} \|\mathbf{A} - \mathbf{U}\mathbf{W}^T\|_F^2, \quad (1)$$

where $\|\cdot\|_F$ denotes the Frobenius norm.

This decomposition of our activations \mathbf{A} yields two matrices: \mathbf{W} containing our Concept Activation Vectors (CAVs) and \mathbf{U} that redefines the data points in our dataset according to this new basis. Moreover, this decomposition in this new basis has some interesting properties that go beyond the simple low-rank factorization – since $r \ll \min(n, p)$. First, NMF can be understood as the joint learning of a dictionary of Concept Activation Vectors – called a “concept bank” in Fig. 4 – that maps a \mathbb{R}^p basis onto \mathbb{R}^r , and \mathbf{U} the coefficients of the vectors \mathbf{A} expressed in this new basis. The minimization of the reconstruction error $\frac{1}{2} \|\mathbf{A} - \mathbf{U}\mathbf{W}\|_F^2$ ensures that the new basis contains (mostly) relevant concepts. Intuitively, the non-negativity constraints $\mathbf{U} \geq 0, \mathbf{W} \geq 0$ encourage (i) \mathbf{W} to be sparse (useful for creating disentangled concepts), (ii) \mathbf{U} to be sparse (convenient for selecting a minimal set of useful concepts) and (iii) missing data to be imputed [56], which corresponds to the sparsity pattern of *post-ReLU* activations \mathbf{A} .

It is worth noting that each input x_i can be expressed

as a linear combination of concepts denoted as $\mathbf{A}_i = \sum_{j=1}^r \mathbf{U}_{(i,j)} \mathbf{W}_j^\top$. This approach is advantageous because it allows us to interpret each input as a composition of the underlying concepts. Furthermore, the strict positivity of each term – NMF is working over the anti-negative semiring, – enhances the interpretability of the decomposition. Another interesting interpretation could be that each input is represented as a superposition of concepts [13].

While other methods in the literature solve a similar problem (such as low-rank factorization using SVD or ICA), the NMF is both fast and effective and is known to yield concepts that are meaningful to humans [20, 74, 77]. Finally, once the concept bank \mathbf{W} has been precomputed, we can associate the concept coefficients \mathbf{u} to any new input \mathbf{x} (e.g., a full image) by solving the underlying Non-Negative Least Squares (NNLS) problem $\min_{\mathbf{u} \geq 0} \frac{1}{2} \|\mathbf{g}(\mathbf{x}) - \mathbf{u} \mathbf{W}^\top\|_F^2$, and therefore recover its decomposition in the concept basis.

In essence, the core of our method can be summarized as follows: using a set of images, the idea is to re-interpret their embedding at a given layer as a composition of concepts that humans can easily understand. In the next section, we show how one can recursively apply concept activation factorizations to preceding layer for an image containing a previously computed concept.

3.2. Ingredient 1: A pinch of recursivity

One of the most apparent issues in previous work [24, 77] is the need for choosing a priori a layer at which the activation maps are computed. This choice will critically affect the concepts that are identified because certain concepts get amalgamated [53] into one at different layers of the neural network, resulting in incoherent and indecipherable clusters, as illustrated in Fig. 3. We posit that this can be solved by iteratively applying our decomposition at different layer depths, and for the concepts that remain difficult to understand, by looking for their sub-concepts in earlier layers by isolating the images that contain them. This allows us to build hierarchies of concepts for each class.

We offer a simple solution consisting of reapplying our method to a concept by performing a second step of concept activation factorization on a set of images that contain the concept \mathcal{C} in order to refine it and create sub-concepts (e.g., decompose \mathcal{C} into $\{\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3\}$) see Fig. 3 for an illustrative example. Note that we generalize current methods in the sense that taking images $(\mathbf{x}_1, \dots, \mathbf{x}_n)$ that are clustered in the logits layer (belonging to the same class) and decomposing them in a previous layer – as done in [24, 77] – is a valid recursive step. For a more general case, let us assume that a set of images that contain a common concept is obtained using the first step of concept activation factorization.

We will then take a subset of the auxiliary dataset points to refine any concept j . To do this, we select the subset of points that contain the concept $\mathcal{C}_j = \{\mathbf{X}_i : \mathbf{U}_{(i,j)} >$

$\lambda_j, 1 \leq i \leq n\}$, where λ_j is the 90th percentile of the values of the concept $\mathbf{U}_{(1,j)}, \dots, \mathbf{U}_{(n,j)}$. In other words, the 10% of images that activate the concept j the most are selected for further refinement into sub-concepts. Given this new set of points, we can then re-apply the Concept Matrix Factorization method to an earlier layer $\mathbf{g}'(\cdot)$ to obtain the sub-concepts decomposition from the initial concept – as illustrated in Fig.3.

3.3. Ingredient 2: A dash of sensitivity analysis

A major concern with concept extraction methods is that concepts that makes sense to humans are not necessarily the same as those being used by a model to classify images. In order to prevent such confirmation bias during our concept analysis phase, a faithful estimate the overall importance of the extracted concepts is crucial. Kim et al. [40] proposed an importance estimator based on directional derivatives: the partial derivative of the model output with respect to the vector of concepts. While this measure is theoretically grounded, it relies on the same principle as gradient-based methods, and thus, suffers from the same pitfalls: neural network models have noisy gradients [66, 71]. Hence, the farther the chosen layer is from the output, the noisier the directional derivative score will be.

Since we essentially want to know which concept has the greatest effect on the output of the model, it is natural to consider the field of sensitivity analysis [9, 33, 34, 67, 68]. In this section, we briefly recall the classic “total Sobol indices” and how to apply them to our problem. The complete derivation of the Sobol-Hoeffding decomposition is presented in Section D of the supplementary materials. Formally, a natural way to estimate the importance of a concept i is to measure the fluctuations of the model’s output $\mathbf{h}(\mathbf{U} \mathbf{W}^\top)$ in response to meaningful perturbations of the concept coefficient $\mathbf{U}_{(1,i)}, \dots, \mathbf{U}_{(n,i)}$. Concretely, we will use perturbation masks $\mathbf{M} = (M_1, \dots, M_r) \sim \mathcal{U}([0, 1]^r)$, here an i.i.d sequence of real-valued random variables, we introduce a concept fluctuation to reconstruct a perturbed activation $\tilde{\mathbf{A}} = (\mathbf{U} \odot \mathbf{M}) \mathbf{W}^\top$ where \odot denote the Hadamard product (e.g., the masks can be used to remove a concept by setting its value to zero). We can then propagate this perturbed activation to the model output $\mathbf{Y} = \mathbf{h}(\tilde{\mathbf{A}})$. Simply put, removing or applying perturbation of an important concept will result in a substantial variation in the output, whereas an unused concept will have minimal effect on the output.

Finally, we can capture the importance that a concept might have as a main effect – along with its interactions with other concepts – on the model’s output by calculating the expected variance that would remain if all the concepts except the i were to be fixed. This yields the general definition of the total Sobol indices.

Definition 3.1 (Total Sobol indices). *The total Sobol index*

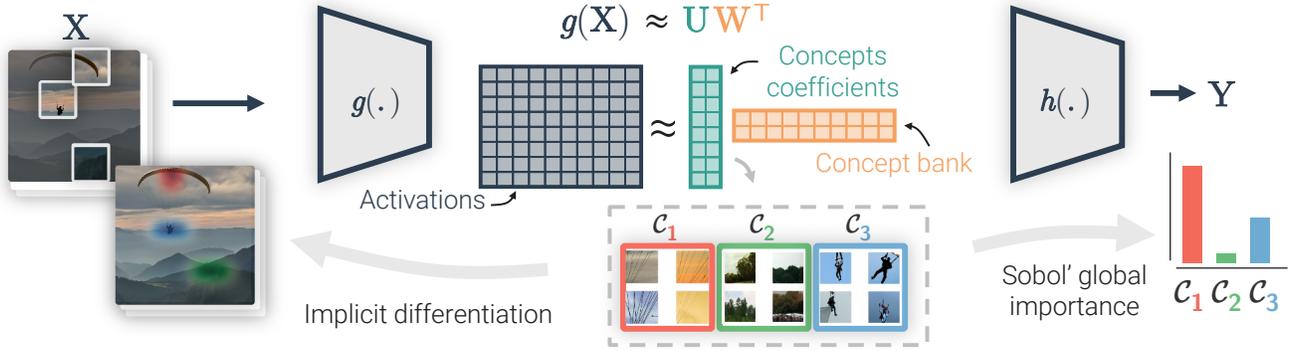


Figure 4. **Overview of CRAFT.** Starting from a set of crops \mathbf{X} containing a concept \mathcal{C} (e.g., crops images of the class “parachute”), we compute activations $g(\mathbf{X})$ corresponding to an intermediate layer from a neural network for random image crops. We then factorize these activations into two lower-rank matrices, (\mathbf{U}, \mathbf{W}) . \mathbf{W} is what we call a “concept bank” and is a new basis used to express the activations, while \mathbf{U} corresponds to the corresponding coefficients in this new basis. We then extend the method with 3 new ingredients: (1) recursivity – by proposing to re-decompose a concept (e.g., take a new set of images containing \mathcal{C}_1) at an earlier layer, (2) a better importance estimation using Sobol indices and (3) an approach to leverage implicit differentiation to generate *concept attribution maps* to localize concepts in an image.

S_i^T , which measures the contribution of a concept i as well as its interactions of any order with any other concepts to the model output variance, is given by:

$$S_i^T = \frac{\mathbb{E}_{\mathbf{M} \sim i}(\mathbb{V}_{M_i}(\mathbf{Y}|\mathbf{M}_{\sim i}))}{\mathbb{V}(\mathbf{Y})} \quad (2)$$

$$= \frac{\mathbb{E}_{\mathbf{M} \sim i}(\mathbb{V}_{M_i}(\mathbf{h}((\mathbf{U} \odot \mathbf{M})\mathbf{W}^T)|\mathbf{M}_{\sim i}))}{\mathbb{V}(\mathbf{h}((\mathbf{U} \odot \mathbf{M})\mathbf{W}^T))}. \quad (3)$$

In practice, this index can be calculated very efficiently [36, 48, 52, 58, 72], more details on the Quasi-Monte Carlo sampling and the estimator used are left in appendix D.

3.4. Ingredient 3: A smidgen of implicit differentiation

Attribution methods are useful for determining the regions deemed important by a model for its decision, but they lack information about what exactly triggered it. We have seen that we can already extract this information from the matrices \mathbf{U} and \mathbf{W} , but as it is, we do not know in what part of an image a given concept is represented. In this section, we will show how we can leverage attribution methods (forward and backward modes) to find where a concept is located in the input image (see Fig. 2). Forward attribution methods do not rely on any gradient computation as they only use inference processes, whereas backward methods require back-propagating through a network’s layers. By application of the chain rule, computing $\partial\mathbf{U}/\partial\mathbf{A}$ requires access to $\partial\mathbf{U}/\partial\mathbf{A}$.

To do so, one could be tempted to solve the linear system $\mathbf{U}\mathbf{W}^T = \mathbf{A}$. However, this problem is ill-posed since \mathbf{W}^T is low rank. A standard approach is to calculate the Moore-Penrose pseudo-inverse $(\mathbf{W}^T)^\dagger$, which solves rank

deficient systems by looking at the minimum norm solution [2]. In practice, $(\mathbf{W}^T)^\dagger$ is computed with the Singular Value Decomposition (SVD) of \mathbf{W}^T . Unfortunately, SVD is also the solution to the *unstructured minimization* of $\frac{1}{2}\|\mathbf{A} - \mathbf{U}\mathbf{W}^T\|_F^2$ by the Eckart-Young-Mirsky theorem [12]. Hence, the non-negativity constraints of the NMF are ignored, which prevents such approaches from succeeding. Other issues stem from the fact that the \mathbf{U}, \mathbf{W} decomposition is generally not unique.

Our third contribution consists of tackling this problem to allow the use of attribution methods, i.e., *concept attribution maps*, by proposing a strategy to differentiate through the NMF block.

Implicit differentiation of NMF block The NMF problem 1 is NP-hard [73], and it is not convex with respect to the input pair (\mathbf{U}, \mathbf{W}) . However, fixing the value of one of the two factors and optimizing the other turns the NMF formulation into a pair of Non-Negative Least Squares (NNLS) problems, which are convex. This ensures that alternating minimization (a standard approach for NMF) of (\mathbf{U}, \mathbf{W}) factors will eventually reach a local minimum. Each of this alternating NNLS problems fulfills the Karush–Kuhn–Tucker (KKT) conditions [38, 45], which can be encoded in the so-called *optimality function* \mathbf{F} from [4], see Eq. 10 Appendix C.2. The implicit function theorem [25] allows us to use implicit differentiation [3, 25, 44] to efficiently compute the Jacobians $\partial\mathbf{U}/\partial\mathbf{A}$ and $\partial\mathbf{W}/\partial\mathbf{A}$ without requiring to back-propagate through each of the iterations of the NMF solver:

$$\frac{\partial(\mathbf{U}, \mathbf{W}, \bar{\mathbf{U}}, \bar{\mathbf{W}})}{\partial\mathbf{A}} = -(\partial_1\mathbf{F})^{-1}\partial_2\mathbf{F}. \quad (4)$$

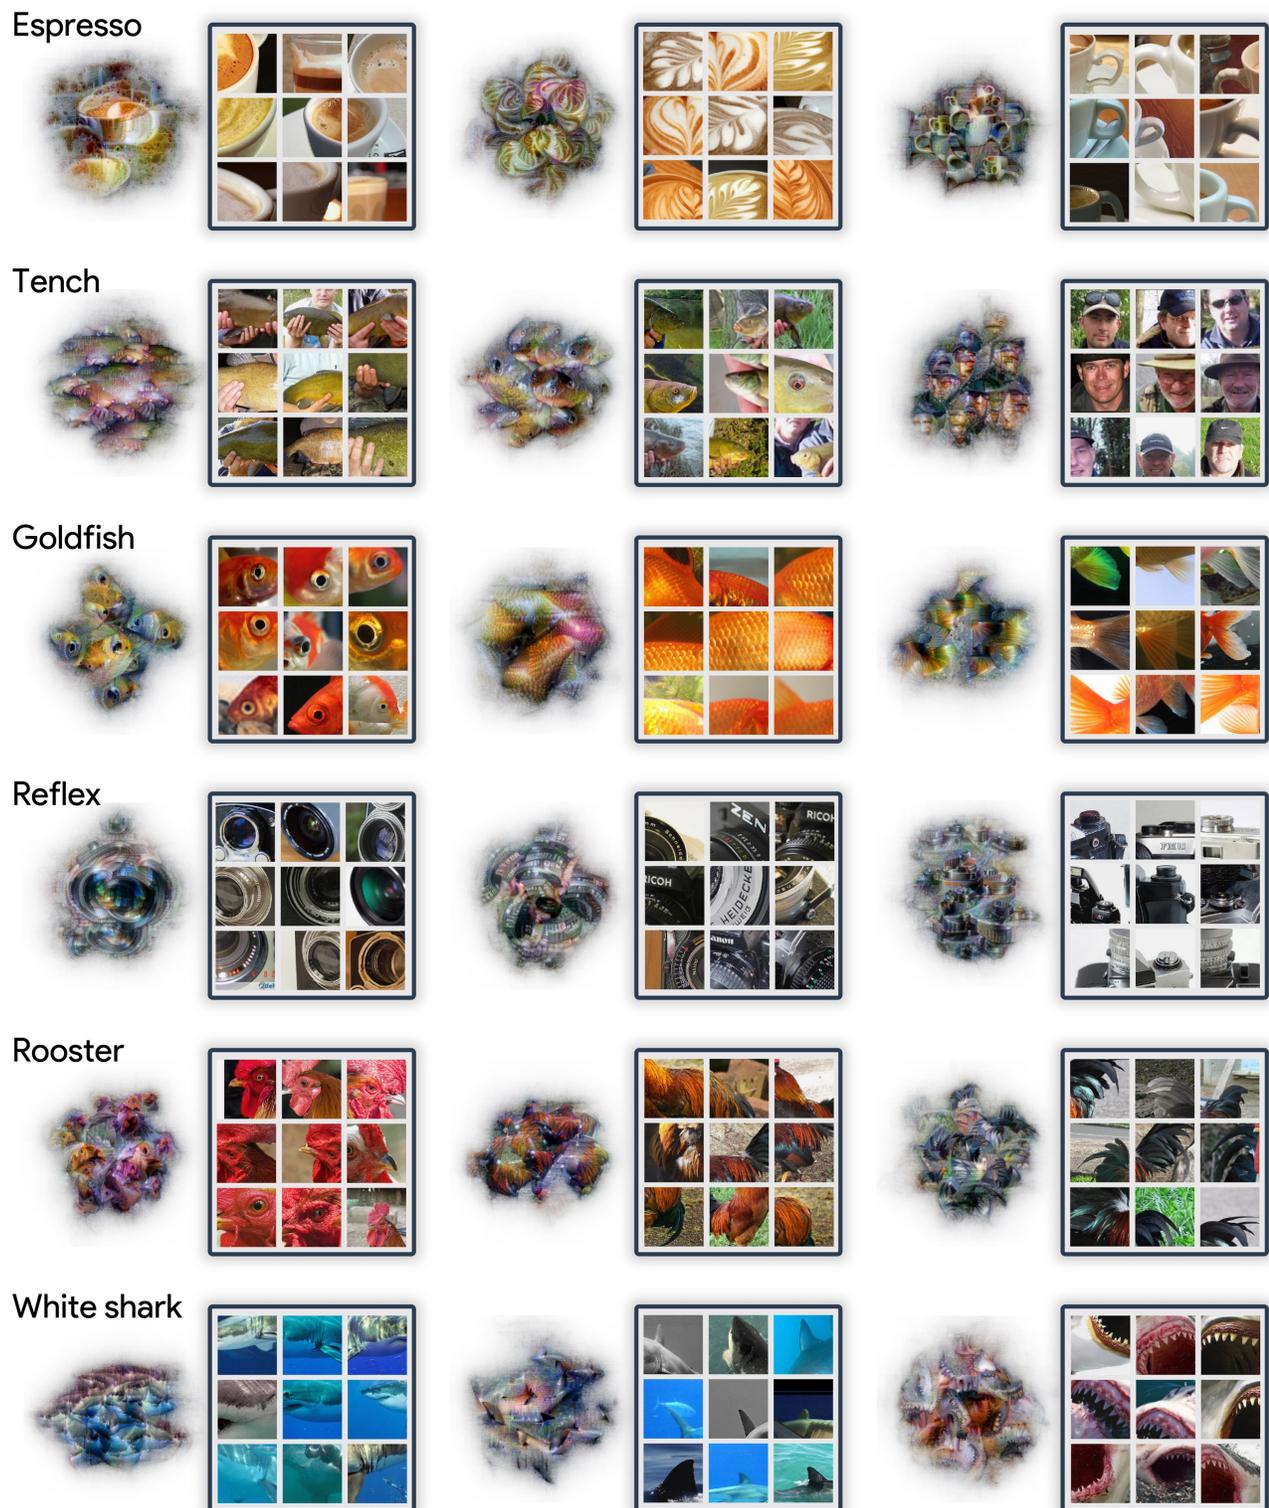


Figure 5. **Qualitative Results:** CRAFT results on 6 classes of ILSVRC2012 [10] for a trained ResNet50V2. The results showcase the top 3 most important concepts for each class. This is done by displaying crop images that activate the concept the most (using **U**) and also feature visualization [51] of the associated CAVs (using **W**).

However, this requires the dual variables $\bar{\mathbf{U}}$ and $\bar{\mathbf{W}}$, which are not computed in scikit-learn’s [54] popular implementation. Consequently, we leverage the work of [32] and we re-implement our own solver with Jaxopt [4] based on ADMM [5], a GPU friendly algorithm (see Appendix C.2).

Concretely, given our concepts bank \mathbf{W} , the concept attribution maps of a new input \mathbf{x} are calculated by solving the NNLS problem $\min_{\mathbf{U} \geq 0} \frac{1}{2} \|\mathbf{g}(\mathbf{x}) - \mathbf{U}\mathbf{W}^T\|_F^2$. The implicit differentiation of the NMF block $\partial\mathbf{U}/\partial\mathbf{A}$ is integrated into the classic back-propagation to obtain $\partial\mathbf{U}/\partial\mathbf{x}$. Most interestingly, this technical advance enables the use of all white-box explainability methods [59, 62, 66, 69, 71, 78] to generate concept-wise attribution maps and trace the part of an image that triggered the detection of the concept by the network. Additionally, it is even possible to employ black-box methods [14, 47, 55, 57] since it only amounts to solving an NNLS problem.

4. Experimental evaluation

In order to evaluate the interest and the benefits brought by CRAFT, we start in Section 4.1 by assessing the practical utility of the method on a human-centered benchmark composed of 3 XAI scenarios.

After demonstrating the usefulness of the method using these human experiments, we independently validate the 3 proposed ingredients. First, we provide evidence that recursivity allows refining concepts, making them more meaningful to humans using two additional human experiments in Section 4.2. Next, we evaluate our new Sobol estimator and show quantitatively that it provides a more faithful assessment of concept importance in Section 4.3. Finally, we run an ablation experiment that measures the interest of local explanations based on concept attribution maps coupled with global explanations. Additional experiments, including a sanity check and an example of deep dreams applied on the concept bank, as well as many other examples of local explanations for randomly picked images from ILSVRC2012, are included in Section B of the supplementary materials. We leave the discussion on the limitations of this method and on the broader impact in appendix A.

4.1. Utility Evaluation

As emphasized by Doshi-Velez et al. [11], the goal of XAI should be to develop methods that help a user better understand the behavior of deep neural network models. An instantiation of this idea was recently introduced by Colin & Fel et al. [8] who described an experimental framework to quantitatively measure the practical usefulness of explainability methods in real-world scenarios. For their initial setup, these authors recruited $n = 1, 150$ online participants (evaluated over 8 unique conditions and 3 AI scenarios) –

Scikit-learn uses a block coordinate descent algorithm [7, 17], with a randomized SVD initialization.

making it the largest benchmark to date in XAI. Here, we follow their framework rigorously to allow for the robust evaluation of the utility of our proposed CRAFT method and the related ACE. The 3 representative real-world scenarios are: (1) identifying bias in an AI system (using Husky vs Wolf dataset from [57]), (2) characterizing the visual strategy that are too difficult for an untrained non-expert human observer (using the Paleobotanical dataset from [75]), (3) understanding complex failure cases (using ImageNet “Red fox” vs “Kit fox” binary classification). Using this benchmark, we evaluate CRAFT, ACE, as well as CRAFT with only the global concepts (CRAFTCO) to allow for a fair comparison with ACE. To the best of our knowledge, we are the first to systematically evaluate concept-based methods against attribution methods.

Results are shown in Table 1 and demonstrate the benefit of CRAFT, which achieves higher scores than all of the attribution methods tested as well as ACE in the first two scenarios. To date, no method appears to exceed the baseline on the third scenario suggesting that additional work is required. We also note that, in the first two scenarios, CRAFTCO is one of the best-performing methods and it always outperforms ACE – meaning that even without the local explanation of the concept attribution maps, CRAFT largely outperforms ACE. Examples of concepts produced by CRAFT are shown in the Appendix E.1.

4.2. Validation of Recursivity

To evaluate the meaningfulness of the extracted high-level concepts, we performed psychophysics experiments with human subjects, whom we asked to answer a survey in two phases. Furthermore, we distinguished two different audiences: on the one hand, experts in machine learning, and on the other hand, people with no particular knowledge of computer vision. Both groups of participants were volunteers and did not receive any monetary compensation. Some examples of the developed interface are available the appendix E. It is important to note that this experiment was carried out independently from the utility evaluation and thus it was setup differently.

Intruder detection experiment First, we ask users to identify the intruder out of a series of five image crops belonging to a certain class, with the odd one being taken from a different concept but still from the same class. Then, we compare the results of this intruder detection with another intruder detection, this time, using a concept (e.g., \mathcal{C}_1) coming from a layer l and one of its sub-concepts (e.g., \mathcal{C}_{12} in Fig.3) extracted using our recursive method. If the concept (or sub-concept) is coherent, then it should be easy for the users to find the intruder. Table 2 summarizes our results, showing that indeed both concepts and sub-concepts are coherent, and that recursivity can lead to a slightly higher understanding of the generated concepts (significant for non-experts, but

Session n ^o	<i>Husky vs. Wolf</i>				<i>Leaves</i>				<i>“Kit Fox” vs “Red Fox”</i>				
	1	2	3	<i>Utility</i>	1	2	3	<i>Utility</i>	1	2	3	<i>Utility</i>	
Baseline	55.7	66.2	62.9		70.1	76.8	78.6		58.8	62.2	58.8		
Control	53.3	61.0	61.4	0.95	72.0	78.0	80.2	1.02	60.7	59.2	48.5	0.94	
Attributions	Saliency [62]	53.9	69.6	73.3	1.06	83.2	88.7	82.4	<u>1.13</u>	61.7	60.2	58.2	1.00
	Integ.-Grad. [71]	67.4	72.8	73.2	1.15	82.5	82.5	85.3	1.11	59.4	58.3	58.3	0.98
	SmoothGrad [66]	68.7	75.3	78.0	1.20	83.0	85.7	86.3	<u>1.13</u>	50.3	55.0	61.4	0.93
	GradCAM [59]	77.6	85.7	84.1	1.34	81.9	83.5	82.4	1.10	54.4	52.5	54.1	0.90
	Occlusion [76]	71.0	75.7	78.1	1.22	78.8	86.1	82.9	1.10	51.0	60.2	55.1	0.92
	Grad.-Input [61]	65.8	63.3	67.9	1.06	76.5	82.9	79.5	1.05	50.0	57.6	62.6	0.95
Concepts	ACE [24]	68.8	71.4	72.7	1.15	79.8	73.8	82.1	1.05	48.4	46.5	46.1	0.78
	CRAFTCO (ours)	82.4	87.0	85.1	<u>1.38</u>	78.8	85.5	89.4	1.12	55.5	49.5	53.3	0.88
	CRAFT (ours)	90.6	97.3	95.5	1.53	86.2	86.6	85.5	1.15	56.5	50.6	49.4	0.87

Table 1. **Utility scores on 3 datasets from [8].** Their *Utility* benchmark evaluates how well explanations help users identify general rules driving classifications that readily transfer to unseen instances. At training time, users are asked to infer rules driving the decisions of the model given a set of images, and their associated predictions and explanations. At test time, the *Utility* metric measures the accuracy of users at predicting the model decision on novel images averaged over 3 sessions, and normalized by the baseline accuracy of users trained without explanations. The higher the *Utility* score, the more useful the explanation, and the more crucial the information provided is for understanding –and thus predicting the model’s output– on novel samples. CRAFTCO stands for “CRAFT Concept Only” and designates an experimental condition where only global concepts are given to users, without local explanations (i.e., the concept attribution maps). The first and second best results above the baseline are in **bold** and underlined, respectively.

	Experts ($n = 36$)	Laymen ($n = 37$)
<i>Intruder</i>		
Acc. Concept	70.19%	61.08%
Acc. Sub-Concept	74.81% ($p = 0.18$)	67.03% ($p = 0.043$)
<i>Binary choice</i>		
Sub-Concept	76.1% ($p < 0.001$)	74.95% ($p < 0.001$)
Odds Ratios	3.53	2.99

Table 2. **Results from the psychophysics experiments to validate the recursivity ingredient.**

not for experts) and might suggest a way to make concepts more interpretable.

Binary choice experiment In order to test the improvement of coherence of the sub-concept generated by recursivity with respect to the larger parent concept, we showed participants an image crop belonging to both a subcluster and a parent cluster (e.g., $\pi(\mathbf{x}) \in \mathcal{C}_{11} \subset \mathcal{C}_1$) and asked them which of the two clusters (i.e., \mathcal{C}_{11} or \mathcal{C}_1) seemed to accommodate the image the best. If our hypothesis is correct, then the concept refinement brought by recursivity should help form more coherent clusters. The results in Table 2 are satisfying since in both the expert and non-expert groups, the participants chose the sub-cluster more than 74% of the time. We measure the significance of our results by fitting a binomial logistic regression to our data, and we find that both groups are more likely to choose the sub-concept cluster (at a $p < 0.001$).

4.3. Fidelity analysis

We propose to simultaneously verify that identified concepts are faithful to the model and that the concept importance estimator performs better than that used in TCAV [40] by using the fidelity metrics introduced in [24, 77]. These metrics are similar to the ones used for attribution methods, which consist of studying the change of the logit score when removing/adding pixels considered important. Here, we do not introduce these perturbations in the pixel space but in the concept space: once \mathbf{U} and \mathbf{W} are computed, we reconstruct the matrix $\mathbf{A} \approx \mathbf{U}\mathbf{W}^T$ using only the most important concept (or removing the most important concept for deletion) and compute the resulting change in the output of the model. As can be seen from Fig. 6, ranking the extracted concepts using Sobol’s importance score results in steeper curves than when they are sorted by their TCAV scores. We confirm that these results generalize with other matrix factorization techniques (PCA, ICA, RCA) in Section F of the Appendix.

5. Conclusion

In this paper, we introduced CRAFT, a method for automatically extracting human-interpretable concepts from deep networks. Our method aims to explain a pre-trained model’s decisions both on a per-class and per-image basis by highlighting both “what” the model saw and “where” it saw it – with complementary benefits. The approach relies on 3 novel ingredients: 1) a recursive formulation of concept discovery to identify the correct level of granularity for which individual concepts are understandable; 2) a novel method for measuring concept importance through Sobol indices to

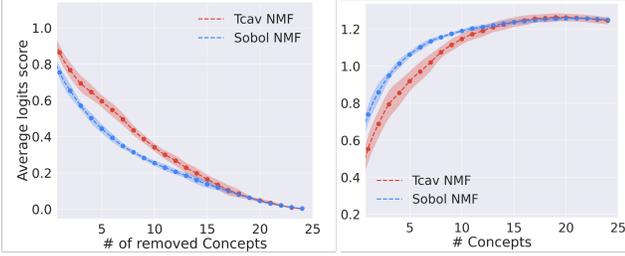


Figure 6. **(Left)** Deletion curves (lower is better). **(Right)** Insertion curves (higher is better). For both the deletion or insertion metrics, Sobol indices lead to better estimates (calculated on >100K images) of important concepts.

more accurately identify which concepts influence a model’s decision for a given class; and 3) the use of implicit differentiation methods to backpropagate through non-negative matrix factorization (NMF) blocks to allow the generation of concept-wise local explanations or *concept attribution maps* independently of the attribution method used. Using a recently introduced human-centered utility benchmark, we conducted psychophysics experiments to confirm the validity of the approach: and that the concepts identified by CRAFT are useful and meaningful to human experimenters. We hope that this work will guide further efforts in the search for concept-based explainability methods.

Acknowledgements

This work was conducted as part of the DEEL project. Funding was provided by ANR-3IA Artificial and Natural Intelligence Toulouse Institute (ANR-19-PI3A-0004). Additional support provided by ONR grant N00014-19-1-2029 and NSF grant IIS-1912280. Support for computing hardware provided by Google via the TensorFlow Research Cloud (TFRC) program and by the Center for Computation and Visualization (CCV) at Brown University (NIH grant S10OD025181).

References

- [1] Julius Adebayo, Justin Gilmer, Michael Muehly, Ian Goodfellow, Moritz Hardt, and Been Kim. Sanity checks for saliency maps. *Advances in neural information processing systems*, 31, 2018. 2, 3, 25
- [2] João Carlos Alves Barata and Mahir Saleh Hussein. The moore–penrose pseudoinverse: A tutorial review of the theory. *Brazilian Journal of Physics*, 42(1):146–165, 2012. 6
- [3] Bradley M Bell and James V Burke. Algorithmic differentiation of implicit functions and optimal values. In *Advances in Automatic Differentiation*, pages 67–77. Springer, 2008. 6, 20
- [4] Mathieu Blondel, Quentin Berthet, Marco Cuturi, Roy Frostig, Stephan Hoyer, Felipe Linares-López, Fabian Pedregosa, and Jean-Philippe Vert. Efficient and modular implicit differentiation. *arXiv preprint arXiv:2105.15183*, 2021. 3, 6, 8, 20
- [5] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, Jonathan Eckstein, et al. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011. 8, 15, 19, 20
- [6] Chaofan Chen, Oscar Li, Daniel Tao, Alina Barnett, Cynthia Rudin, and Jonathan K Su. This looks like that: deep learning for interpretable image recognition. *Advances in neural information processing systems*, 32, 2019. 15
- [7] Andrzej Cichocki and Anh-Huy Phan. Fast local algorithms for large scale nonnegative matrix and tensor factorizations. *IEICE transactions on fundamentals of electronics, communications and computer sciences*, 92(3):708–721, 2009. 8, 20
- [8] Julien Colin, Thomas Fel, Rémi Cadène, and Thomas Serre. What i cannot predict, i do not understand: A human-centered evaluation framework for explainability methods. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 2, 3, 8, 9, 22
- [9] RI Cukier, CM Fortuin, Kurt E Shuler, AG Petschek, and J Ho Schaibly. Study of the sensitivity of coupled reaction systems to uncertainties in rate coefficients. i theory. *The Journal of chemical physics*, 59(8):3873–3878, 1973. 5
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 2, 7, 16, 24
- [11] Finale Doshi-Velez and Been Kim. Towards a rigorous science of interpretable machine learning. *ArXiv e-print*, 2017. 2, 3, 8
- [12] Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936. 6
- [13] 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. Toy models of superposition. *Transformer Circuits Thread*, 2022. 5
- [14] Thomas Fel, Rémi Cadène, Mathieu Chalvidal, Matthieu Cord, David Vigouroux, and Thomas Serre. Look at the variance! efficient black-box explanations with sobol-based sensitivity analysis. *Advances in Neural Information Processing Systems*, 34, 2021. 2, 3, 8, 20
- [15] Thomas Fel, Mélanie Ducoffe, David Vigouroux, Rémi Cadène, Mikael Capelle, Claire Nicodème, and Thomas Serre. Don’t lie to me! robust and efficient explainability with verified perturbation analysis. *arXiv preprint arXiv:2202.07728*, 2022. 3
- [16] Thomas Fel, Lucas Hervier, David Vigouroux, Antonin Poche, Justin Plakoo, Remi Cadene, Mathieu Chalvidal, Julien Colin, Thibaut Boissin, Louis Béthune, Agustin Picard, Claire Nicodeme, Laurent Gardes, Gregory Flandin, and Thomas Serre. Xplique: A deep learning explainability toolbox. *Workshop, Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 18
- [17] Cédric Févotte and Jérôme Idier. Algorithms for nonnegative matrix factorization with the β -divergence. *Neural computation*, 23(9):2421–2456, 2011. 8, 20
- [18] Ruth C Fong and Andrea Vedaldi. Interpretable explanations of black boxes by meaningful perturbation. In *Proceedings of the IEEE international conference on computer vision*, pages 3429–3437, 2017. 3
- [19] Ruth C Fong and Andrea Vedaldi. Interpretable explanations of black boxes by meaningful perturbation. In *Proceedings of the IEEE international conference on computer vision*, pages 3429–3437, 2017. 15
- [20] Xiao Fu, Kejun Huang, Nicholas D Sidiropoulos, and Wing-Kin Ma. Nonnegative matrix factorization for signal and data analytics: Identifiability, algorithms, and applications. *IEEE Signal Process. Mag.*, 36(2):59–80, 2019. 5
- [21] Yunhao Ge, Yao Xiao, Zhi Xu, Meng Zheng, Srikrishna Karanam, Terrence Chen, Laurent Itti, and Ziyang Wu. A peek into the reasoning of neural networks: Interpreting with structural visual concepts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2195–2204, 2021. 15
- [22] Sahra Ghalebikesabi, Lucile Ter-Minassian, Karla DiazOrdaz, and Chris C Holmes. On locality of local explanation models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 3
- [23] Amirata Ghorbani, Abubakar Abid, and James Zou. Interpretation of neural networks is fragile. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2017. 3
- [24] Amirata Ghorbani, James Wexler, James Y Zou, and Been Kim. Towards automatic concept-based explanations. *Advances in Neural Information Processing Systems*, 32, 2019. 2, 3, 4, 5, 9, 15, 16, 24
- [25] Andreas Griewank and Andrea Walther. *Evaluating derivatives: principles and techniques of algorithmic differentiation*. SIAM, 2008. 3, 6, 20
- [26] Peter Hase and Mohit Bansal. Evaluating explainable ai: Which algorithmic explanations help users predict model behavior? In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL Short Papers)*, 2020. 2, 3
- [27] Peter Hase, Harry Xie, and Mohit Bansal. The out-of-distribution problem in explainability and search methods for feature importance explanations. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 2

- [28] Johannes Haug, Stefan Zürn, Peter El-Jiz, and Gjergji Kasneci. On baselines for local feature attributions. *arXiv preprint arXiv:2101.00905*, 2021. 3
- [29] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016. 24
- [30] Magnus R Hestenes and Eduard Stiefel. Methods of conjugate gradients for solving. *Journal of research of the National Bureau of Standards*, 49(6):409, 1952. 19, 20
- [31] Cheng-Yu Hsieh, Chih-Kuan Yeh, Xuanqing Liu, Pradeep Ravikumar, Seungyeon Kim, Sanjiv Kumar, and Cho-Jui Hsieh. Evaluations and methods for explanation through robustness analysis. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021. 3
- [32] Kejun Huang, Nicholas D Sidiropoulos, and Athanasios P Liavas. A flexible and efficient algorithmic framework for constrained matrix and tensor factorization. *IEEE Transactions on Signal Processing*, 64(19):5052–5065, 2016. 8, 19
- [33] Marouane El Idrissi, Vincent Chabridon, and Bertrand Iooss. Developments and applications of shapley effects to reliability-oriented sensitivity analysis with correlated inputs. *Environmental Modelling & Software*, 2021. 5
- [34] Bertrand Iooss and Paul Lemaître. A review on global sensitivity analysis methods. In *Uncertainty management in simulation-optimization of complex systems*, pages 101–122. Springer, 2015. 3, 5
- [35] Alon Jacovi, Ana Marasović, Tim Miller, and Yoav Goldberg. Formalizing trust in artificial intelligence: Prerequisites, causes and goals of human trust in ai. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 624–635, 2021. 2
- [36] Alexandre Janon, Thierry Klein, Agnes Lagnoux, Maëlle Nodet, and Clémentine Prieur. Asymptotic normality and efficiency of two sobol index estimators. *ESAIM: Probability and Statistics*, 18:342–364, 2014. 6, 21
- [37] Margot E Kaminski and Jennifer M Urban. The right to contest ai. *Columbia Law Review*, 121(7):1957–2048, 2021. 2
- [38] William Karush. Minima of functions of several variables with inequalities as side constraints. *M. Sc. Dissertation. Dept. of Mathematics, Univ. of Chicago*, 1939. 6, 19
- [39] Been Kim, Rajiv Khanna, and Oluwasanmi O Koyejo. Examples are not enough, learn to criticize! criticism for interpretability. *Advances in Neural Information Processing Systems (NeurIPS)*, 2016. 3
- [40] Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viegas, et al. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). In *International conference on machine learning*, pages 2668–2677. PMLR, 2018. 2, 3, 5, 9, 18, 24
- [41] Sunnie S. Y. Kim, Nicole Meister, Vikram V. Ramaswamy, Ruth Fong, and Olga Russakovsky. HIVE: Evaluating the human interpretability of visual explanations. In *European Conference on Computer Vision (ECCV)*, 2022. 2, 3
- [42] Pieter-Jan Kindermans, Sara Hooker, Julius Adebayo, Maximilian Alber, Kristof T. Schütt, Sven Dähne, Dumitru Erhan, and Been Kim. *The (Un)reliability of Saliency Methods*, pages 267–280. Springer International Publishing, Cham, 2019. 3
- [43] Mauritz Kop. Eu artificial intelligence act: The european approach to ai. In *Stanford - Vienna Transatlantic Technology Law Forum, Transatlantic Antitrust and IPR Developments, Stanford University, Issue No. 2/2021*. <https://law.stanford.edu/publications/eu-artificial-intelligence-act-the-european-approach-to-ai/>. Stanford-Vienna Transatlantic Technology Law Forum, Transatlantic Antitrust . . . , 2021. 2
- [44] Steven George Krantz and Harold R Parks. *The implicit function theorem: history, theory, and applications*. Springer Science & Business Media, 2002. 3, 6, 20
- [45] Harold W Kuhn and Albert W Tucker. Nonlinear programming proceedings of the second berkeley symposium on mathematical statistics and probability. *Neyman*, pages 481–492, 1951. 6, 19
- [46] Daniel D Lee and H Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999. 2
- [47] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30, 2017. 3, 8
- [48] Amandine Marrel, Bertrand Iooss, Beatrice Laurent, and Olivier Roustant. Calculations of sobol indices for the gaussian process metamodel. *Reliability Engineering & System Safety*, 94(3):742–751, 2009. 6
- [49] Giang Nguyen, Daeyoung Kim, and Anh Nguyen. The effectiveness of feature attribution methods and its correlation with automatic evaluation scores. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 2, 3
- [50] Paul Novello, Thomas Fel, and David Vigouroux. Making sense of dependence: Efficient black-box explanations using dependence measure. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 2, 3
- [51] Chris Olah, Alexander Mordvintsev, and Ludwig Schubert. Feature visualization. *Distill*, 2017. <https://distill.pub/2017/feature-visualization>. 7, 15, 18
- [52] Art B Owen. Better estimation of small sobol’sensitivity indices. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 23(2):1–17, 2013. 6
- [53] Vardan Papyan, XY Han, and David L Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40):24652–24663, 2020. 5
- [54] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *Journal of machine learning research*, 12(Oct):2825–2830, 2011. 8
- [55] Vitali Petsiuk, Abir Das, and Kate Saenko. Rise: Randomized input sampling for explanation of black-box models. *arXiv preprint arXiv:1806.07421*, 2018. 1, 3, 8, 20, 24
- [56] Bin Ren, Laurent Pueyo, Christine Chen, Élodie Choquet, John H Debes, Gaspard Duchêne, François Ménard, and Marshall D Perrin. Using data imputation for signal separation in high-contrast imaging. *The Astrophysical Journal*, 892(2):74, 2020. 4

- [57] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. "why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144, 2016. 3, 8, 20
- [58] Andrea Saltelli, Paola Annoni, Ivano Azzini, Francesca Campolongo, Marco Ratto, and Stefano Tarantola. Variance based sensitivity analysis of model output. design and estimator for the total sensitivity index. *Computer physics communications*, 181(2):259–270, 2010. 6
- [59] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017. 2, 3, 8, 9
- [60] Hua Shen and Ting-Hao Huang. How useful are the machine-generated interpretations to general users? a human evaluation on guessing the incorrectly predicted labels. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 8, pages 168–172, 2020. 2, 3
- [61] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating activation differences. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2017. 2, 9
- [62] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *In Workshop at International Conference on Learning Representations*. Citeseer, 2014. 2, 3, 8, 9
- [63] Leon Sixt, Maximilian Granz, and Tim Landgraf. When explanations lie: Why many modified bp attributions fail. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2020. 2
- [64] Leon Sixt, Martin Schuessler, Oana-Iuliana Popescu, Philipp Weiß, and Tim Landgraf. Do users benefit from interpretable vision? a user study, baseline, and dataset. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2022. 2
- [65] Dylan Slack, Sophie Hilgard, Emily Jia, Sameer Singh, and Himabindu Lakkaraju. Fooling lime and shap: Adversarial attacks on post hoc explanation methods. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pages 180–186, 2020. 2, 3
- [66] Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. Smoothgrad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825*, 2017. 2, 3, 5, 8, 9
- [67] Ilya M Sobol. Sensitivity analysis for non-linear mathematical models. *Mathematical modelling and computational experiment*, 1:407–414, 1993. 3, 5, 21
- [68] Ilya M Sobol. Global sensitivity indices for nonlinear mathematical models and their monte carlo estimates. *Mathematics and computers in simulation*, 55(1-3):271–280, 2001. 3, 5
- [69] Jost Tobias Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin Riedmiller. Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806*, 2014. 2, 3, 8
- [70] Pascal Sturmfels, Scott Lundberg, and Su-In Lee. Visualizing the impact of feature attribution baselines. *Distill*, 2020. 3, 15
- [71] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *International conference on machine learning*, pages 3319–3328. PMLR, 2017. 2, 3, 5, 8, 9
- [72] Stefano Tarantola, Debora Gatelli, and Thierry Alex Mara. Random balance designs for the estimation of first order global sensitivity indices. *Reliability Engineering & System Safety*, 91(6):717–727, 2006. 6
- [73] Stephen A Vavasis. On the complexity of nonnegative matrix factorization. *SIAM Journal on Optimization*, 20(3):1364–1377, 2010. 6
- [74] Yu-Xiong Wang and Yu-Jin Zhang. Nonnegative matrix factorization: A comprehensive review. *IEEE Transactions on Knowledge and Data Engineering*, 25(6):1336–1353, 2013. 5
- [75] Peter Wilf, Shengping Zhang, Sharat Chikkerur, Stefan A Little, Scott L Wing, and Thomas Serre. Computer vision cracks the leaf code. *Proceedings of the National Academy of Sciences*, 113(12):3305–3310, 2016. 8
- [76] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer, 2014. 2, 3, 9
- [77] Ruihan Zhang, Prashan Madumal, Tim Miller, Krista A Ehinger, and Benjamin IP Rubinstein. Invertible concept-based explanations for cnn models with non-negative concept activation vectors. *arXiv preprint arXiv:2006.15417*, 2020. 2, 3, 5, 9, 24
- [78] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016. 3, 8

A Limitations and broader impact	15	C.2. Implicit differentiation	19
A.1. Limitations	15	D Sobol indices for concepts	20
A.2. Broader impact	15	E Human experiments	22
B More results of CRAFT	15	E.1. Utility evaluation	22
B.1. Qualitative comparison with ACE	15	E.2. Validation of Recursivity	22
B.2. Most important concepts.	15	F. Fidelity experiments	24
B.3. Feature Visualization validation	18	G Sanity Check	25
C Backpropagating through the NMF block	19		
C.1. Alternating Direction Method of Multipliers (ADMM) for NMF	19		

A. Limitations and broader impact

A.1. Limitations

Although we believe concept-based XAI to be a promising research direction, it isn't without pitfalls. It is capable of producing explanations that are ideally easy to understand by humans, but to what extent is a question that remains unanswered. The fact that there is no way to mathematically measure this prevents researchers from easily comparing the different techniques in the literature other than through time consuming and expensive experiments with human subjects. We think that developing a metric should be one of the field's priorities.

With CRAFT, we address the question of *what* by showing a cluster of the images that better represent each concept. However, we recognize that it's not perfect: in some cases, concepts are difficult to clearly define – put a label on what it represents –, and might induce some confirmation and selection bias. Feature visualization [51] might help in better illustrating the specific concept (as done in appendix B.3), but we believe there's still space for improvement. For instance, an interesting idea could be to leverage image captioning methods to describe the clusters of image crops, as textual information could help humans in better understanding clusters.

Although we believe CRAFT to be a considerable step in the good direction for the field of concept-based XAI, it also have some pitfalls. Namely, we chose the NMF as the activation factorization, which, while drastically improving the quality of extracted concepts, also comes with it's own caveats. For instance, it is known to be NP-hard to compute exactly, and in order to make it scalable, we had to use a tractable approximation by alternating the optimization of \mathbf{U} and \mathbf{W} through ADMM [5]. This approach might indeed yield non-unique solutions. Our experiments (section 4.3), have shown a low variance on between the runs, which comforts us about the stability of our results. However the absence of formal guarantee for uniqueness must be kept in mind: this subject is still an active topic of research and improvement could be expected in the near future. Namely, sparsity constraints and regularization seem to be promising paths. Naturally, we also need enough samples of the class under study to be available for the factorization to construct a relevant concept bank, which might affect the quality of the explanations on frugal applications where data is very

scarce.

A.2. Broader impact

We do hope that CRAFT helps in the transition to more human-understandable ways of explaining neural network models. It's capacity to find easily understandable concepts inside complex architectures and providing an indication of *where* they are located in the image is – to the best of our knowledge – unmatched. We also think that this method's structure is a step towards reducing confirmation bias: for instance dataset's labels are never used in this method, only the model's predictions. Without claiming to remove confirmation bias, the method focuses on what *the model sees* rather than what *we expect the model to see*. We believe this can help end-users build trust on computer vision models, and at the same time, provide ML practitioners with insights into potential sources of bias in the dataset (e.g. the ski pants in the astronaut/shovel example). Other methods in the literature obtaining similar results require very specific architectures [6] or to train another model to generate the explanations [21], so CRAFT provides a considerable advantage in the matter of flexibility in comparison.

B. More results of CRAFT

B.1. Qualitative comparison with ACE

Figure S1 compares the examples of concepts found by CRAFT against those found by ACE [24] for 3 classes of Imagenette. For each class the concepts are ordered by importance (the highest being the most important). ACE uses a clustering technique and TCAV to estimate importance, while CRAFT uses the method introduced in 3 and Sobol to estimate importance. These examples illustrate one of the weaknesses of ACE: the segmentation used can introduce biases through the baseline value used [19,70]. The concepts found by CRAFT seem distinct: (vault, cross, stained glass) for the Church class, (dumpster, truck door, two-wheeler) for the garbage truck, and (eyes, nose, fluffy ears) for the English Springer.

B.2. Most important concepts.

We show more example of the 4 most importants concepts for 6 classes: 'Chain saw', 'English springer', 'Gas pump', 'Golf ball', 'French horn' and 'Garbage Truck' (Figure S2).

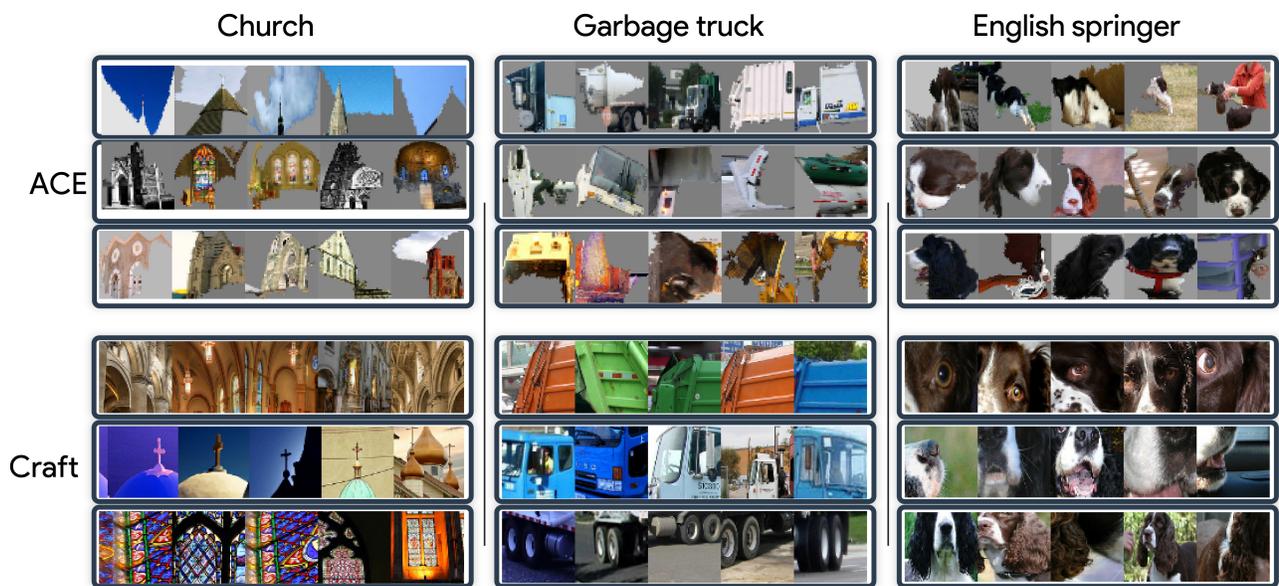


Figure S1. **Qualitative comparison.** We compare concepts found by our method (top) to those extracted with ACE [24] (bottom) for the classes *Church*, *Garbage truck* and *English springer* from ILSVRC2012 [10].

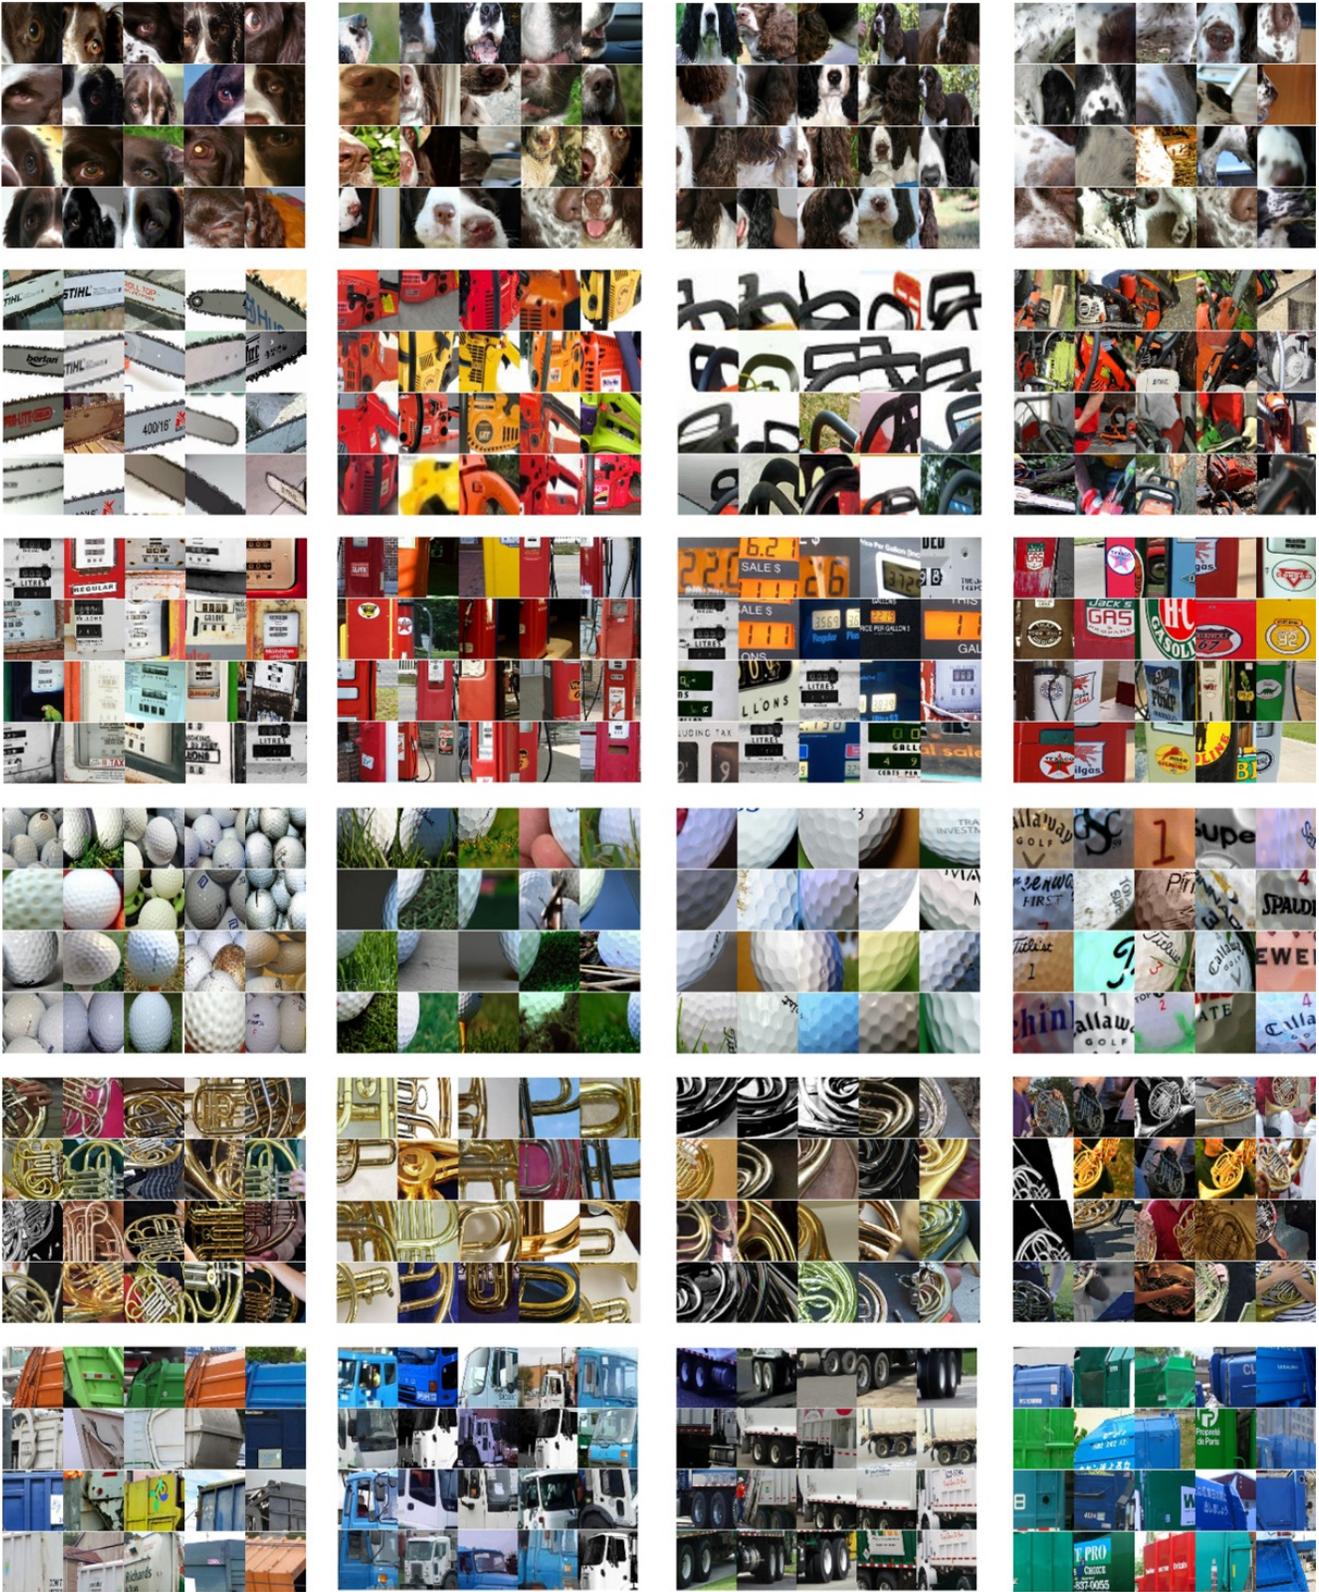


Figure S2. **CRAFT most important concepts.** The 4 most important concepts ranked by importance (left to right) for the following classes: 'English springer', 'Chain saw', 'Gas pump', 'Golf ball', 'French horn', and 'Garbage truck'.

B.3. Feature Visualization validation

Another way of interpreting concepts – as per [40] – is to employ feature visualization methods: through optimization, find an image that maximizes an activation pattern. In our case, we used the set of regularization and constraints proposed by [51], which allow us to successfully obtain realistic images. In Figures [S3-S5], we showcase these synthetic images obtained through feature visualization, along with the segments that maximize the target concept. We observe that they do reflect the underlying concepts of interest.

Concretely, to produce those feature visualization, we are looking for an image x^* that is optimized to correspond to a concept from the concept bank \mathbf{W}_i . We use the so called ‘dot-cosim’ loss proposed by [51], which give the following objective:

$$x^* = \arg \max_{x \in \mathcal{X}} \langle g(x), \mathbf{W}_i \rangle \frac{\langle g(x), \mathbf{W}_i \rangle^2}{\|g(x)\| \|\mathbf{W}_i\|} - \mathcal{R}(x)$$

With $\mathcal{R}(\cdot)$, the regularizations applied to x – the default regularizations in the **Xplique** library [16]. As for the specific parameters, we used Fourier preconditioning on the image with a decay rate of 0.8 and an Adam optimizer ($lr = 1e - 1$).

here

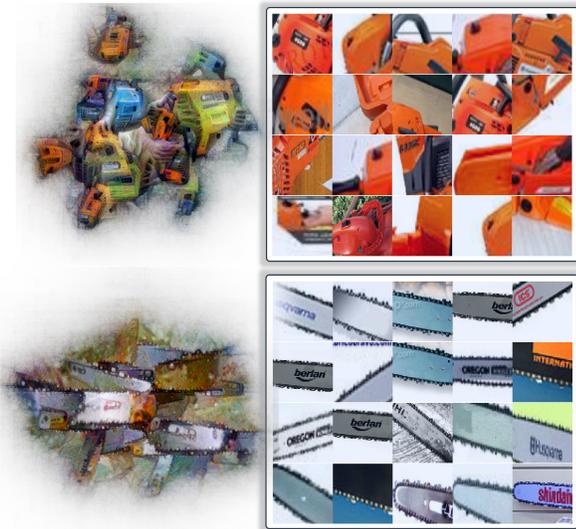


Figure S3. Feature visualization for chainsaw CRAFT concepts.



Figure S4. Feature visualization for english springer CRAFT concepts.

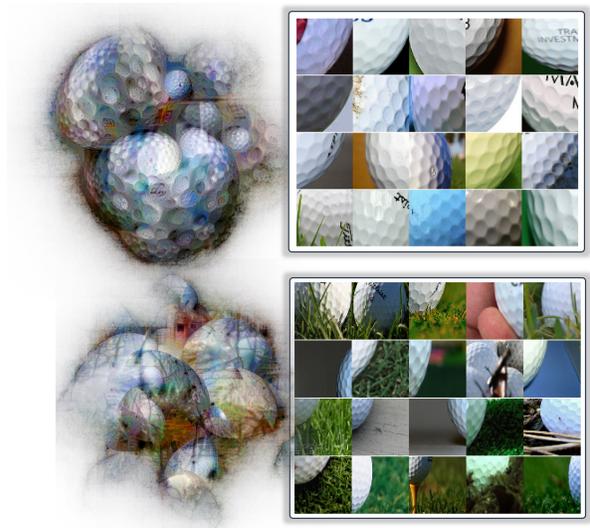


Figure S5. Feature visualization for golf CRAFT concepts.

C. Backpropagating through the NMF block

C.1. Alternating Direction Method of Multipliers (ADMM) for NMF

We recall that NMF decomposes the positive features vector $\mathbf{A} \in \mathbb{R}^{n \times p}$ of n examples lying in dimension p , into a product of positive low rank matrices $\mathbf{U}(\mathbf{A}) \in \mathbb{R}^{n \times r}$ and $\mathbf{W}(\mathbf{A}) \in \mathbb{R}^{p \times r}$ (with $r \ll \min(n, p)$), i.e the solution to the problem:

$$\min_{\mathbf{U} \geq 0, \mathbf{W} \geq 0} \frac{1}{2} \|\mathbf{A} - \mathbf{U}\mathbf{W}^T\|_F^2. \quad (5)$$

For simplicity we used a non-regularized version of the NMF objective, following Algorithms 1 and 3 in paper [32], based on ADMM [5]. This algorithm transforms the non-linear equality constraints into indicator functions δ . Auxiliary variables $\tilde{\mathbf{U}}, \tilde{\mathbf{W}}$ are also introduced to separate the optimization of the objective on the one side, and the satisfaction of the constraint on \mathbf{U}, \mathbf{W} on the other side. The equality constraints $\tilde{\mathbf{U}} = \mathbf{U}, \tilde{\mathbf{W}} = \mathbf{W}$ are linear and easily handled by the ADMM framework through the associated dual variables $\bar{\mathbf{U}}, \bar{\mathbf{W}}$. In our case, the problem in Equation 5 is transformed into:

$$\begin{aligned} \min_{\mathbf{U}, \tilde{\mathbf{U}}, \mathbf{W}, \tilde{\mathbf{W}}} \quad & \frac{1}{2} \|\mathbf{A} - \tilde{\mathbf{U}}\tilde{\mathbf{W}}^T\|_F^2 + \delta(\mathbf{U}) + \delta(\mathbf{W}), \\ \text{s.t.} \quad & \tilde{\mathbf{U}} = \mathbf{U}, \tilde{\mathbf{W}} = \mathbf{W} \\ & \text{with } \delta(\mathbf{H}) = \begin{cases} 0 & \text{if } \mathbf{H} \geq 0, \\ +\infty & \text{otherwise.} \end{cases} \end{aligned} \quad (6)$$

Note that $\tilde{\mathbf{U}}$ and \mathbf{U} (resp. $\tilde{\mathbf{W}}$ and \mathbf{W}) seem redundant: they are meant to be equal thanks to constraints $\tilde{\mathbf{U}} = \mathbf{U}, \tilde{\mathbf{W}} = \mathbf{W}$. This is standard practice within ADMM framework: introducing redundancies allows to disentangle the (unconstrained) optimization of the objective on one side (with $\tilde{\mathbf{U}}$ and $\tilde{\mathbf{W}}$) and constraint satisfaction on the other side with \mathbf{U} and \mathbf{W} . During the optimization process the variables $\tilde{\mathbf{U}}, \mathbf{U}$ (resp. $\tilde{\mathbf{W}}, \mathbf{W}$) are different, and only become equal in the limit at convergence. The dual variables $\bar{\mathbf{U}}, \bar{\mathbf{W}}$ control the balance between optimization of the objective $\frac{1}{2} \|\mathbf{A} - \tilde{\mathbf{U}}\tilde{\mathbf{W}}^T\|_F^2$ and constraint satisfaction $\tilde{\mathbf{U}} = \mathbf{U}, \tilde{\mathbf{W}} = \mathbf{W}$. The constraints are simplified at the cost of a non-smooth (and even a non-finite) objective function $\frac{1}{2} \|\mathbf{A} - \tilde{\mathbf{U}}\tilde{\mathbf{W}}^T\|_F^2 + \delta(\mathbf{U}) + \delta(\mathbf{W})$ due to the term $\delta(\mathbf{U}) + \delta(\mathbf{W})$. ADMM proceeds to create a so-called *augmented Lagrangian* with l_2 regularization $\rho > 0$:

$$\begin{aligned} \mathcal{L}(\mathbf{A}, \mathbf{U}, \mathbf{W}, \tilde{\mathbf{U}}, \tilde{\mathbf{W}}, \bar{\mathbf{U}}, \bar{\mathbf{W}}) = & \\ \frac{1}{2} \|\mathbf{A} - \tilde{\mathbf{U}}\tilde{\mathbf{W}}^T\|_F^2 + \delta(\mathbf{U}) + \delta(\mathbf{W}) & \\ + \bar{\mathbf{U}}^T(\tilde{\mathbf{U}} - \mathbf{U}) + \bar{\mathbf{W}}^T(\tilde{\mathbf{W}} - \mathbf{W}) & \\ + \frac{\rho}{2} \left(\|\tilde{\mathbf{U}} - \mathbf{U}\|_2^2 + \|\tilde{\mathbf{W}} - \mathbf{W}\|_2^2 \right). & \end{aligned} \quad (7)$$

This regularization ensures that the dual problem is well posed and that it remain convex, even with the non smooth and infinite terms $\delta(\mathbf{U}) + \delta(\mathbf{W})$. Once again, this is standard practice within ADMM framework. The (regularized) problem associated to this Lagrangian is decomposed into a sequence of convex problems that alternate minimization over the $\mathbf{U}, \tilde{\mathbf{U}}, \bar{\mathbf{U}}$ and the $\mathbf{W}, \tilde{\mathbf{W}}, \bar{\mathbf{W}}$ triplets.

$$\mathbf{U}_{t+1} = \arg \min_{\mathbf{U} = \tilde{\mathbf{U}}} \frac{1}{2} \|\mathbf{A} - \tilde{\mathbf{U}}\mathbf{W}_t^T\|_F^2 + \delta(\mathbf{U}) + \frac{\rho}{2} \|\tilde{\mathbf{U}} - \mathbf{U}\|_2^2. \quad (8)$$

$$\mathbf{W}_{t+1} = \arg \min_{\mathbf{W} = \tilde{\mathbf{W}}} \frac{1}{2} \|\mathbf{A} - \mathbf{U}_t\tilde{\mathbf{W}}^T\|_F^2 + \delta(\mathbf{W}) + \frac{\rho}{2} \|\tilde{\mathbf{W}} - \mathbf{W}\|_2^2. \quad (9)$$

This guarantees a monotonic decrease of the objective function $\|\mathbf{A} - \tilde{\mathbf{U}}_t\tilde{\mathbf{W}}_t^T\|_F^2$. Each of these sub-problems is thus solved with ADMM separately, by alternating minimization steps of $\frac{1}{2} \|\mathbf{A} - \tilde{\mathbf{U}}\mathbf{W}_t^T\|_F^2 + \tilde{\mathbf{U}}^T(\tilde{\mathbf{U}} - \mathbf{U}) + \frac{\rho}{2} \|\mathbf{U} - \tilde{\mathbf{U}}\|_2^2$ over $\tilde{\mathbf{U}}$ (*i*), with minimization steps of $\delta(\mathbf{U}) + \frac{\rho}{2} \|\mathbf{U} - \tilde{\mathbf{U}}\|_2^2$ over \mathbf{U} (*ii*), and gradient ascent steps (*iii*) on the dual variable $\bar{\mathbf{U}} \leftarrow \bar{\mathbf{U}} + (\tilde{\mathbf{U}} - \mathbf{U})$. A similar scheme is used for \mathbf{W} updates. Step (*i*) is a simple convex quadratic program with equality constraints, whose KKT [38, 45] conditions yield a linear system with a Positive Semi-Definite (PSD) matrix. Step (*ii*) is a simple projection of $\tilde{\mathbf{U}}$ onto the convex set $\delta^{-1}(\mathbf{0})$. Finally, step (*iii*) is inexpensive.

Concretely, we solved the quadratic program using Conjugate Gradient [30], from *jax.scipy.sparse.linalg.cg*. This indirect method only involves *matrix-vector* products and can be more GPU-efficient than methods that are based on matrix factorization (such as Cholesky decomposition). Also, we re-implemented the pseudo code of [32] in *Jax* for a fully GPU-compatible program. We used the primal variables $\mathbf{U}_0, \mathbf{W}_0$ returned by *sklearn.decompose.nmf* as a *warm start* for ADMM and observe that the high quality initialization of these primal variables considerably speeds up the convergence of the dual variables.

C.2. Implicit differentiation

The Lagrangian of the NMF problem reads $\mathcal{L}(\mathbf{U}, \mathbf{W}, \bar{\mathbf{U}}, \bar{\mathbf{W}}) = \frac{1}{2} \|\mathbf{A} - \mathbf{U}\mathbf{W}^T\|_F^2 - \bar{\mathbf{U}}^T\mathbf{U} - \bar{\mathbf{W}}^T\mathbf{W}$, with dual variables $\bar{\mathbf{U}}$ and $\bar{\mathbf{W}}$ associated to the constraints $\mathbf{U} \geq 0, \mathbf{W} \geq 0$. It yields a function F based on the KKT conditions [38, 45] whose optimal tuple $\mathbf{U}, \mathbf{W}, \bar{\mathbf{U}}, \bar{\mathbf{W}}$ is a root.

For single NNLS problem (for example, with optimization over \mathbf{U}) the KKT conditions are:

$$\begin{cases} \nabla_{\mathbf{U}} \left(\frac{1}{2} \|\mathbf{A} - \tilde{\mathbf{U}}\tilde{\mathbf{W}}^T\|_F^2 + \tilde{\mathbf{U}}^T(-\mathbf{U}) \right) = 0, \text{ stationarity,} \\ -\mathbf{U} \leq 0, \text{ primal feasibility,} \\ \tilde{\mathbf{U}} \odot \mathbf{U} = 0, \text{ complementary slackness,} \\ \tilde{\mathbf{U}} \geq 0, \text{ dual feasibility.} \end{cases} \quad (10)$$

By stacking the KKT conditions of the NNLS problems the we obtain the so-called *optimality function* \mathbf{F} :

$$\mathbf{F}((\mathbf{U}, \mathbf{W}, \tilde{\mathbf{U}}, \tilde{\mathbf{W}}), \mathbf{A}) = \begin{cases} (\mathbf{U}\mathbf{W}^T - \mathbf{A})\mathbf{W} - \tilde{\mathbf{U}}, \\ (\mathbf{W}\mathbf{U}^T - \mathbf{A}^T)\mathbf{U} - \tilde{\mathbf{W}}, \\ \tilde{\mathbf{U}} \odot \mathbf{U}, \\ \tilde{\mathbf{W}} \odot \mathbf{W}. \end{cases} \quad (11)$$

The implicit function theorem [25] allows us to use implicit differentiation [3, 25, 44] to efficiently compute the Jacobians $\frac{\partial \mathbf{U}}{\partial \mathbf{A}}$ and $\frac{\partial \mathbf{W}}{\partial \mathbf{A}}$ without requiring to back-propagate through each of the iterations of the NMF solver:

$$\frac{\partial(\mathbf{U}, \mathbf{W}, \tilde{\mathbf{U}}, \tilde{\mathbf{W}})}{\partial \mathbf{A}} = -(\partial_1 \mathbf{F})^{-1} \partial_2 \mathbf{F}. \quad (12)$$

Implicit differentiation requires access to the dual variables of the optimization problem in equation 1, which are not computed by Scikit-learn’s popular implementation. Scikit-learn uses Block coordinate descent algorithm [7, 17], with a randomized SVD initialization. Consequently, we leverage our implementation in Jax based on ADMM [5].

Concretely, we perform a two-stage backpropagation *Jax* (2) \rightarrow *Tensorflow* (1) to leverage the advantage of each framework. The lower stage (1) corresponds to feature extraction $\mathbf{A} = \mathbf{h}_l(\mathbf{X})$ from crops of images \mathbf{X} , and upper stage (2) computes NMF $\mathbf{A} \approx \mathbf{U}\mathbf{W}^T$.

We use the *Jaxopt* [4] library that allows efficient computation of $\frac{\partial(\mathbf{U}, \mathbf{W}, \tilde{\mathbf{U}}, \tilde{\mathbf{W}})}{\partial \mathbf{A}} = -(\partial_1 \mathbf{F})^{-1} \partial_2 \mathbf{F}$. The matrix $(\partial_1 \mathbf{F})^{-1}$ is never explicitly computed – that would be too costly. Instead, the system $\partial_1 \mathbf{F} \frac{\partial(\mathbf{U}, \mathbf{W}, \tilde{\mathbf{U}}, \tilde{\mathbf{W}})}{\partial \mathbf{A}} = -\partial_2 \mathbf{F}$ is solved with Conjugate Gradient [30] through the use of Jacobian Vector Products (JVP) $\mathbf{v} \mapsto (\partial_1 \mathbf{F})\mathbf{v}$.

The chain rule yields:

$$\frac{\partial \mathbf{U}}{\partial \mathbf{X}} = \frac{\partial \mathbf{A}}{\partial \mathbf{X}} \frac{\partial \mathbf{U}}{\partial \mathbf{A}}.$$

Usually, most Autodiff frameworks (e.g Tensorflow, Pytorch, Jax) handle it automatically. Unfortunately, combining two of those framework raises a new difficulty since they are not compatible. Hence, we re-implement manually the two stages auto-differentiation.

Since r is far smaller ($r = 25$ in all our experiments) than input dimension \mathbf{X} (typically 224×244 for ImageNet

images), back-propagation is the preferred algorithm in this setting over forward-propagation. We start by computing sequentially the gradients $\nabla_{\mathbf{X}} \mathbf{U}_i$ for all concepts $1 \leq i \leq r$. This amounts to compute $\mathbf{v} = \nabla_{\mathbf{A}} \mathbf{U}_i$ with Implicit Differentiation in Jax, convert the Jax array \mathbf{v} into Tensorflow tensor, and then to compute $\nabla_{\mathbf{X}} \mathbf{U}_i = \frac{\partial \mathbf{A}}{\partial \mathbf{X}} \nabla_{\mathbf{A}} \mathbf{U}_i = \nabla_{\mathbf{X}} (\mathbf{h}_l(\mathbf{X}) \cdot \mathbf{v})$. The latter is easily done in Tensorflow. Finally we stack the gradients $\nabla_{\mathbf{X}} \mathbf{U}_i$ to obtain the Jacobian $\frac{\partial \mathbf{U}}{\partial \mathbf{X}}$.

D. Sobol indices for concepts

We propose to formally derive the Sobol indices for the estimation of the importance of concepts. Let us define a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ of possible concept perturbations. In order to build these concept perturbations, we start from an original vector of concepts coefficient $\tilde{\mathbf{U}} \in \mathbb{R}^r$ and use i.i.d. stochastic masks $\mathbf{M} = (M_1, \dots, M_r) \sim \mathcal{U}([0, 1]^r)$, as well as a perturbation operator $\tau(\cdot)$ to create stochastic perturbation of $\tilde{\mathbf{U}}$ that we call concept perturbation $\mathbf{U} = \tau(\tilde{\mathbf{U}}, \mathbf{M})$.

Concretely, to create our concept perturbation we consider the inpainting function as our perturbation operator (as in [14, 55, 57]): $\tau(\mathbf{U}, \mathbf{M}) = \mathbf{U} \odot \mathbf{M} + (\mathbf{1} - \mathbf{M})\mu$ with \odot the Hadamard product and $\mu \in \mathbb{R}$ a baseline value, here zero. For the sake of notation, we will note $\mathbf{f} : \mathcal{F} \rightarrow \mathbb{R}$ the function mapping a random concept perturbation \mathbf{U} from an intermediate layer to the output. We denote the set $\mathcal{U} = \{1, \dots, r\}$, \mathbf{u} a subset of \mathcal{U} , its complementary $\sim \mathbf{u}$ and $\mathbb{E}(\cdot)$ the expectation over the perturbation space. Finally, we assume that $\mathbf{f} \in \mathbb{L}^2(\mathcal{F}, \mathbb{P})$ i.e. $|\mathbb{E}(\mathbf{f}(\mathbf{U}))| < +\infty$.

The Hoeffding decomposition allows us to express the function \mathbf{f} into summands of increasing dimension, denoting $\mathbf{f}_{\mathbf{u}}$ the partial contribution of the concepts $\mathbf{U}_{\mathbf{u}} = (U_i)_{i \in \mathbf{u}}$ to the score $\mathbf{f}(\mathbf{U})$:

$$\begin{aligned} \mathbf{f}(\mathbf{U}) &= \mathbf{f}_{\emptyset} \\ &+ \sum_i^r \mathbf{f}_i(U_i) \\ &+ \sum_{1 \leq i < j \leq r} \mathbf{f}_{i,j}(U_i, U_j) + \dots \\ &+ \mathbf{f}_{1, \dots, r}(U_1, \dots, U_r) \\ &= \sum_{\mathbf{u} \subseteq \mathcal{U}} \mathbf{f}_{\mathbf{u}}(\mathbf{U}_{\mathbf{u}}). \end{aligned} \quad (13)$$

Eq. 13 consists of 2^r terms and is unique under the following orthogonality constraint:

$$\forall (\mathbf{u}, \mathbf{v}) \subseteq \mathcal{U}^2 \text{ s.t. } \mathbf{u} \neq \mathbf{v}, \quad \mathbb{E}(\mathbf{f}_{\mathbf{u}}(\mathbf{U}_{\mathbf{u}}) \mathbf{f}_{\mathbf{v}}(\mathbf{U}_{\mathbf{v}})) = 0. \quad (14)$$

Furthermore, orthogonality yields the characterization $\mathbf{f}_{\mathbf{u}}(\mathbf{U}_{\mathbf{u}}) = \mathbb{E}(\mathbf{f}(\mathbf{U}) | \mathbf{U}_{\mathbf{u}}) - \sum_{\mathbf{v} \subsetneq \mathbf{u}} \mathbf{f}_{\mathbf{v}}(\mathbf{U}_{\mathbf{v}})$ and allows us

to decompose the model variance as:

$$\begin{aligned}
\mathbb{V}(\mathbf{f}(\mathbf{U})) &= \sum_i^r \mathbb{V}(\mathbf{f}_i(U_i)) \\
&+ \sum_{1 \leq i < j \leq r} \mathbb{V}(\mathbf{f}_{i,j}(U_i, U_j)) \\
&+ \dots + \mathbb{V}(\mathbf{f}_{1,\dots,r}(U_1, \dots, U_r)) \\
&= \sum_{\mathbf{u} \subseteq \mathcal{U}} \mathbb{V}(\mathbf{f}_{\mathbf{u}}(\mathbf{U}_{\mathbf{u}})).
\end{aligned} \tag{15}$$

Building from Eq. 15, it is natural to characterize the influence of any subset of concepts \mathbf{u} as its own variance w.r.t. the total variance. This yields, after normalization by $\mathbb{V}(\mathbf{f}(\mathbf{U}))$, the general definition of Sobol' indices.

Definition D.1 (Sobol indices [67]). The sensitivity index $\mathcal{S}_{\mathbf{u}}$ which measures the contribution of the concept set $\mathbf{U}_{\mathbf{u}}$ to the model response $\mathbf{f}(\mathbf{U})$ in terms of fluctuation is given

by:

$$\begin{aligned}
\mathcal{S}_{\mathbf{u}} &= \frac{\mathbb{V}(\mathbf{f}_{\mathbf{u}}(\mathbf{U}_{\mathbf{u}}))}{\mathbb{V}(\mathbf{f}(\mathbf{U}))} \\
&= \frac{\mathbb{V}(\mathbb{E}(\mathbf{f}(\mathbf{U})|\mathbf{U}_{\mathbf{u}})) - \sum_{\mathbf{v} \subset \mathbf{u}} \mathbb{V}(\mathbb{E}(\mathbf{f}(\mathbf{U})|\mathbf{U}_{\mathbf{v}}))}{\mathbb{V}(\mathbf{f}(\mathbf{U}))}.
\end{aligned} \tag{16}$$

Sobol indices give a quantification of the importance of any subset of concepts with respect to the model decision, in the form of a normalized measure of the model output deviation from $\mathbf{f}(\mathbf{U})$. Thus, Sobol indices sum to one : $\sum_{\mathbf{u} \subseteq \mathcal{U}} \mathcal{S}_{\mathbf{u}} = 1$.

Furthermore, the framework of Sobol' indices enables us to easily capture higher-order interactions between features. Thus, we can view the Total Sobol indices defined in 2 as the sum of of all the Sobol indices containing the concept i : $\mathcal{S}_i^T = \sum_{\mathbf{u} \subseteq \mathcal{U}, i \in \mathbf{u}} \mathcal{S}_{\mathbf{u}}$. Concretely, we estimate the total Sobol indices using the Jansen estimator [36] and Quasi-Monte carlo Sequence (Sobol LP_{τ} sequence).

E. Human experiments

We first describe how participants were enrolled in our studies, then the general experimental design they went through.

E.1. Utility evaluation

Participants The participants that went through our experiments are users from the online platform Amazon Mechanical Turk (AMT), specifically, we recruit users with high qualifications (number of HIT completed = 5000 and HIT accepted > 98%). All participants provided informed consent electronically in order to perform the experiment ($\sim 5 - 8$ min), for which they received 1.4\$.

For the *Husky vs. Wolf* scenario, $n = 84$ participants passed all our screening and filtering process, respectively $n = 32$ for CRAFT, $n = 22$ for ACE and $n = 22$ for CRAFTCO.

For the *Leaves* scenario, after filtering, we analyzed data from $n = 87$ participants, respectively $n = 32$ for CRAFT, $n = 24$ for ACE and $n = 31$ for CRAFTCO.

For the *"Kit Fox" vs. "Red Fox"* scenario, the results come from $n = 79$ participants who passed all our screening processes, respectively $n = 22$ for CRAFT, $n = 31$ for ACE and $n = 26$ for CRAFTCO.

General study design We followed the experimental design proposed by Colin and Fel et al. [8], in which explanations are evaluated according to their ability to help training participants at getting better at predicting their models' decisions on unseen images.

Each of those participants are only tested on a single condition to avoid possible experimental confounds.

The main experiment is divided into 3 training sessions (with 5 training samples in each) each followed by a brief test. In each individual training trial, an image was presented with the associated prediction of the model, together with an explanation. After a brief training phase (5 samples), participants' ability to predict the classifier's output was evaluated on 7 new samples during a test phase. During the test phase, no explanation was provided. We also use the

reservoir that subjects can refer to during the testing phase to minimize memory load as a confounding factor.

We implement the same 3-stage screening process as in [8]: First we filter participants not successful at the practice session done prior to the main experiment used to teach them the task, then we have them go through a quiz to make sure they understood the instructions. Finally, we add a catch trial in each testing phase –that users paying attention are expected to be correct on– allowing us to catch uncooperative participants.

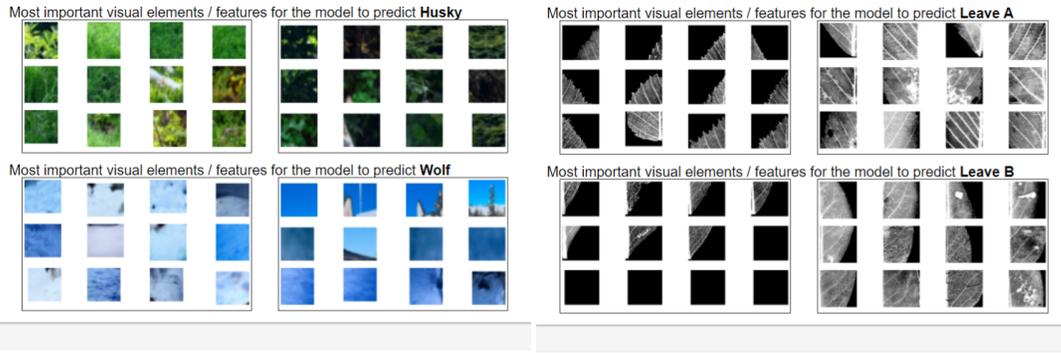
E.2. Validation of Recursivity

Participants Behavioral accuracy data were gathered from $n = 73$ participants. All participants provided informed consent electronically in order to perform the experiment ($\sim 4 - 6$ min). The protocol was approved by the University IRB and was carried out in accordance with the provisions of the World Medical Association Declaration of Helsinki. For each of the 2 experiment tested, we had prepared filtering criteria for uncooperative people (namely based on time), but all participants passed these filters.

General study design For the first experiment – consisting in finding the intruder among elements of the same concept and an element from a different concept (but of the same class, see Figure S7b) – the order of presentation is randomized across participants so that it does not bias the results. Moreover, in order to avoid any bias coming from the participants themselves (one group being more successful than the other) all participants went through both conditions of finding intruders in batches of images coming from either concepts or sub-concepts. Concerning experiment 2, the order was also randomized (see Figure S7c).

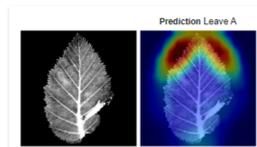
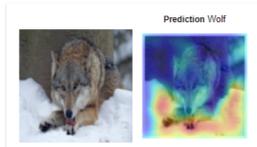
The participants had to successively find 30 intruders (15 block concepts and 15 block sub-concepts) for experiment 1 and then make 15 choices (sub-concept vs concept) for experiment 2, see Figure S7a.

The expert participants are people working in machine learning (researchers, software developers, engineers) and have participated in the study following an announcement in the authors' laboratory/company. The other participants (Laymen) have no expertise in machine learning.

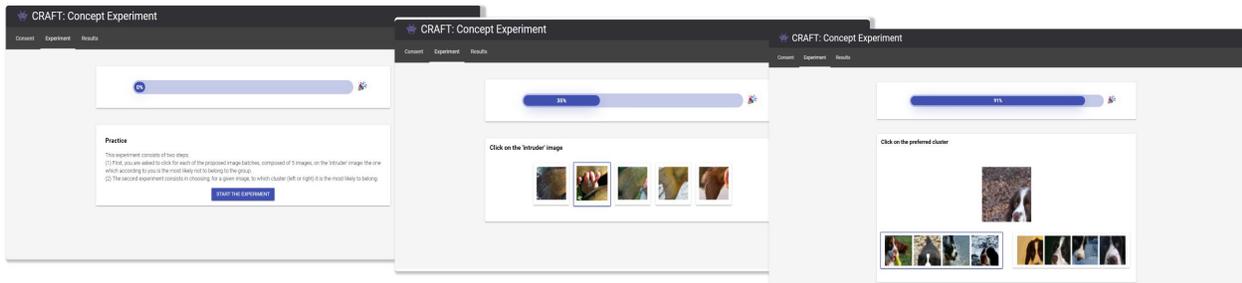


Study examples of classification of the model, try to understand its behavior.

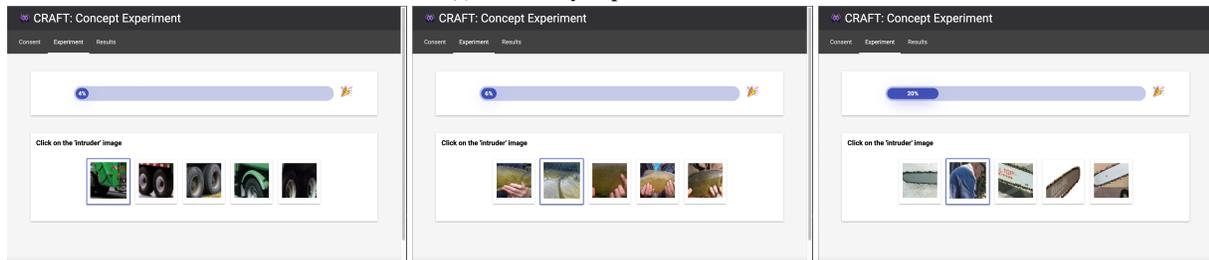
Study examples of classification of the model, try to understand its behavior.



(a) **Utility experiment.** Training trials taken from the *Husky vs. Wolf* scenario (left) and the *Leaves* scenario (right).



(a) **Recursivity Experiment Website.**



(b) **Binary choice experiment.**



(c) **Intruder experiment.**

F. Fidelity experiments

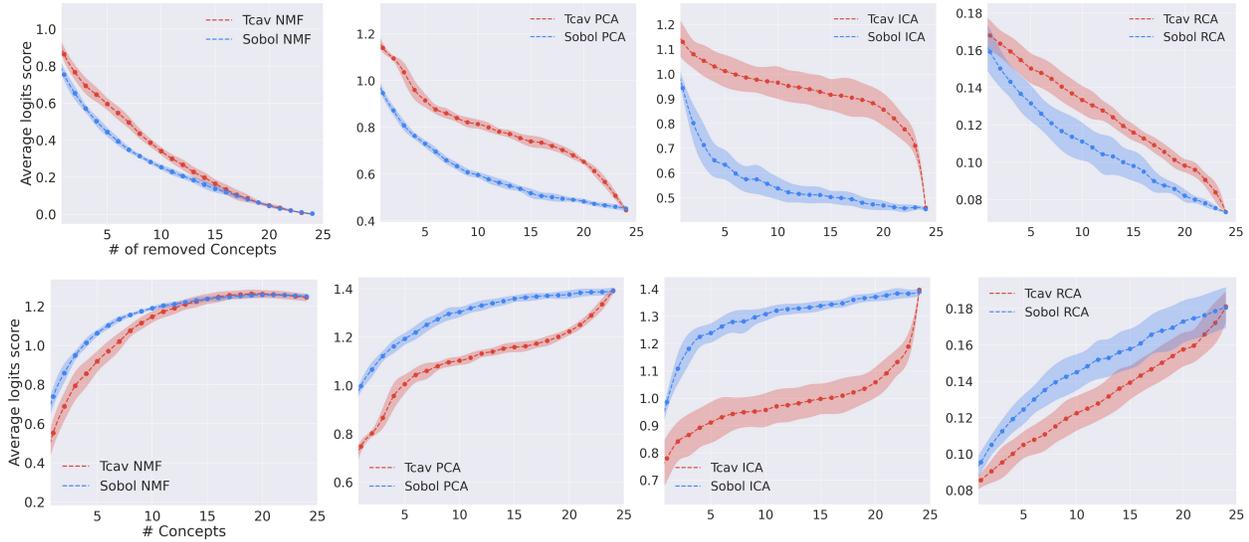


Figure S8. **(1) Deletion curves** for different concept extraction methods, Sobol outperforms TCAV not only for NMF to correctly estimate concept importance (lower is better). **(2) Insertion curves** for different concept extraction methods, Sobol outperforms TCAV to correctly estimate concept importance (higher is better).

For our experiments on the concept importance measure, we focused on certain classes of ILSRVC2012 [10] and used a ResNet50V2 [29] that had already been trained on this dataset. Just like in [24, 77], we measure the insertion and deletion metrics for our concept extraction technique – as well as concepts vectors extracted using PCA, ICA and RCA as dimensionality reduction algorithms, see Figure S8 – and we compare them when we add/remove the concepts as ranked by the TCAV score [40] and by the Sobol importance score. As originally explained in [55], the objective of these metrics is to add/remove parts of the input according to how much an explainability method considers that it is influential and looking at the speed at which the logit for the predicted class increases/decreases.

In particular, for our experimental evaluations, we have randomly chosen 100000 images from ILSVRC2012 [10] and computed the deletion and insertion metrics for 5 different seeds – for a total of half a million images. In Figure S8, the shade around the curves represent the standard deviation over these 5 experiments.

G. Sanity Check

Following the work from [1], we performed a sanity check on our method, by running the concept extraction pipeline on a randomized model. This procedure was performed on a ResNet-50v2 model with randomized weights. As showcased in Figure S9, the concepts drastically differ from trained models, thus proving that CRAFT passes the sanity check.

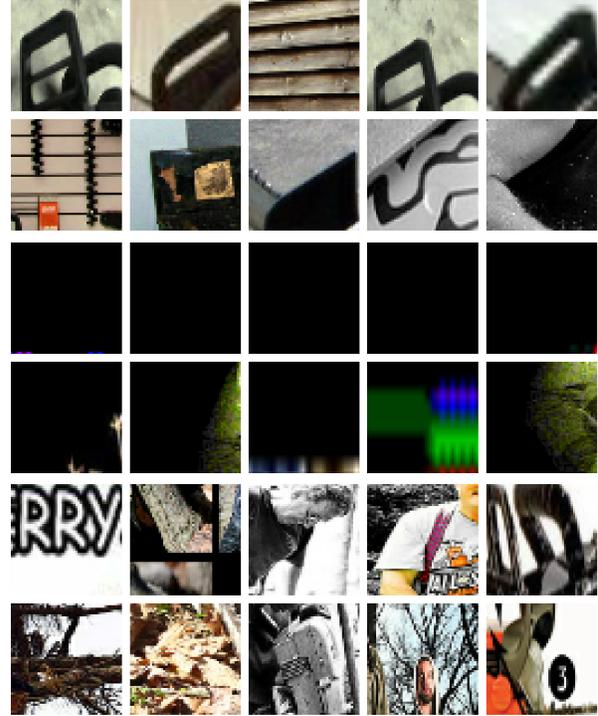


Figure S9. **Sanity check of the method:** we ran the method on a Resnet50 with randomized weights, and extracted the 3 most relevant concepts for the class 'Chain saw'. When weights are randomized, concepts are mainly based on color histograms.