



HAL
open science

What does LIME really see in images?

Damien Garreau, Dina Mardaoui

► **To cite this version:**

Damien Garreau, Dina Mardaoui. What does LIME really see in images?. ICML 2021 - 38th International Conference on Machine Learning, Jul 2021, Virtual Conference, United States. 10.48550/arXiv.2102.06307 . hal-03233014

HAL Id: hal-03233014

<https://hal.science/hal-03233014v1>

Submitted on 23 May 2021

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.

What does LIME really see in images?

Damien Garreau¹ Dina Mardaoui²

Abstract

The performance of modern algorithms on certain computer vision tasks such as object recognition is now close to that of humans. This success was achieved at the price of complicated architectures depending on millions of parameters and it has become quite challenging to understand how particular predictions are made. Interpretability methods propose to give us this understanding. In this paper, we study LIME, perhaps one of the most popular. On the theoretical side, we show that when the number of generated examples is large, LIME explanations are concentrated around a limit explanation for which we give an explicit expression. We further this study for elementary shape detectors and linear models. As a consequence of this analysis, we uncover a connection between LIME and integrated gradients, another explanation method. More precisely, the LIME explanations are similar to the sum of integrated gradients over the superpixels used in the preprocessing step of LIME.

1. Introduction

Deep neural networks and deep convolutional neural network (CNN) in particular have changed the way computers look at images (Schmidhuber, 2015). Many specific tasks in computer vision such as character recognition and object recognition are now routinely achieved by personal computers with human-like accuracy. The success of these algorithms seems partly due to the great complexity of the models they encode, the most recent relying on hundreds of layers and millions of parameters.

While the accuracy is often the only relevant metric for practitioners, there are numerous situations where one is not satisfied if the model is making good predictions for the wrong reasons. We would like to know *why* the model makes a particular prediction. Responding to this emerging need,

¹Université Côte d'Azur, Inria, CNRS, LJAD, France
²Polytech Nice. Correspondence to: Damien Garreau <damien.garreau@univ-cotedazur.fr>.

many *interpretability methods* have appeared in the last five years. Among them, *model agnostic* methods aim to provide to the user meaningful insights on the inner working of a specific algorithm without making any specific assumption on the architecture of the model. We refer to Adadi and Berrada (2018) and Guidotti et al. (2018) for a recent review.




Figure 1: Explaining a prediction with LIME. In this example, the function to be explained f is the likelihood, according to the InceptionV3 network, that the input image ξ contains a lion. After a run of LIME with default parameters, the top five positive coefficients are highlighted in the right panel.

In this paper, we study the image version of LIME (Local Interpretable Model-agnostic Explanations, Ribeiro et al., 2016). Let us recall briefly how it operates: in order to explain the prediction of a model f for example ξ , LIME

1. decomposes ξ in d superpixels, that is, small homogeneous image patches;
2. creates a number of new images x_1, \dots, x_n by *randomly turning on and off* these superpixels;
3. queries the model, getting predictions $y_i = f(x_i)$;
4. builds a local weighted surrogate model $\hat{\beta}_n$ fitting the y_i s to the presence or absence of superpixels.

Each coefficient of $\hat{\beta}_n$ is associated to a superpixel of the original image ξ and, intuitively, the more positive the more important the superpixel is for the prediction at ξ according to LIME. Generally, the user visualizes $\hat{\beta}_n$ by highlighting the superpixels associated to the top positive coefficients (usually five, see Figure 1).

The central question underlying this work is that of the soundness of LIME for explaining simple models: before using LIME on deep neural networks, are we sure that the explanations provided make sense for the most simple models? Can we guarantee it theoretically?

Contributions. Our contributions are the following:

- when the number of perturbed examples is large, the interpretable coefficients **concentrate with high probability around a vector** β that depends only on the model and the example to explain;
- we provide an **explicit expression for** β , from which we gain some reassurance on LIME. In particular, the explanations are **linear** in the model;
- for simple **shape detectors**, we can be more precise in the computation of β and we show that **LIME provides meaningful explanations** in that case;
- we can also compute β for **linear models**. The limit explanation takes a very simple form: β_j is **the sum of coefficients multiplied by pixel values on each superpixel**;
- as a consequence, we show experimentally that for smooth models the outputs of LIME are similar to the sum over superpixels of **integrated gradients**, another interpretability method.

All experiments are realized with the default settings of LIME for images unless otherwise specified.¹

Related work. This work follows the line of ideas initiated by Garreau and von Luxburg (2020a;b) for the tabular data version of LIME and later extended to text data by Mar-daoui and Garreau (2021). In particular, our main result and its proof are similar to the theory laid out in these papers. The interesting differences come from the sampling procedure of LIME for images: there is no superpixel creation step in the text and tabular data version of the algorithm. Therefore, the exact expression of the limit coefficients and the associated conclusions differ.

Organization of the paper. We start by presenting LIME for images in Section 2. Section 3 contains our main results, which are further developed for simple models in Section 4. Finally, we investigate the link between LIME and integrated gradients in Section 5.

¹The code of all experiments is available at https://github.com/dgarreau/image_lime_theory

2. LIME for images

From now on, we consider a model $f : [0, 1]^D \rightarrow \mathbb{R}$ as well as a fixed example to explain $\xi \in [0, 1]^D$. Hence D denotes the number of pixels of the images on which f operates. In practice, the inputs of f are always 2- or 3-dimensional arrays. Of particular interest, grayscale images are usually encoded as $h \times w$ arrays, whereas RGB images are $h \times w \times 3$, with each channel corresponding to a primary color. We will see that it does not make a difference and our results can be read *channel-wise* if there is more than one color channel.

2.1. Superpixels

The first step of the LIME operation is to split ξ into *superpixels*. These are contiguous patches of the image that share color and / or brightness similarities. We refer to Figure 2 for an illustration. In the text version of LIME, the counterpart of this superpixel decomposition is a local dictionary where each interpretable feature is a unique word of the text, whereas in the tabular version a complicated discretization procedure is needed.

By default, LIME uses the *quickshift* algorithm to produce these superpixels (Vedaldi and Soatto, 2008). In a nutshell, quickshift is a mode-seeking algorithm that considers the pixels as samples over a 5-dimensional space (3 color dimensions and 2 space dimensions).

For any $1 \leq k \leq d$, we denote the k th superpixel associated to ξ by J_k . Therefore, the d subsets J_1, \dots, J_d form a partition of the pixels, that is,

$$J_1 \cup \dots \cup J_d = \{1, \dots, D\} \quad \text{and} \quad J_k \cap J_\ell = \emptyset \quad \forall k \neq \ell.$$

Note that, even though the superpixels are generally contiguous patches of the image, we do not make this assumption.

2.2. Sampling

As we have seen in Section 1, one of LIME’s key ideas is to create *new examples* from ξ by randomly replacing some superpixels of the image. By default, these chosen superpixels are replaced by the mean color of the superpixel, a procedure that we call *mean replacement*. It is also possible to choose a specific color as a replacement image. We demonstrate the sampling procedure in Figure 2 as well as the two possible choices for the replacement image.

Let us be more precise and let us assume that ξ is fixed and J_1, \dots, J_d are given. The first step of the sampling scheme is to compute the replacement image $\bar{\xi} \in [0, 1]^D$. If a given color c is provided, then $\bar{\xi}_u = c$ for all $1 \leq u \leq D$. If no color is provided, then the mean image is computed: for any superpixel J_k , we define $\bar{\xi} \in [0, 1]^D$ by

$$\forall u \in J_k, \quad \bar{\xi}_u = \frac{1}{|J_k|} \sum_{u \in J_k} \xi_u. \quad (1)$$




Figure 2: Sampling procedure of LIME for images. The image to explain, ξ , is first split into $d = 72$ superpixels (*lower left corner*). A replacement image $\bar{\xi}$ is computed, which is by default the mean of ξ on each superpixel (*top row*), see Eq. (1). This replacement image can also be filled uniformly with a pre-determined color (*bottom row*: replacement with the color black). Then, for each new example x_i with $1 \leq i \leq n$, the superpixels are randomly switched depending on the throw of d independent Bernoulli $1/2$ random variables.

Of course, if the input images have several channels, the mean is computed on each channel.

Then, for each $1 \leq i \leq n$, LIME samples a random vector $z_i \in \{0, 1\}^d$ where each coordinate of z_i is i.i.d. Bernoulli with parameter $1/2$. Each $z_{i,j}$ corresponds to the activation ($z_{i,j} = 1$) or inactivation ($z_{i,j} = 0$) of superpixel j . We call the z_i s the *interpretable features*. To be precise, for any given $i \in \{1, \dots, n\}$, the new example $x_i \in [0, 1]^D$ has pixel values given by

$$\forall u \in J_j, \quad x_{i,u} = z_{i,j} \xi_u + (1 - z_{i,j}) \bar{\xi}_u. \quad (2)$$

Again, if ξ has several color channels, Eq. (2) is written channel-wise. Note that ξ corresponds to the vector $\mathbf{1} = (1, \dots, 1)^\top$ (all the superpixels of the image are activated).

2.3. Weights

Of course, the new examples x_i can be quite different from the original image. For instance, if most of the $z_{i,j}$ are zero, then x_i is close to $\bar{\xi}$. Some care is taken when building the surrogate model, and new examples are given a positive weight π_i that takes this proximity into account. By default, these weights are defined by

$$\forall 1 \leq i \leq n, \quad \pi_i := \exp\left(\frac{-d_{\cos}(\mathbf{1}, z_i)^2}{2\nu^2}\right), \quad (3)$$

where $\nu > 0$ is a positive *bandwidth parameter* equal to 0.25 by default and d_{\cos} is the *cosine distance*. Namely,

$$\forall u, v \in \mathbb{R}^d, \quad d_{\cos}(u, v) := 1 - \frac{u^\top v}{\|u\| \cdot \|v\|}.$$

We see that $d_{\cos}(z_i, \mathbf{1})$ takes near zero values if most of the superpixels are activated, and values near 1 in the opposite

scenario, as expected.

An important remark is that the weights π_i **depend only on the number of inactivated superpixels**. Indeed, conditionally to z_i having exactly s elements equal to zero, we have $z_i^\top \mathbf{1} = d - s$ and $\|z_i\| = \sqrt{d - s}$. Since $\|\mathbf{1}\| = \sqrt{d}$, using Eq. (3), we deduce that $\pi_i = \psi(s/d)$, where we defined

$$\forall t \in [0, 1], \quad \psi(t) := \exp\left(\frac{-(1 - \sqrt{1 - t})^2}{2\nu^2}\right). \quad (4)$$

2.4. Surrogate model

The next stage of LIME is to build a *surrogate model*. More precisely, LIME builds a linear model with the interpretable features z_i as input and the model predictions $y_i := f(x_i)$ as responses. This linear model, in the default implementation, is obtained by (weighted) ridge regression (Hoerl and Kennard, 1970). Formally, the outputs of LIME for model f and image ξ are given by

$$\hat{\beta}_n^\lambda \in \arg \min_{\beta \in \mathbb{R}^{d+1}} \left\{ \sum_{i=1}^n \pi_i (y_i - \beta^\top z_i)^2 + \lambda \|\beta\|^2 \right\}, \quad (5)$$

where $\lambda > 0$ is a regularization parameter. We call the coordinates of $\hat{\beta}_n^\lambda$ the *interpretable coefficients*. By convention, the 0th coordinate of $\hat{\beta}_n^\lambda$ is the intercept of the model. Some feature selection procedure can be used: we do not consider such extensions in our analysis and keep to the default implementation, which is ridge.

Another important remark is the following: as in the text and tabular cases, LIME uses the default setting of `sklearn` for the regularization parameter, that is, $\lambda = 1$. Hence the first term in Eq. (13) is roughly of order n and the

second term of order d . Since we experiment in the large n regime ($n = 1000$ is default) and with images split up in ≈ 100 superpixels, we are in a situation where $n \gg d$. Therefore, **we can consider that $\lambda = 0$ in our analysis and still recover meaningful results**. We will denote by $\hat{\beta}_n$ the solution of Eq. (13) with $\lambda = 0$, that is, ordinary least-squares.

The final step of LIME for images is to display the superpixels associated to the top *positive* coefficients of $\hat{\beta}_n^\lambda$ (usually five, see Figure 1). Part of what makes the method attractive to the practitioner is the ease with which one can read the results from one run of LIME just by looking at the highlighted part of the image. Note that it is also possible to highlight the superpixels associated to the top *negative* coefficients in another color, to see which parts of the image have a *negative* influence on the prediction.

3. Main results

In this section we present our main results. Namely, the concentration of $\hat{\beta}_n$ around β^f (Section 3.1) and the expression of β^f as a function of f and other quantities (Section 3.2).

3.1. Concentration of $\hat{\beta}_n$

When the number of new samples n is large, we expect the empirical explanations provided by LIME to stabilize. Our first result formalizes this intuition.

Theorem 1 (Concentration of $\hat{\beta}_n$). *Assume that f is bounded by a constant $M > 0$ on $[0, 1]^D$. Let $\epsilon > 0$ and $\eta \in (0, 1)$. Let d be the number of superpixels. Then, there exists $\hat{\beta}^f \in \mathbb{R}^{d+1}$ such that, for every*

$$n \gtrsim \max(M, M^2) \epsilon^{-2} d^7 e^{\frac{4}{\nu^2}} \log \frac{8d}{\eta},$$

we have $\mathbb{P}(\|\hat{\beta}_n - \beta^f\| \geq \epsilon) \leq \eta$.

We refer to the appendix for a complete statement (we omitted numerical constants and the intercept for clarity). Intuitively, Theorem 2 means that when n is large, $\hat{\beta}_n$ stabilizes around β^f . Thus we can focus on β^f to study LIME. The main limitation of Theorem 2 is the dependency on d and ν : the control that we achieve on $\|\hat{\beta}_n - \beta^f\|$ is quite poor whenever d is too large or ν is too small. Note also that $\hat{\beta}_n$ is given by the *non-regularized* version of LIME.

Theorem 2 is quite similar to Theorem 1 in Garreau and von Luxburg (2020b) and Theorem 1 in Mardaoui and Garreau (2021), which are essentially the same result for the tabular data and the text data version of LIME. The rate of convergence is slightly better here, but this seems to be an artifact of the proof and we do not think that one should sample less when dealing with images.

3.2. Expression of β^f

In this section we obtain the explicit expression of β^f . Before doing so, we need to introduce additional notation. From now on, we introduce the random variable $z \in \{0, 1\}^d$ such that z_1, \dots, z_n are i.i.d. samples of z ; it is the only source of randomness in the sampling and all expectations are taken with respect to it. We denote by π and x the associated weights and examples.

Definition 1 (α coefficients). Define $\alpha_0 := \mathbb{E}[\pi]$ and, for any $1 \leq p \leq d$, $\alpha_p := \mathbb{E}[\pi z_1 \cdots z_p]$.

Intuitively, when ν is large, α_p corresponds to the probability that exactly p superpixels of ξ are turned *on*. Since the sampling scheme of LIME for images is completely symmetrical as well as the definition of the weights, we see that this probability does not depend on the exact set of indices, hence the definition of the α coefficients. We show in appendix that the expected covariance matrix of problem (13) can be written with the first three α coefficients. Though Definition 1 is identical to Definition 3 in Mardaoui and Garreau (2021), the exact expression of the α coefficients is different in this case since the sampling procedure differs.

Proposition 1 (Computation of the α coefficients). *Let $d \geq 2$ and $0 \leq p \leq d$. For any $\nu > 0$, it holds that*

$$\alpha_p = \frac{1}{2^d} \sum_{s=0}^d \binom{d-p}{s} \psi\left(\frac{s}{d}\right),$$

where ψ is defined as in Eq. (14).

We prove Proposition 1 in the appendix. From the α coefficients, we then form the normalization constant

$$c_d := (d-1)\alpha_0\alpha_2 - d\alpha_1^2 + \alpha_0\alpha_1,$$

and the σ coefficients:

Definition 2 (σ coefficients). For any $d \geq 2$ and $\nu > 0$, define

$$\begin{cases} \sigma_1 & := -\alpha_1, \\ \sigma_2 & := \frac{(d-2)\alpha_0\alpha_2 - (d-1)\alpha_1^2 + \alpha_0\alpha_1}{\alpha_1 - \alpha_2}, \\ \sigma_3 & := \frac{\alpha_1^2 - \alpha_0\alpha_2}{\alpha_1 - \alpha_2}. \end{cases}$$

We show in appendix that the *inverse of the expected covariance matrix* associated to problem (13) can be expressed with the help of the σ coefficients and c_d . With these notation in hand, we have:

Proposition 2 (Expression of β^f). *Under the assumptions of Theorem 2, we have $c_d > 0$ and, for any $1 \leq j \leq d$,*

$$\beta_j^f = c_d^{-1} \left[\sigma_1 \mathbb{E}[\pi f(x)] + \sigma_2 \mathbb{E}[\pi z_j f(x)] + \sigma_3 \sum_{\substack{k=1 \\ k \neq j}}^d \mathbb{E}[\pi z_k f(x)] \right].$$

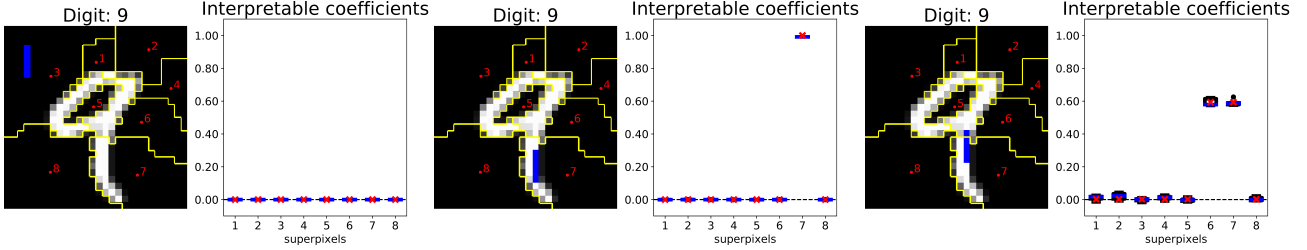


Figure 3: In this figure, we show how the theoretical predictions of Proposition 15 compare to practice. We considered a digit from the MNIST dataset (LeCun et al., 1998). The function to explain takes value 1 if all pixels marked in blue have value greater than $\tau = 0.5$, 0 otherwise. In each case, we ran LIME with $n = 1000$ examples, default regularization $\lambda = 1$ and zero replacement. We repeated the experiment five times, which gave the boxplot corresponding to the empirical values of the interpretable coefficients for each superpixel. The red crosses correspond to the predictions given by Proposition 15. We see that when the shape is split among p superpixels, each one receives a coefficient approximately equal to $1/2^p$.

We provide a detailed proof of Proposition 2 as well as the expression of the intercept β_0^f in the appendix. Let us note that Proposition 2 is quite similar to Eq. (6) in Garreau and von Luxburg (2020b) and Eq. (9) in Mardaoui and Garreau (2021). We will see that many properties of β^f are similar to the tabular and text case. Let us now present some immediate consequences of Proposition 2.

Linearity of explanations. As in the tabular and the text setting, the mapping $f \mapsto \beta^f$ is **linear**. Thus for any model that can be decomposed as a sum, the explanations provided by LIME are **the sum of the explanations of individual models**. This is true up to noise coming from the sampling (quantified by Theorem 2) and a small error due to the regularization, which is not taken into account in Theorem 2.

Large bandwidth. Because of the weights π and their complex dependency in the bandwidth ν , it can be difficult to make sense of Proposition 2 in the general case. It is somewhat easier when $\nu \rightarrow +\infty$. Indeed, we show in the appendix that $c_d \rightarrow 1/4$, $\sigma_1 \rightarrow -1/2$, $\sigma_2 \rightarrow 1$, and $\sigma_3 \rightarrow 0$. Moreover, $\pi \rightarrow 1$ almost surely. Therefore, by dominated convergence, the expression of β^f simplifies to

$$(\beta_\infty^f)_j = 2 (\mathbb{E}[f(x)|z_j = 1] - \mathbb{E}[f(x)]) , \quad (6)$$

for any $1 \leq j \leq d$. In other words, the explanation provided by LIME is proportional to the difference between the mean value of the model conditioned to superpixel j being activated and the mean value of the model for x sampled as explained previously. It seems that Eq. (6) encompasses a desirable trait of LIME for images: when the bandwidth is large, the interpretable coefficient for superpixel j takes large positive values if **in the vicinity of ξ , the model takes significantly larger values when this superpixel is present in the image.**

4. Expression of β^f for simple models

In this section, we get meaningful insights on the explanations provided by LIME for simple models.

4.1. Shape detectors

We start with a very simple model: a *fixed shape detector* for gray scale images. To this extent, let $\mathcal{S} := \{u_1, \dots, u_q\}$ be a set of q distinct pixels indices, the *shape*. Let $\tau \in (0, 1)$ be a positive threshold. We define the associated shape detector

$$\forall x \in [0, 1]^D, \quad f(x) := \prod_{u \in \mathcal{S}} \mathbf{1}_{x_u > \tau}. \quad (7)$$

Readily, $f(x)$ takes the value 1 if the pixels of shape \mathcal{S} are lit up, and 0 otherwise.

It is possible to compute β^f in closed-form in this case. In fact, the result does not depend on the exact shape to be detected, but rather on how it intersect the LIME superpixels. Let us define

$$E := \{j \in \{1, \dots, d\} \text{ s.t. } J_j \cap \mathcal{S} \neq \emptyset\},$$

the set of superpixels intersecting the shape \mathcal{S} . We separate this set in two parts,

$$E_+ := \{j \in E \text{ s.t. } \bar{\xi}_j > \tau\}, \text{ and } E_- := \{j \in E \text{ s.t. } \bar{\xi}_j \leq \tau\}.$$

Intuitively, E_+ (resp. E_-) contains superpixels that are brighter than τ on average (resp. darker). We also need to define

$$\mathcal{S}_+ := \{u \in \mathcal{S} \text{ s.t. } \xi_u > \tau\}, \text{ and } \mathcal{S}_- := \{u \in \mathcal{S} \text{ s.t. } \xi_u \leq \tau\}.$$

Namely, \mathcal{S}_+ (resp. \mathcal{S}_-) contains pixels in the shape that have a value greater (resp. smaller) than the threshold.

We now make the following assumption:

$$\forall j \in E_+, \quad J_j \cap \mathcal{S}_- = \emptyset. \quad (8)$$

Intuitively, Eq. (8) means that there is no superpixel intersecting the shape such that the average value of the superpixel activates our detector without the detector being activated. This could happen for instance if the shape intersection with the superpixels is very dark but pixels around are much brighter. It is a reasonable assumption since superpixels are quite homogeneous in color and shape. Moreover, in the case of grayscale images with zero replacement, Eq. (8) is always satisfied since $\tau > 0$. Nevertheless, we provide a result that does not rely on Eq. (8) in the appendix, which specializes into:

Proposition 3 (Computation of β^f , shape detector).

Let f be defined as in Eq. (7). Assume that Eq. (8) holds and let $p := |E_-|$. Then, if there exists $j \in E_-$ such that $J_j \cap \mathcal{S}_- \neq \emptyset$, $\beta^f = 0$. Otherwise, for any $j \in E_-$,

$$\beta_j^f = c_d^{-1} \{ \sigma_1 \alpha_p + \sigma_2 \alpha_p + (p-1) \sigma_3 \alpha_p + (d-p) \sigma_3 \alpha_{p+1} \}$$

and for any $j \in \{1, \dots, d\} \setminus E_-$,

$$\beta_j^f = c_d^{-1} \{ \sigma_1 \alpha_p + \sigma_2 \alpha_{p+1} + p \sigma_3 \alpha_p + (d-p-1) \sigma_3 \alpha_{p+1} \}$$

Proposition 15 is reminiscent of Proposition 3 in Mardaoui and Garreau (2021). This is not a surprise, since both these results study LIME for models with analogous structures. However, the result differ since one has to consider the intersections of the superpixels in the present case. We now make two remarks, focusing on grayscale images with zero replacement for simplicity.

Shape included in a single superpixel. In the simple case where the shape is detected and is totally included in superpixel k , then it is straightforward from the definitions of the α coefficients and c_d that $\beta_k^f = 1$ and $\beta_j^f = 0$ otherwise. This is a good property of LIME for images since superpixel k is the only relevant superpixel in this particular situation (see Figure 3).

Shape included in several superpixels. The situation is slightly more complicated when \mathcal{S} intersects several superpixels (that is, $p > 1$). Reading Proposition 15, we see that when ν is large, the coefficients for an intersecting superpixel is approximately $1/2^p$ and 0 otherwise. That is, the importance given by LIME is evenly divided between superpixels that contain part of the shape. Again, this makes sense since they are the only relevant superpixels in this case (see Figure 3).

4.2. Linear model

We now turn to *linear models*. That is,

$$\forall x \in [0, 1]^D, \quad f(x) = \sum_{u=1}^D \lambda_u x_u, \quad (9)$$

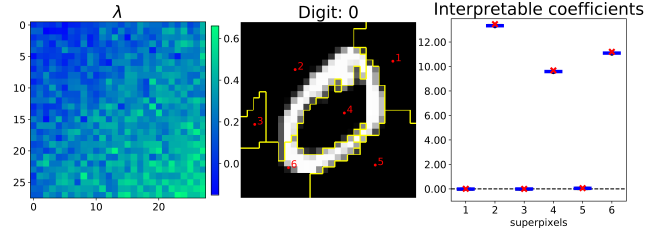


Figure 4: In this figure, we show how the theoretical predictions of Proposition 4 fare in practice. We consider a digit from the MNIST dataset. The function to explain is linear, with $\lambda_{i,j}$ proportional to $i+j$ with added white noise (heatmap in the left panel). We ran LIME 5 times with zero replacement and default parameters, the superpixels used are displayed in the middle panel. We see that our predictions match perfectly. As predicted, J_5 receives a coefficient equal to zero whereas f has high coefficients in this area, since the pixel values are equal to zero in this superpixel.

with $\lambda_1, \dots, \lambda_D \in \mathbb{R}$ arbitrary coefficients. Note that we can adapt this definition if there is more than one color channel, by considering $3 \times D$ coefficients.

Let us compute β^f when f is linear. We show in the appendix the following:

Proposition 4 (Computation of β^f , linear case). Assume that f is defined as in Eq. (32). Then, for any $1 \leq j \leq d$,

$$\beta_j^f = \sum_{u \in J_j} \lambda_u \cdot (\xi_u - \bar{\xi}_u).$$

When $\bar{\xi} = 0$, the coefficients take a very simple expression. Namely, the interpretable coefficient associated to superpixel J_j is **the sum of the coefficients of f multiplied by the pixel values on the superpixel**. If another replacement is chosen, then the *normalized* values of the pixel is taken into account in this product. This seems to make a lot of sense: let us say that the coefficients of f take large positive values on superpixel j . Then LIME will give a high interpretable coefficient to this superpixel, unless the pixel values are small (or very close to the replacement color if another replacement is chosen). We illustrate the accuracy of Proposition 4 in Figure 4. It is also interesting to see that there is no bandwidth dependency in Proposition 4. In a sense, this is to be expected since LIME is doing averages that are scale invariant if the function to explain is linear.

Proposition 4 is similar in spirit to Eq. (12) in Mardaoui and Garreau (2021), where the interpretable coefficients provided by LIME for a linear model were also found to be (approximately) equal to the product of the coefficients of the linear model and the feature value.

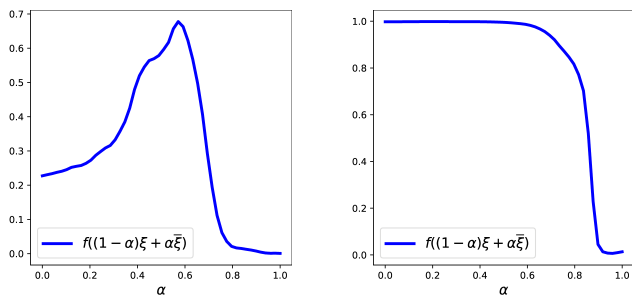


Figure 5: Predictions given by the InceptionV3 network (Szegedy et al., 2016) on a linear path between ξ ($\alpha = 0$) and $\bar{\xi}$ ($\alpha = 1$). *Left panel*: we see how the predicted values can vary along the path and why only considering the gradient at ξ or $\bar{\xi}$ may not be a good idea to build a linear approximation. *Right panel*: we see how the gradient can saturate when the network is confident in the prediction.

5. Approximated explanations

An interesting question in light of the results of the previous question is the following: if we approximate f by a linear function with coefficients λ between ξ and $\bar{\xi}$, are the explanations given by Proposition 4 close to the LIME explanations? We discuss linear approximations of an arbitrary function in Section 5.1 before investigating empirically in Section 5

5.1. Linear approximation

The most natural linear approximation of a function is given by its Taylor expansion truncated at order one. Since we want to approximate $f(x)$, where x is somewhere between ξ and $\bar{\xi}$, we could write, for instance, that $f(x) \approx f(\xi) + \nabla f(\xi)^\top (x - \xi)$. There are two main objection in doing so in the present case. First, we do not expect f to be linear between ξ and $\bar{\xi}$, and taking just one gradient would lead to a poor approximation. We illustrate this behavior in Figure 5 by computing the predictions across a line between ξ and $\bar{\xi}$.

Second, it is a well-known phenomenon in modern architecture that the gradient of the model with respect to the input can *saturate* when the network is confident in the prediction for certain activation functions (see, for instance, Krizhevsky et al. (2012)). Since from our point of view f is a black-box model, we do not have information on the activation functions (in fact, we do not even assume that f is a neural network). Therefore gradients taken at ξ or $\bar{\xi}$ can be zero, giving us essentially no information on the behavior of f (see Figure 5).

For both these reasons, we build a linear approximation of f between ξ and $\bar{\xi}$ using the *averaged gradients on a linear*

path between ξ and $\bar{\xi}$. Formally, we define

$$g_u := \int_0^1 \frac{\partial f((1-\alpha)\xi + \alpha\bar{\xi})}{\partial x_u} d\alpha \quad (10)$$

the averaged gradient at pixel u . We approximate this integral by a Riemann sum, that is,

$$g_u^{\text{approx}} := \frac{1}{m} \sum_{k=1}^m \frac{\partial f((1 - \frac{k}{m})\xi + \frac{k}{m}\bar{\xi})}{\partial x_u}. \quad (11)$$

Subsequently, we approximate $f(x)$ by $(x - \bar{\xi})^\top g^{\text{approx}} + f(\bar{\xi})$. Applying Proposition 4 to this approximation we obtain the approximate explanations

$$\forall 1 \leq j \leq d, \quad \beta_j^{\text{approx}} = \sum_{u \in J_j} (\xi_u - \bar{\xi}_u) \cdot g_u^{\text{approx}}. \quad (12)$$

Inside the sum, we recognize the definition of *integrated gradients* between ξ and $\bar{\xi}$ (Sundararajan et al., 2017), another interpretability method. Eq. (12) therefore corresponds to **the sum of integrated gradients over superpixel j** .

5.2. Experiments

In this section, we show experimentally that LIME explanations are similar to the approximated explanations derived in the previous section.

Setting. We first considered images from the CIFAR10 dataset (Krizhevsky et al., 2009), that is, 32×32 RGB images belonging to ten categories. For a subset of 1000 images of the test set, we computed the explanations given by LIME with default settings, with the exception of the kernel size used by the quickshift algorithm which we decreased to 1 to get wider superpixels. We compared these explanations with the approximated explanations of Section 5 for four different models. First, we started with a very simple one-hidden-layer neural network, trained to 35% accuracy. We then moved to VGG-like architectures (Simonyan and Zisserman, 2014), progressively increasing the number of blocks in the model (from 1 to 3). For each model, we considered the function corresponding to the most likely class for ξ . We then collected the indices of the superpixels associated to the top five and top ten positive average coefficients. For the sum of integrated gradients, we considered $m = 20$ steps in Eq. (11) as in Sundararajan et al. (2017). The results are presented in Table 1.

We then moved to more realistic images coming from the test set of the 2017 large scale visual recognition challenge (LSVRC, Russakovsky et al., 2015). We used three pretrained models from the Keras framework: MobileNetV2 (Sandler et al., 2018), DenseNet121 (Huang et al., 2017), and InceptionV3 (Szegedy et al., 2016) with default input shape (299, 299, 3). Again, we compared the LIME




Figure 6: Comparing the explanations given by LIME vs approximate explanations obtained by summing the integrated gradient over the LIME superpixels. Here we explain the top predicted class for images of the ILSVRC2017 test data with the InceptionV3 network. In both cases, we showcase the top five positive coefficients. Qualitatively, the explanations obtained are quite similar, identifying close superpixels when they are not matching exactly.

Table 1: Comparison between LIME and approximated explanations for CIFAR-10. For each model, we report JX, the Jaccard similarity between the top X positive coefficients.

Model	# param.	# layers	acc.	J5	J10
1-layer	330K	1	0.35	0.99	0.99
VGG1	1M	4	0.67	0.81	0.85
VGG2	600K	8	0.69	0.75	0.81
VGG3	550K	12	0.70	0.71	0.76

default explanations to the approximated explanations for 1000 of these images. Qualitative results are presented in Figure 6 for the InceptionV3 network, while Table 2 contains the quantitative results.

Metric. We compared the list of the top 5 and 10 positive coefficients *via* the *Jaccard similarity coefficient*, that is, the size of the intersection divided by the size of the union of the two lists. Hence a Jaccard index of 1 means perfect match between the identified superpixels whereas a Jaccard index of 0 complete disagreement. Note that the Jaccard similarity between a fixed set of size 5 (resp. 10) and a random subset of $\{1, \dots, 60\}$ is equal to 0.05 (resp. 0.06). Here, 60 is the observed average number of superpixels produced by the quickshift algorithms for the images at hand.

Results. Without being a perfect match, we observe a **substantial overlap between the LIME explanations and the approximated explanations** for all the models and datasets that we tried. This is particularly striking for simple models. More precisely, the Jaccard similarities observed are several times higher than what a random guess would produce. This is surprising since we are considering a linear

Table 2: Comparison between LIME and approximated explanations for LSVRC images.

Model	# param.	# layers	acc.	J5	J10
MobileNetV2	3.5M	88	0.90	0.43	0.54
DenseNet121	8.0M	121	0.92	0.42	0.44
InceptionV3	23.9M	159	0.94	0.35	0.36

approximation of highly non-linear functions. As a matter of fact, the exact values of the interpretable coefficients are quite different. Nevertheless, they are sufficiently close so that the sets of superpixels identified by both methods are consistently overlapping.

We notice that this link seems to weaken when the models become too complex, while still a third of identified superpixels are common for InceptionV3. However, visual inspection reveals that the superpixels identified by both methods remain close from each other even when they are distinct (see Figure 6 and additional qualitative results in appendix).

We want to emphasize that, if the model is not smooth, the link between approximate explanations in the sense of Eq. (12) and LIME does not exist anymore. For instance, a random forest model based on CART trees has gradient equal to zero everywhere. Therefore, the integrated gradient is also zero, and $\beta_j^{\text{approx}} = 0$ for any j . We also want to point out that we did not evaluate β^{approx} as an interpretability method. In particular, it could be the case that the associated explanations are of a lesser quality than LIME’s.

Computation time. Setting aside the segmentation step, each run of LIME requires $n = 1000$ queries of the model,

whereas the averaged gradient estimation requires $m = 20$ queries. In the favorable scenario where getting a gradient is as costly as a model query, computing the approximated explanations is much faster than LIME.

6. Conclusion

In this paper, we proposed the first theoretical analysis of LIME for images. We showed that the explanations provided make sense for elementary shape detectors and linear models. As a consequence of this analysis, we discovered that for smooth models the interpretable coefficients of LIME for images resemble to the sum of integrated gradients over the LIME superpixels.

As future work, we plan on tackling more complex models. A starting point is the study of polynomial functions: obtaining an analogous to Proposition 4 would open the door to more precise approximate statements depending on the higher derivatives of f .

References

- A. Adadi and M. Berrada. Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). *IEEE Access*, 6:52138–52160, 2018.
- D. Garreau and U. von Luxburg. Explaining the explainer: A first theoretical analysis of LIME. In *Proceedings of the 33rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1287–1296, 2020a.
- D. Garreau and U. von Luxburg. Looking Deeper into Tabular LIME. *arXiv preprint arXiv:2008.11092*, 2020b.
- R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi. A survey of methods for explaining black box models. *ACM computing surveys (CSUR)*, 51(5):1–42, 2018.
- A. E. Hoerl and R. W. Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.
- G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- A. Krizhevsky, G. Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25: 1097–1105, 2012.
- Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- D. Mardaoui and D. Garreau. An Analysis of LIME for Text Data. In *AISTATS (to appear)*, 2021.
- M. T. Ribeiro, S. Singh, and C. 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.
- O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4510–4520, 2018.
- J. Schmidhuber. Deep learning in neural networks: An overview. *Neural networks*, 61:85–117, 2015.
- K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- J. M. Steele. *The Cauchy-Schwarz master class: an introduction to the art of mathematical inequalities*. Cambridge University Press, 2004.
- M. Sundararajan, A. Taly, and Q. Yan. Axiomatic attribution for deep networks. In *International Conference on Machine Learning*, pages 3319–3328. PMLR, 2017.
- C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- A. Vedaldi and S. Soatto. Quick shift and kernel methods for mode seeking. In *European Conference on Computer Vision*, pages 705–718, 2008.

Supplementary material for the paper: “What does LIME really see in images?”

Organization of the supplementary material

In this appendix, we present the detailed proof of our main results (Theorem 1 and Proposition 2) and additional qualitative results. We follow the proof scheme of Garreau and von Luxburg (2020b). In a nutshell, when $\lambda = 0$, the main problem

$$\hat{\beta}_n^\lambda \in \arg \min_{\beta \in \mathbb{R}^{d+1}} \left\{ \sum_{i=1}^n \pi_i (y_i - \beta^\top z_i)^2 + \lambda \|\beta\|^2 \right\} \quad (13)$$

reduces to least squares, with $\hat{\beta}_n$ given in closed-form by

$$\hat{\beta}_n = (Z^\top W Z)^{-1} Z^\top W y,$$

with $Z \in \{0, 1\}^{n \times d}$ the matrix whose lines are given by the z_i s and W the diagonal matrix such that $W_{i,i} = \pi_i$. Setting $\hat{\Sigma}_n := \frac{1}{n} Z^\top W Z$ and $\hat{\Gamma}_n := \frac{1}{n} Z^\top W y$, the study of $\hat{\beta}_n$ can be split in two parts: the examination of $\hat{\Sigma}_n$ (Section 1), and then that of $\hat{\Gamma}_n$ (Section 2). We put everything together in Section 3, proving the concentration of $\hat{\beta}_n$ and providing the expression of β^f . All technical results are collected in Section 4. Finally, additional qualitative results are presented in Section 5.

1. Study of $\hat{\Sigma}_n$

We start by the study of $\hat{\Sigma}_n$, first computing its limit Σ when $n \rightarrow +\infty$ (Section 1.1). We show that Σ is invertible in closed-form in Section 1.2. We then proceed to show that $\hat{\Sigma}_n$ is concentrated around Σ in Section 1.3. We conclude this section by obtaining a control on the operator norm of Σ^{-1} (Section 1.4), a technical requirement for the proof of the main result.

1.1. Computation of Σ

By definition of Z and W , the matrix $\hat{\Sigma}_n$ can be written

$$\hat{\Sigma} = \begin{pmatrix} \frac{1}{n} \sum_{i=1}^n \pi_i & \frac{1}{n} \sum_{i=1}^n \pi_i z_{i,1} & \cdots & \frac{1}{n} \sum_{i=1}^n \pi_i z_{i,d} \\ \frac{1}{n} \sum_{i=1}^n \pi_i z_{i,1} & \frac{1}{n} \sum_{i=1}^n \pi_i z_{i,1}^2 & \cdots & \frac{1}{n} \sum_{i=1}^n \pi_i z_{i,1} z_{i,d} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{1}{n} \sum_{i=1}^n \pi_i z_{i,d} & \frac{1}{n} \sum_{i=1}^n \pi_i z_{i,1} z_{i,d} & \cdots & \frac{1}{n} \sum_{i=1}^n \pi_i z_{i,d}^2 \end{pmatrix} \in \mathbb{R}^{(d+1) \times (d+1)}.$$

Recall that we defined the random variable z such that z_i is i.i.d. z for any i , as well as π and x the associated weights and perturbed samples. For any $p \geq 0$, we also defined $\alpha_p = \mathbb{E}[\pi \prod_{i=1}^p z_i]$ (Definition 1). Taking the expectation with respect to z in the previous display, we obtain

$$\Sigma_{j,k} = \begin{cases} \alpha_0 & \text{if } j = k = 0, \\ \alpha_1 & \text{if } j = 0 \text{ and } k > 0 \text{ or } j > 0 \text{ and } k = 0 \text{ or } j = k > 0, \\ \alpha_2 & \text{otherwise.} \end{cases}$$

As promised, it is possible to compute the α coefficients in closed-form. Let us denote by S the number of superpixel deletions. Since the coordinates of z are i.i.d. Bernoulli with parameter $1/2$, we deduce that S is a *binomial* random variable

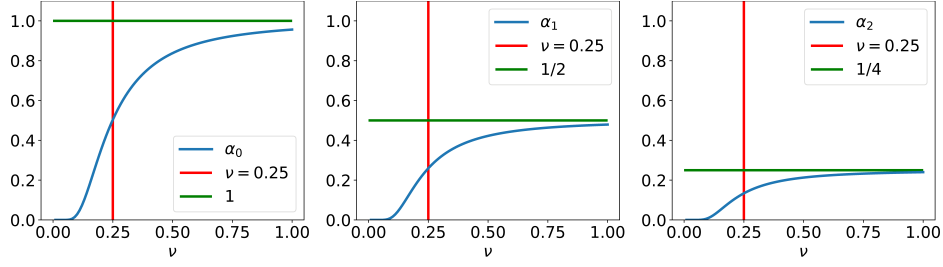


Figure 7: The first three α coefficients as a function of the bandwidth ν for $d = 50$. In green the limit value given by Lemma 1.

of parameters d and $1/2$. Note that, conditionally to $S = s$, $\sum_j z_j = d - s$ and therefore $\pi = \psi(s/d)$ with

$$\forall t \in [0, 1], \quad \psi(t) := \exp\left(\frac{-(1 - \sqrt{1-t})^2}{2\nu^2}\right) \quad (14)$$

as in the paper. As a consequence of these observations, we have:

Proposition 5 (Computation of the α coefficients). *Let $p \geq 0$ be an integer. Then*

$$\alpha_p = \frac{1}{2^d} \sum_{s=0}^d \binom{d-p}{s} \psi(s/d).$$

Proof. We write

$$\begin{aligned} \alpha_p &= \mathbb{E}[\pi z_1 \cdots z_p] \\ &= \sum_{s=0}^d \mathbb{E}_s[\pi z_1 \cdots z_p] \mathbb{P}(S = s) && \text{(law of total expectation)} \\ &= \frac{1}{2^d} \sum_{s=0}^d \binom{d}{s} \mathbb{E}_s[\pi z_1 \cdots z_p | z_1 = 1, \dots, z_p = 1] \mathbb{P}_s(z_1 = 1, \dots, z_p = 1) && (S \sim \mathcal{B}(n, 1/2)) \\ &= \frac{1}{2^d} \sum_{s=0}^d \binom{d}{s} \psi(s/d) \mathbb{P}_s(z_1 = 1, \dots, z_p = 1) && \text{(definition of } \psi) \\ \alpha_p &= \frac{1}{2^d} \sum_{s=0}^d \binom{d}{s} \frac{(d-p)!}{d!} \cdot \frac{(d-s)!}{(d-s-p)!} \psi(s/d) && \text{(Lemma 3)} \end{aligned}$$

We conclude by some algebra. □

It is quite straightforward to compute the limits of the α coefficients when $\nu \rightarrow +\infty$. In fact, since $e^{-1/(2\nu^2)} \leq \psi(t) \leq 1$ for any $\nu > 0$, we have the following bounds on α_p :

Lemma 1 (Bounding the α coefficients). *For any $p \geq 0$, we have*

$$\frac{e^{-\frac{1}{2\nu^2}}}{2^p} \leq \alpha_p \leq \frac{1}{2^p}.$$

In particular, when $\nu \rightarrow +\infty$, we have $\alpha_p \rightarrow \frac{1}{2^p}$ for any $p \geq 0$.

We demonstrate these approximations in Figure 7.

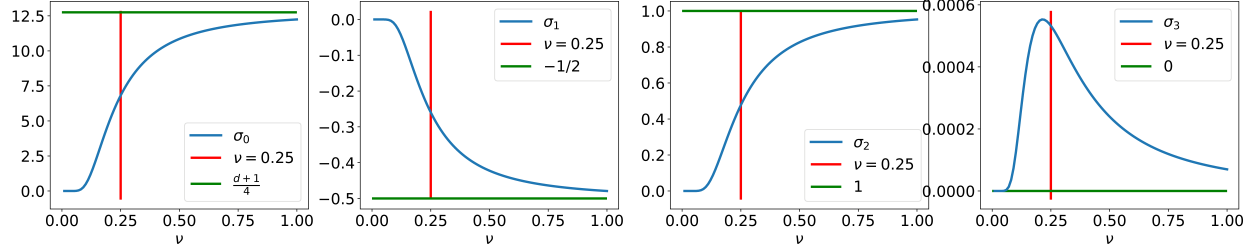


Figure 8: The first four σ coefficients as a function of the bandwidth ν for $d = 50$. In green, the limit values given by Eq. (16).

1.2. σ coefficients

Since the structure of Σ is the same as in the text case (Mardaoui and Garreau, 2021), we can invert it similarly.

Proposition 6 (Inverse of Σ). For any $d \geq 1$, recall that we defined

$$\begin{cases} \sigma_1 &= -\alpha_1, \\ \sigma_2 &= \frac{(d-2)\alpha_0\alpha_2 - (d-1)\alpha_1^2 + \alpha_0\alpha_1}{\alpha_1 - \alpha_2}, \\ \sigma_3 &= \frac{\alpha_1^2 - \alpha_0\alpha_2}{\alpha_1 - \alpha_2}, \end{cases}$$

and $c_d = (d-1)\alpha_0\alpha_2 - d\alpha_1^2 + \alpha_0\alpha_1$. Let us further define $\sigma_0 := (d-1)\alpha_2 + \alpha_1$. Assume that $c_d \neq 0$ and $\alpha_1 \neq \alpha_2$. Then Σ is invertible, and it holds that

$$\Sigma^{-1} = \frac{1}{c_d} \begin{pmatrix} \sigma_0 & \sigma_1 & \sigma_1 & \cdots & \sigma_1 \\ \sigma_1 & \sigma_2 & \sigma_3 & \cdots & \sigma_3 \\ \sigma_1 & \sigma_3 & \sigma_2 & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & \sigma_3 \\ \sigma_1 & \sigma_3 & \cdots & \sigma_3 & \sigma_2 \end{pmatrix} \in \mathbb{R}^{(d+1) \times (d+1)}. \quad (15)$$

From Lemma 1, we deduce

$$\sigma_0 \rightarrow \frac{d+1}{4}, \quad \sigma_1 \rightarrow \frac{-1}{2}, \quad \sigma_2 \rightarrow 1, \quad \sigma_3 \rightarrow 0, \quad \text{and} \quad c_d \rightarrow 1/4. \quad (16)$$

when $\nu \rightarrow +\infty$. We illustrate this in Figure 8. Now, Proposition 6 requires $\alpha_1 \neq \alpha_2$ and $c_d \neq 0$ in order for Σ to be invertible. One of the consequences of the following result is that these conditions are always satisfied.

Proposition 7 (Σ is invertible). Let $d \geq 1$ and $\nu > 0$. Then $\alpha_1 - \alpha_2 \geq e^{\frac{-1}{2\nu^2}}/4$ and $c_d \geq e^{\frac{-1}{2\nu^2}}/4$.

Note that in this case the lower bound obtained on c_d is tight. We show the evolution of c_d with respect to the bandwidth in Figure 9.

Proof. By definition of the α coefficients and Pascal identity, it holds that

$$\alpha_p - \alpha_{p+1} = \frac{1}{2^d} \sum_{s=0}^d \binom{d-p-1}{s-1} \psi\left(\frac{s}{d}\right), \quad (17)$$

for any $p \geq 0$. Since $e^{-1/(2\nu^2)} \leq \psi(t) \leq 1$ for any $1 \leq t \leq d$, we deduce from Eq. (17) that, for any $p \geq 0$,

$$\frac{e^{\frac{-1}{2\nu^2}}}{2^{p+1}} \leq \alpha_p - \alpha_{p+1} \leq \frac{1}{2^{p+1}}. \quad (18)$$

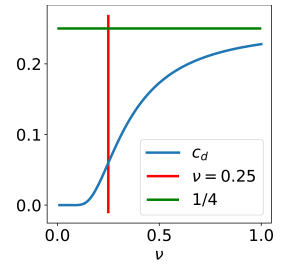


Figure 9: Evolution of c_d with respect to ν when $d = 50$.

We deduce the lower bound on $\alpha_1 - \alpha_2$ by setting $p = 1$ in the previous display.

Let us turn to c_d . We write

$$\begin{aligned} c_d &= d\alpha_1(\alpha_0 - \alpha_1) - (d-1)\alpha_0(\alpha_1 - \alpha_2) \\ &= \frac{1}{4^d} \left[d \cdot \sum_{s=0}^d \binom{d-1}{s} \psi\left(\frac{s}{d}\right) \cdot \sum_{s=0}^d \binom{d-1}{s-1} \psi\left(\frac{s}{d}\right) - (d-1) \cdot \sum_{s=0}^d \binom{d}{s} \psi\left(\frac{s}{d}\right) \cdot \sum_{s=0}^d \binom{d-2}{s-1} \psi\left(\frac{s}{d}\right) \right] \\ &\hspace{20em} \text{(using Eq. (17))} \\ c_d &= \frac{1}{4^d} \left[\sum_{s=0}^d \binom{d-1}{s} \psi\left(\frac{s}{d}\right) \cdot \sum_{s=0}^d s \binom{d}{s} \psi\left(\frac{s}{d}\right) - \sum_{s=0}^d \binom{d}{s} \psi\left(\frac{s}{d}\right) \cdot \sum_{s=0}^d s \binom{d-1}{s} \psi\left(\frac{s}{d}\right) \right], \end{aligned}$$

where we used elementary properties of the binomial coefficients in the last display. For any $0 \leq s \leq d$, let us set

$$A_s := \binom{d-1}{s} \sqrt{\psi\left(\frac{s}{d}\right)}, B_s := s \sqrt{\psi\left(\frac{s}{d}\right)}, C_s := \sqrt{\psi\left(\frac{s}{d}\right)}, \text{ and } D_s := \binom{d}{s} \sqrt{\psi\left(\frac{s}{d}\right)}.$$

With these notation,

$$c_d = \frac{1}{4^d} \left[\sum_s A_s C_s \cdot \sum_s B_s D_s - \sum_s A_s B_s \cdot \sum_s C_s D_s \right].$$

According to the four-letter identity (Proposition 17), we can rewrite c_d as

$$\begin{aligned} c_d &= \frac{1}{4^d} \sum_{s < t} (A_s D_t - A_t D_s)(C_s B_t - C_t B_s) \\ &= \frac{1}{4^d} \sum_{s < t} (t-s) \left(\binom{d-1}{s} \binom{d}{t} - \binom{d-1}{t} \binom{d}{s} \right) \psi\left(\frac{s}{d}\right) \psi\left(\frac{t}{d}\right) \\ c_d &= \frac{1}{d \cdot 4^d} \sum_{s < t} \binom{d}{s} \binom{d}{t} (s-t)^2 \psi\left(\frac{s}{d}\right) \psi\left(\frac{t}{d}\right). \end{aligned}$$

Since $e^{-1/(2\nu^2)} \leq \psi(t) \leq 1$ for any $1 \leq t \leq d$, all that is left to do is to control the double sum. According to Proposition 18, we have

$$\sum_{s < t} \binom{d}{s} \binom{d}{t} (s-t)^2 = d \cdot 4^{d-1}.$$

We deduce that

$$\frac{e^{-\frac{1}{2\nu^2}}}{4} \leq c_d \leq \frac{1}{4}. \tag{19}$$

□

We conclude this section with useful relationships between α and σ coefficients.

Proposition 8 (Useful equalities). *Let α_p, σ_p , and c_d be defined as above. Then it holds that*

$$\sigma_0 \alpha_1 + \sigma_1 \alpha_1 + (d-1) \sigma_1 \alpha_2 = 0, \tag{20}$$

$$\sigma_1 \alpha_1 + \sigma_2 \alpha_1 + (d-1) \sigma_3 \alpha_2 = c_d, \tag{21}$$

$$\sigma_1 \alpha_1 + \sigma_2 \alpha_2 + \sigma_3 \alpha_1 + (d-2) \sigma_3 \alpha_2 = 0, \tag{22}$$

$$\sigma_1 \alpha_0 + \sigma_2 \alpha_1 + (d-1) \sigma_3 \alpha_1 = 0, \tag{23}$$

$$\sigma_0 \alpha_0 + d \sigma_1 \alpha_1 = c_d. \tag{24}$$

Proof. Straightforward from the definitions. □

1.3. Concentration of $\hat{\Sigma}_n$

We now turn to the concentration of $\hat{\Sigma}_n$ around Σ . More precisely, we show that $\hat{\Sigma}_n$ is close to Σ in operator norm, with high probability. Since the definition of $\hat{\Sigma}_n$ is identical to the one in the Tabular LIME case, we can use the proof machinery of Garreau and von Luxburg (2020b).

Proposition 9 (Concentration of $\hat{\Sigma}_n$). *For any $t \geq 0$,*

$$\mathbb{P} \left(\left\| \hat{\Sigma}_n - \Sigma \right\|_{\text{op}} \geq t \right) \leq 4d \cdot \exp \left(\frac{-nt^2}{32d^2} \right).$$

Proof. We can write $\hat{\Sigma} = \frac{1}{n} \sum_i \pi_i Z_i Z_i^\top$. The summands are bounded i.i.d. random variables, thus we can apply the matrix version of Hoeffding inequality. More precisely, the entries of $\hat{\Sigma}_n$ belong to $[0, 1]$ by construction, and Lemma 1 guarantees that the entries of Σ also belong to $[0, 1]$. Therefore, if we set $M_i := \frac{1}{n} \pi_i Z_i Z_i^\top - \Sigma$, then the M_i satisfy the assumptions of Theorem 21 in Garreau and von Luxburg (2020b) and we can conclude since $\frac{1}{n} \sum_i M_i = \hat{\Sigma}_n - \Sigma$. \square

1.4. Control of $\|\Sigma^{-1}\|_{\text{op}}$

In this section, we obtain a control on the operator norm of the inverse covariance matrix. Our strategy is to bound the norm of the σ coefficients. We start with the control of $\alpha_1^2 - \alpha_0 \alpha_2$, a quantity appearing in σ_2 and σ_3 .

Lemma 2 (Control of $\alpha_1^2 - \alpha_0 \alpha_2$). *For any $d \geq 2$, we have*

$$|\alpha_1^2 - \alpha_0 \alpha_2| \leq \frac{1}{2d}.$$

Proof. By definition of the α coefficients, we know that

$$\alpha_1^2 - \alpha_0 \alpha_2 = \frac{1}{4^d} \left[\left(\sum_{s=0}^d \binom{d-1}{s} \psi \left(\frac{s}{d} \right) \right)^2 - \left(\sum_{s=0}^d \binom{d}{s} \psi \left(\frac{s}{d} \right) \right) \cdot \left(\sum_{s=0}^d \binom{d-2}{s} \psi \left(\frac{s}{d} \right) \right) \right].$$

Let us ignore the $1/4^d$ normalization for now, and set $w_s := \binom{d}{s} \psi \left(\frac{s}{d} \right)$. Elementary manipulations of the binomial coefficients allow us to rewrite the previous display as

$$\left(\sum_{s=0}^d \frac{d-s}{d} w_s \right)^2 - \left(\sum_{s=0}^d w_s \right) \cdot \left(\sum_{s=0}^d \frac{d-s}{d} \cdot \frac{d-s-1}{d-1} w_s \right). \quad (25)$$

Let us notice that

$$\frac{d-s}{d} - \frac{d-s-1}{d-1} = \frac{s}{d(d-1)}.$$

Thus we can split Eq. (25) in two parts.

The first part is reminiscent of the Cauchy-Schwarz-like expression that appears in the proof of Proposition 7:

$$\left(\sum_{s=0}^d \frac{d-s}{d} w_s \right)^2 - \left(\sum_{s=0}^d w_s \right) \cdot \left(\sum_{s=0}^d \frac{(d-s)^2}{d^2} w_s \right). \quad (26)$$

We use, again, the four letter identity (Proposition 17) to bound this term. Namely, for any $0 \leq s \leq d$, let us set

$$A_s = B_s := \frac{d-s}{d} \sqrt{w_s}, \quad \text{and} \quad C_s = D_s := \sqrt{w_s}.$$

Then we can rewrite Eq. (26) as

$$\sum_{s < t} (A_s D_t - A_t D_s)(C_s B_t - C_t B_s) = \frac{-1}{d^2} \sum_{s < t} (t-s)^2 \binom{d}{s} \binom{d}{t} \psi \left(\frac{s}{d} \right) \psi \left(\frac{t}{d} \right). \quad (27)$$

According to Proposition 18, Eq. (27) is bounded by $d \cdot 4^{d-1}/d^2 = 4^{d-1}/d$.

The second part of Eq. (25) reads

$$\left(\sum_{s=0}^d w_s \right) \cdot \left(\sum_{s=0}^d \frac{d-s}{d} \cdot \frac{s}{d(d-1)} w_s \right).$$

Since ψ is bounded by 1, coming back to the definition of the w_s , it is straightforward to show that $|\sum_s w_s| \leq 2^d$ and that $|\sum_s s(d-s)w_s| \leq d(d-1)2^{d-2}$. We deduce that (the absolute value of) this second term is upper bounded by $4^{d-1}/d$.

Putting together the bounds obtained on both terms and renormalizing by 4^d , we obtain that

$$|\alpha_1^2 - \alpha_0\alpha_2| \leq \frac{1}{4^d} \left[\frac{4^{d-1}}{d} + \frac{4^{d-1}}{d} \right] = \frac{1}{2d}.$$

□

We now have everything we need to provide reasonably tight upper bounds for the σ coefficients.

Proposition 10 (Bounding the σ coefficients). *Let $d \geq 2$. Then the following holds:*

$$|\sigma_0| \leq \frac{3d}{4}, \quad |\sigma_1| \leq \frac{1}{2}, \quad |\sigma_2| \leq 2e^{\frac{1}{2\nu^2}}, \quad \text{and} \quad |\sigma_3| \leq \frac{2e^{\frac{1}{2\nu^2}}}{d}.$$

Proof. From Lemma 1 and the definition of σ_0 , we have

$$|\sigma_0| = |(d-1)\alpha_2 + \alpha_1| \leq \frac{d-1}{4} + \frac{1}{2} = \frac{d+3}{4}.$$

We deduce the first result since $d \geq 2$. Next, since $\sigma_1 = -\alpha_1$, we obtain $|\sigma_1| \leq 1/2$ directly from Lemma 1. Regarding the last two coefficients, recall that Proposition 7 guarantees that their common denominator $\alpha_1 - \alpha_2$ is lower bounded by $e^{\frac{-1}{2\nu^2}}/4$. Since

$$(d-2)\alpha_0\alpha_2 - (d-1)\alpha_1^2 + \alpha_0\alpha_1 = c_d + \alpha_1^2 - \alpha_0\alpha_2,$$

we can write $\sigma_2 = (c_d + \alpha_1^2 - \alpha_0\alpha_2)/(\alpha_1 - \alpha_2)$ and deduce that

$$|\sigma_2| \leq \frac{1/4 + 1/(2d)}{e^{\frac{-1}{2\nu^2}}/4} \leq 2e^{\frac{1}{2\nu^2}},$$

since, according to Eq. (19), $c_d \leq 1/4$ and $\alpha_1^2 - \alpha_0\alpha_2 \leq 1/(2d)$ according to Lemma 2. Finally, we write

$$|\sigma_3| = \left| \frac{\alpha_1^2 - \alpha_0\alpha_2}{\alpha_1 - \alpha_2} \right| \leq \frac{1/(2d)}{e^{\frac{-1}{2\nu^2}}/4} = \frac{2e^{\frac{1}{2\nu^2}}}{d}.$$

□

The bounds obtained in Proposition 10 immediately translate into a control of the Frobenius norm of Σ^{-1} , which in turn yields a control over the operator norm of Σ^{-1} , as promised.

Corollary 1 (Control of $\|\Sigma^{-1}\|_{\text{op}}$). *Let $d \geq 2$. Then $\|\Sigma^{-1}\|_{\text{op}} \leq 8de^{\frac{1}{\nu^2}}$.*

Proof. Using Proposition 10, we write

$$\begin{aligned} \|\Sigma^{-1}\|_{\text{F}}^2 &= \frac{1}{c_d^2} [\sigma_0^2 + 2d\sigma_1^2 + d\sigma_2^2 + (d^2 - d)\sigma_3^2] \\ &\leq 16e^{\frac{1}{\nu^2}} \left[\frac{9d^2}{16} + \frac{2d}{4} + 4de^{\frac{1}{\nu^2}} + 4e^{\frac{1}{\nu^2}} \right] \\ &\leq 61d^2e^{\frac{2}{\nu^2}}, \end{aligned}$$

where we used $d \geq 2$ in the last display. Since the operator norm is upper bounded by the Frobenius norm, we conclude observing that $\sqrt{61} \leq 8$. □

2. Study of $\hat{\Gamma}_n$

We now turn to the study of $\hat{\Gamma}_n$. We start by computing the limiting expression. Recall that we defined $\hat{\Gamma}_n = \frac{1}{n} Z^\top W y$, where $y \in \mathbb{R}^{d+1}$ is the random vector defined coordinate-wise by $y_i = f(x_i)$. From the definition of $\hat{\Gamma}_n$, it is straightforward that

$$\hat{\Gamma}_n = \begin{pmatrix} \frac{1}{n} \sum_{i=1}^n \pi_i f(x_i) \\ \frac{1}{n} \sum_{i=1}^n \pi_i z_{i,1} f(x_i) \\ \vdots \\ \frac{1}{n} \sum_{i=1}^n \pi_i z_{i,d} f(x_i) \end{pmatrix} \in \mathbb{R}^{d+1}.$$

As a consequence, if we define $\Gamma^f := \mathbb{E}[\hat{\Gamma}_n]$, it holds that

$$\Gamma^f = \begin{pmatrix} \mathbb{E}[\pi f(x)] \\ \mathbb{E}[\pi z_1 f(x)] \\ \vdots \\ \mathbb{E}[\pi z_d f(x)] \end{pmatrix}. \quad (28)$$

We specialize Eq. (28) to shape detectors in Section 2.1 and linear models in Section 2.2. The concentration of $\hat{\Gamma}_n$ around Γ is obtained in Section 2.3.

2.1. Shape detectors

Recall that we defined

$$\forall x \in [0, 1]^D, \quad f(x) = \prod_{u \in \mathcal{S}} \mathbf{1}_{x_u > \tau}, \quad (29)$$

with $\mathcal{S} = \{u_1, \dots, u_q\}$ a fixed set of pixels indices and $\tau \in (0, 1)$ a threshold. As in the paper, let us define $E = \{j \text{ s.t. } J_j \cap \mathcal{S} \neq \emptyset\}$ denote the set of superpixels intersecting the shape, and

$$E_+ = \{j \in E \text{ s.t. } \bar{\xi}_j > \tau\} \quad \text{and} \quad E_- = \{j \in E \text{ s.t. } \bar{\xi}_j \leq \tau\}.$$

We also defined

$$\mathcal{S}_+ = \{u \in \mathcal{S} \text{ s.t. } \xi_u > \tau\} \quad \text{and} \quad \mathcal{S}_- = \{u \in \mathcal{S} \text{ s.t. } \xi_u \leq \tau\}.$$

In the main paper, we made the following simplifying assumption:

$$\forall j \in E_+, \quad J_j \cap \mathcal{S}_- = \emptyset. \quad (30)$$

This is not the case here. Unfortunately, without this assumption, the expression of Γ^f is slightly more complicated and we need to generalize the definition of the α coefficients.

Definition 3 (Generalized α coefficients). For any p, q such that $p + q \leq d$, we define

$$\alpha_{p,q} := \mathbb{E}[\pi z_1 \cdots z_p \cdot (1 - z_{p+1}) \cdots (1 - z_{p+q})]. \quad (31)$$

We notice that, for any $1 \leq p \leq d$, $\alpha_{p,0} = \alpha_p$. As it is the case with α coefficients, the generalized α coefficients can be computed in closed-form:

Proposition 11 (Computation of the generalized α coefficients). Let p, q such that $p + q \leq d$. Then

$$\alpha_{p,q} = \frac{1}{2^d} \sum_{s=0}^d \binom{d-p-q}{s-q} \psi\left(\frac{s}{d}\right).$$

Proof. We follow the proof of Proposition 5.

$$\begin{aligned}
 \alpha_{p,q} &= \mathbb{E} [\pi z_1 \cdots z_p \cdot (1 - z_{p+1}) \cdots (1 - z_{p+q})] \\
 &= \sum_{s=0}^d \mathbb{E}_s [\pi z_1 \cdots z_p \cdot (1 - z_{p+1}) \cdots (1 - z_{p+q})] \cdot \mathbb{P}(S = s) \\
 &= \frac{1}{2^d} \sum_{s=0}^d \binom{d}{s} \psi\left(\frac{s}{d}\right) \mathbb{P}_s(z_1 = \cdots = z_p = 1, z_{p+1} = \cdots = z_{p+q} = 0) \\
 &= \frac{1}{2^d} \sum_{s=0}^d \binom{d}{s} \psi\left(\frac{s}{d}\right) \binom{d-p-q}{s-q} \binom{d}{s} \quad (\text{Lemma 4}) \\
 \alpha_{p,q} &= \frac{1}{2^d} \sum_{s=0}^d \binom{d-p-q}{s-q} \psi\left(\frac{s}{d}\right).
 \end{aligned}$$

□

Notice that the expression of $\alpha_{p,q}$ coincide with that of α_p when $q = 0$. We can now give the expression of Γ^f for an elementary shape detector in the general case.

Proposition 12 (Computation of Γ^f , elementary shape detector). *Assume that f is written as in Eq. (29). Assume that for any $j \in E_-$, $J_j \cap \mathcal{S}_- = \emptyset$ (otherwise $\Gamma^f = 0$). Let $p := |E_-|$ and $q := |\{j \in E_+, J_j \cap \mathcal{S}_- \neq \emptyset\}|$. Then $\mathbb{E}[\pi f(x)] = \alpha_{p,q}$ and*

$$\mathbb{E}[\pi z_j f(x)] = \begin{cases} 0 & \text{if } j \in \{j \in E_+ \text{ s.t. } J_j \cap \mathcal{S}_- \neq \emptyset\}, \\ \alpha_{p,q} & \text{if } j \in E_-, \\ \alpha_{p+1,q} & \text{otherwise.} \end{cases}$$

Taking $q = 0$ (a consequence of Eq. (30)) in Proposition 12 directly yields $\mathbb{E}[\pi f(x)] = \alpha_p$ and

$$\mathbb{E}[\pi z_j f(x)] = \begin{cases} \alpha_p & \text{if } j \in E_-, \\ \alpha_{p+1} & \text{otherwise.} \end{cases}$$

Proof. We notice that, for any $u \in J_j$,

$$x_u = z_j \xi_u + (1 - z_j) \bar{\xi}_u.$$

There are four cases to consider when deciding whether $x_u > \tau$ or not:

- $\xi_u > \tau$ and $\bar{\xi}_u > \tau$, that is, $j \in E_+$ and $u \in J_j \cap \mathcal{S}_+$. Then $x_u > \tau$ a.s.;
- $\xi_u \leq \tau$ and $\bar{\xi}_u > \tau$, that is, $j \in E_+$ and $u \in J_j \cap \mathcal{S}_-$. Then $x_u > \tau$ if, and only if, $z_j = 0$;
- $\xi_u > \tau$ and $\bar{\xi}_u \leq \tau$, that is, $j \in E_-$ and $u \in J_j \cap \mathcal{S}_+$. Then $x_u > \tau$ if, and only if, $z_j = 1$;
- $\xi_u \leq \tau$ and $\bar{\xi}_u \leq \tau$, that is, $j \in E_-$ and $u \in J_j \cap \mathcal{S}_-$. Then $x_u \leq \tau$ a.s., but this last case cannot happen since we assume that for any $j \in E_-$, $J_j \cap \mathcal{S}_- = \emptyset$.

This case separation allows us to rewrite $f(x)$ as

$$\begin{aligned}
 f(x) &= \prod_{u \in \mathcal{S}} \mathbf{1}_{x_u > \tau} \quad (\text{Eq. (29)}) \\
 &= \prod_{j \in E_+} \prod_{u \in J_j \cap \mathcal{S}_-} (1 - z_j) \cdot \prod_{j \in E_-} \prod_{u \in J_j \cap \mathcal{S}_+} z_j
 \end{aligned}$$

Since we assumed that for any $j \in E_-$, $J_j \cap \mathcal{S}_- = \emptyset$, then for any $j \in E_-$, $J_j \cap \mathcal{S}_+ \neq \emptyset$. Thus the rightmost inner products are never empty, and since $z_j \in \{0, 1\}$ a.s., we deduce that there are p terms in the rightmost product. By definition of q ,

and again since $1 - z_j \in \{0, 1\}$ a.s., there are q terms in the leftmost product. By definition of E_+ and E_- , these products do not have any common terms. We deduce that $\mathbb{E}[\pi f(x)] = \alpha_{p,q}$ by definition of the generalized α coefficients.

When computing $\mathbb{E}[\pi z_j f(x)]$, there are several possibilities. First, if $j \in \{j \in E_+ \text{ s.t. } J_j \cap S_- \neq \emptyset\}$, since $z_j(1 - z_j) = 0$ a.s., we deduce that $\mathbb{E}[\pi z_j f(x)] = 0$. Second, if $j \in E_-$, since $z_j^2 = z_j$, we recover $\mathbb{E}[\pi z_j f(x)] = \mathbb{E}[\pi f(x)] = \alpha_{p,q}$. Finally, if j does not belong to one of these sets, then the rightmost product gains one additional term and we obtain $\alpha_{p+1,q}$. \square

2.2. Linear model

In this section, we compute Γ^f for a linear f . As in the paper, we define

$$f(x) = \sum_{u=1}^D \lambda_u x_u, \quad (32)$$

with $\lambda_1, \dots, \lambda_D \in \mathbb{R}$ arbitrary coefficients. By linearity, we just have to look into the case $f : x \mapsto x_u$ where $u \in \{1, \dots, D\}$ is a fixed pixel index.

Proposition 13 (Computation of Γ^f , linear case). *Assume that f is defined as in Eq. (32) and $u \in J_j$. Then*

$$\mathbb{E}[\pi x_u] = \alpha_1(\xi_u - \bar{\xi}_u) + \alpha_0 \bar{\xi}_u,$$

$$\mathbb{E}[\pi z_j x_u] = \alpha_1(\xi_u - \bar{\xi}_u) + \alpha_1 \bar{\xi}_u,$$

and, for any $k \neq j$,

$$\mathbb{E}[\pi z_k x_u] = \alpha_2(\xi_u - \bar{\xi}_u) + \alpha_1 \bar{\xi}_u.$$

Proof. As in the proof of Proposition 12, we notice that

$$x_u = z_j \xi_u + (1 - z_j) \bar{\xi}_u.$$

Then we write

$$\begin{aligned} \mathbb{E}[\pi x_u] &= \mathbb{E}[\pi(z_j \xi_u + (1 - z_j) \bar{\xi}_u)] \\ &= \mathbb{E}[\pi z_j(\xi_u - \bar{\xi}_u) + \pi \bar{\xi}_u] \\ \mathbb{E}[\pi x_u] &= \alpha_1(\xi_u - \bar{\xi}_u) + \alpha_0 \bar{\xi}_u, \end{aligned}$$

where we used the definition of the α coefficients. Now let us compute $\mathbb{E}[\pi z_j f(x)]$:

$$\begin{aligned} \mathbb{E}[\pi z_j x_u] &= \mathbb{E}[\pi z_j(z_j \xi_u + (1 - z_j) \bar{\xi}_u)] \\ &= \mathbb{E}[\pi z_j((\xi_u - \bar{\xi}_u) z_j + \bar{\xi}_u)] \quad (z_j \in \{0, 1\} \text{ a.s.}) \\ \mathbb{E}[\pi z_j x_u] &= \alpha_1(\xi_u - \bar{\xi}_u) + \alpha_1 \bar{\xi}_u. \end{aligned}$$

And finally, for any $k \neq j$,

$$\begin{aligned} \mathbb{E}[\pi z_k x_u] &= \mathbb{E}[\pi z_k((\xi_u - \bar{\xi}_u) z_j + \bar{\xi}_u)] \\ &= \alpha_2(\xi_u - \bar{\xi}_u) + \alpha_1 \bar{\xi}_u. \end{aligned}$$

\square

2.3. Concentration of $\hat{\Gamma}_n$

We now show that $\hat{\Gamma}_n$ is concentrated around Γ^f . Since the expression of $\hat{\Gamma}_n$ is the same than in the tabular case, and we assume that f is bounded on the support of x , the same reasoning as in the proof of Proposition 24 in Garreau and von Luxburg (2020b) can be applied.

Proposition 14 (Concentration of $\hat{\Gamma}_n$). *Assume that f is bounded by $M > 0$ on $\text{Supp}(x)$. Then, for any $t > 0$, it holds that*

$$\mathbb{P}\left(\|\hat{\Gamma}_n - \Gamma^f\| \geq t\right) \leq 4d \exp\left(\frac{-nt^2}{32Md^2}\right).$$

Proof. Since f is bounded by M on $\text{Supp}(x)$, it holds that $|f(x)| \leq M$ almost surely. We can then proceed as in the proof of Proposition 24 in Garreau and von Luxburg (2020b). \square

3. The study of β^f

3.1. Concentration of $\hat{\beta}_n$

In this section we show the concentration of $\hat{\beta}_n$ (Theorem 1 in the paper). The proof scheme follows closely that of Garreau and von Luxburg (2020b).

Theorem 2 (Concentration of $\hat{\beta}_n$). *Assume that f is bounded by a constant M on the unit cube $[0, 1]^D$. Let $\epsilon > 0$ and $\eta \in (0, 1)$. Let d be the number of superpixels used by LIME. Then, there exists $\beta^f \in \mathbb{R}^{d+1}$ such that, for every*

$$n \geq \left\lceil \max\left(2^{15}d^4e^{\frac{2}{\nu^2}}, \frac{2^{21}d^7 \max(M, M^2)e^{\frac{4}{\nu^2}}}{\epsilon^2}\right) \log \frac{8d}{\eta} \right\rceil,$$

we have $\mathbb{P}(\|\hat{\beta}_n - \beta^f\| \geq \epsilon) \leq \eta$.

Proof. As in Garreau and von Luxburg (2020b), the key idea of the proof is to notice that

$$\|\hat{\beta}_n - \beta^f\| \leq 2\|\Sigma^{-1}\|_{\text{op}}\|\hat{\Gamma} - \Gamma^f\| + 2\|\Sigma^{-1}\|_{\text{op}}^2\|\Gamma^f\|\|\hat{\Sigma} - \Sigma\|_{\text{op}}, \quad (33)$$

provided that (i) $\|\Sigma^{-1}(\hat{\Sigma} - \Sigma)\|_{\text{op}} \leq 0.32$ (this is Lemma 27 in Garreau and von Luxburg (2020b)). We are going to build an event of probability at least $1 - \eta$ such that $\hat{\Sigma}_n$ is close to Σ and $\hat{\Gamma}_n$ is close from Γ^f . The deterministic bound obtained on $\|\Sigma^{-1}\|_{\text{op}}$ together with the boundedness of f will allow us to show that (ii) $\|\Sigma^{-1}\|_{\text{op}}\|\hat{\Gamma} - \Gamma^f\| \leq \epsilon/4$ and (iii) $\|\Sigma^{-1}\|_{\text{op}}^2\|\Gamma^f\|\|\hat{\Sigma} - \Sigma\|_{\text{op}} \leq \epsilon/4$.

We first show (i). Let us set $n_1 := \left\lceil 2^{15}d^4e^{\frac{2}{\nu^2}} \log \frac{8d}{\eta} \right\rceil$ and $t_1 := \frac{1}{25de^{\frac{1}{\nu^2}}}$. According to Proposition 9, for any $n \geq n_1$,

$$\mathbb{P}\left(\left\|\hat{\Sigma}_n - \Sigma\right\|_{\text{op}} \geq t_1\right) \leq 4d \cdot \exp\left(\frac{-n_1 t_1^2}{32d^2}\right) \leq \frac{\eta}{2}.$$

Moreover, we know that $\|\Sigma^{-1}\|_{\text{op}} \leq 8de^{\frac{1}{\nu^2}}$ (Corollary 1). Since the operator norm is sub-multiplicative, with probability greater than $1 - \eta/2$, we have

$$\left\|\Sigma^{-1}(\hat{\Sigma}_n - \Sigma)\right\|_{\text{op}} \leq \|\Sigma^{-1}\|_{\text{op}} \cdot \left\|\hat{\Sigma}_n - \Sigma\right\|_{\text{op}} \leq 8de^{\frac{1}{\nu^2}} \cdot t_1 = 0.32.$$

Now let us show (ii). Let us define $n_2 := \left\lceil \frac{2^{15}Md^4e^{\frac{2}{\nu^2}}}{\epsilon^2} \log \frac{8d}{\eta} \right\rceil$ and $t_2 := \frac{\epsilon}{32de^{\frac{1}{\nu^2}}}$. According to Proposition 14, for any $n \geq n_2$, we have

$$\mathbb{P}\left(\left\|\hat{\Gamma}_n - \Gamma\right\| \geq t_2\right) \leq 4d \cdot \exp\left(\frac{-n_2 t_2^2}{32Md^2}\right) \leq \frac{\eta}{2}.$$

Recall that $\|\Sigma^{-1}\|_{\text{op}} \leq 8de^{\frac{1}{\nu^2}}$ (Corollary 1): with probability higher than $1 - \eta/2$,

$$\|\Sigma^{-1}\|_{\text{op}} \cdot \left\|\hat{\Gamma}_n - \Gamma^f\right\| \leq 8de^{\frac{1}{\nu^2}} \cdot t_2 = \frac{\epsilon}{4}.$$

Finally let us show (iii). Let us define $n_3 := \left\lceil \frac{2^{21} d^7 M^2 e^{\frac{4}{\nu^2}}}{\epsilon^2} \log \frac{8d}{\eta} \right\rceil$ and $t_3 := \frac{\epsilon}{2^8 M d^{5/2} e^{\frac{2}{\nu^2}}}$. According to Proposition 9, for any $n \geq n_3$, we have

$$\mathbb{P} \left(\left\| \hat{\Sigma}_n - \Sigma \right\|_{\text{op}} \geq t_3 \right) \leq 4d \cdot \exp \left(\frac{-n_3 t_3^2}{32d^2} \right) \leq \frac{\eta}{2}.$$

Since f is bounded by M , it is straightforward to show that $\left\| \hat{\Gamma}^f \right\| \leq M \cdot d^{1/2}$. Moreover, recall that $\left\| \Sigma^{-1} \right\|_{\text{op}}^2 \leq 64d^2 e^{\frac{2}{\nu^2}}$. We deduce that, with probability at least $\eta/2$,

$$\left\| \Sigma^{-1} \right\|_{\text{op}}^2 \cdot \left\| \Gamma^f \right\| \cdot \left\| \hat{\Sigma}_n - \Sigma \right\|_{\text{op}} \leq 64d^2 e^{\frac{2}{\nu^2}} \cdot M d^{1/2} \cdot t_3 = \frac{\epsilon}{4}.$$

Finally, we notice that both n_2 and n_3 are smaller than

$$n_4 := \left\lceil \frac{2^{21} d^7 \max(M, M^2) e^{\frac{4}{\nu^2}}}{\epsilon^2} \log \frac{8d}{\eta} \right\rceil.$$

Thus (ii) and (ii) simultaneously happen on an event of probability greater than $\eta/2$ when n is larger than n_4 . We conclude by a union bound argument. \square

3.2. General expression of β^f

We are now able to recover Proposition 2 of the paper: the expression of β^f is obtained simply by multiplying Eq. (15) and (28). We also give the value of the intercept (β_0 with our notation), which is omitted in the paper for simplicity's sake.

Corollary 2 (Computation of β^f). *Under the assumptions of Theorem 2.*

$$\beta_0^f = c_d^{-1} \left\{ \sigma_0 \mathbb{E} [\pi f(x)] + \sigma_1 \sum_{j=1}^d \mathbb{E} [\pi z_j f(x)] \right\}, \quad (34)$$

and, for any $1 \leq j \leq d$,

$$\beta_j^f = c_d^{-1} \left\{ \sigma_1 \mathbb{E} [\pi f(x)] + \sigma_2 \mathbb{E} [\pi z_j f(x)] + \sigma_3 \sum_{\substack{k=1 \\ k \neq j}}^d \mathbb{E} [\pi z_k f(x)] \right\}. \quad (35)$$

3.3. Shape detectors

We now specialize Corollary 2 to the case of elementary shape detectors.

Proposition 15 (Expression of β^f , shape detector). *Let f be written as in Eq. (29). Assume that for any $j \in E_-$, $J_j \cap \mathcal{S}_- = \emptyset$ (otherwise $\beta^f = 0$). Let p and q as before. Then*

$$\beta_0^f = c_d^{-1} \{ \sigma_0 \alpha_{p,q} + p \sigma_1 \alpha_{p,q} + (d - p - q) \alpha_{p+1,q} \},$$

for any $j \in E_-$,

$$\beta_j^f = c_d^{-1} \{ \sigma_1 \alpha_{p,q} + \sigma_2 \alpha_{p,q} + (p - 1) \sigma_2 \alpha_{p,q} + (d - p - q) \sigma_3 \alpha_{p+1,q} \},$$

for any $j \in E_+$ such that $J_j \cap \mathcal{S}_- \neq \emptyset$,

$$\beta_j^f = c_d^{-1} \{ \sigma_1 \alpha_{p,q} + p \sigma_3 \alpha_{p,q} + (d - p - q) \alpha_{p+1,q} \},$$

and

$$\beta_j^f = c_d^{-1} \{ \sigma_1 \alpha_{p,q} + \sigma_2 \alpha_{p+1,q} + p \sigma_3 \alpha_{p,q} + (d - p - q - 1) \sigma_3 \alpha_{p+1,q} \}$$

otherwise.

Proof. Straightforward from Corollary 2 and Proposition 12. \square

Note that taking $q = 0$ in Proposition 15 yields Proposition XX of the paper.

3.4. Linear models

We deduce from Proposition 13 the expression of β^f for linear models. Let us define M_j the binary mask associated to superpixel J_j and let \circ be the termwise product.

Proposition 16 (Computation of β^f , linear case). *Assume that f is defined as in Eq. (32). Then*

$$\beta_0^f = \sum_{u=1}^D \lambda_u \bar{\xi}_u = f(\bar{\xi}),$$

and, for any $1 \leq j \leq d$,

$$\beta_j^f = \sum_{u \in J_j} \lambda_u (\xi_u - \bar{\xi}_u) = f(M_j \circ (\xi - \bar{\xi})).$$

It is interesting to compute prediction of the surrogate model at ξ :

$$\beta_0^f + \beta_1^f + \dots + \beta_d^f = f(\bar{\xi}) + f(M_1 \circ (\xi - \bar{\xi})) + \dots + f(M_d \circ (\xi - \bar{\xi})) = f(\xi).$$

Thus in the case of linear models, the limit explanation is faithful.

Proof. By linearity, we can start by computing β^f for the function $x \mapsto x_u$. Assume that $j \in \{1, \dots, d\}$ is such that $u \in J_j$. According to Corollary 2 and Proposition 13,

$$\begin{aligned} \beta_0^f &= \frac{1}{c_d} \left\{ \sigma_0 \mathbb{E}[\pi f(x)] + \sigma_1 \sum_{j=1}^d \mathbb{E}[\pi z_j f(x)] \right\} \\ &= \frac{1}{c_d} \left\{ \sigma_0 (\alpha_1 (\xi_u - \bar{\xi}_u) + \alpha_0 \bar{\xi}_u) + \sigma_1 (\alpha_1 (\xi_u - \bar{\xi}_u) + \alpha_1 \bar{\xi}_u) + (d-1) \sigma_1 (\alpha_2 (\xi_u - \bar{\xi}_u) + \alpha_1 \bar{\xi}_u) \right\} \\ &= \frac{1}{c_d} \left\{ (\sigma_0 \alpha_1 + \sigma_1 \alpha_1 + (d-1) \sigma_1 \alpha_2) (\xi_u - \bar{\xi}_u) + (\sigma_0 \alpha_0 + d \sigma_1 \alpha_1) \bar{\xi}_u \right\} \\ \beta_0^f &= \bar{\xi}_u, \end{aligned}$$

where we used Eqs. (20) and (24) in the last display.

$$\begin{aligned} \beta_j^f &= \frac{1}{c_d} \left\{ \sigma_1 \mathbb{E}[\pi f(x)] + \sigma_2 \mathbb{E}[\pi z_j f(x)] + \sigma_3 \sum_{\substack{k=1 \\ k \neq j}}^d \mathbb{E}[\pi z_k f(x)] \right\} \\ &= \frac{1}{c_d} \left\{ \sigma_1 (\alpha_1 (\xi_u - \bar{\xi}_u) + \alpha_0 \bar{\xi}_u) + \sigma_2 (\alpha_1 (\xi_u - \bar{\xi}_u) + \alpha_1 \bar{\xi}_u) + (d-1) \sigma_3 (\alpha_2 (\xi_u - \bar{\xi}_u) + \alpha_1 \bar{\xi}_u) \right\} \\ &= \frac{1}{c_d} \left\{ (\sigma_1 \alpha_1 + \sigma_2 \alpha_1 + (d-1) \sigma_3 \alpha_2) (\xi_u - \bar{\xi}_u) + (\sigma_1 \alpha_0 + \sigma_2 \alpha_1 + (d-1) \sigma_3 \alpha_1) \bar{\xi}_u \right\} \\ \beta_j^f &= \xi_u - \bar{\xi}_u, \end{aligned}$$

where we used Eqs. (21) and (23) in the last display. Finally, let $k \neq j$:

$$\begin{aligned} \beta_k^f &= \frac{1}{c_d} \left\{ \sigma_1 \mathbb{E}[\pi f(x)] + \sigma_2 \mathbb{E}[\pi z_k f(x)] + \sigma_3 \sum_{\substack{k'=1 \\ k' \neq j, k}}^d \mathbb{E}[\pi z_{k'} f(x)] \right\} \\ &= \frac{1}{c_d} \left\{ \sigma_1 (\alpha_1 (\xi_u - \bar{\xi}_u) + \alpha_0 \bar{\xi}_u) + \sigma_2 (\alpha_2 (\xi_u - \bar{\xi}_u) + \alpha_1 \bar{\xi}_u) + \sigma_3 (\alpha_1 (\xi_u - \bar{\xi}_u) + \alpha_1 \bar{\xi}_u) \right. \\ &\quad \left. + (d-2) \sigma_3 (\alpha_2 (\xi_u - \bar{\xi}_u) + \alpha_1 \bar{\xi}_u) \right\} \\ &= \frac{1}{c_d} \left\{ (\sigma_1 \alpha_1 + \sigma_2 \alpha_2 + \sigma_3 \alpha_1 + (d-2) \sigma_3 \alpha_2) (\xi_u - \bar{\xi}_u) + (\sigma_1 \alpha_0 + \sigma_2 \alpha_1 + (d-1) \sigma_3 \alpha_1) \bar{\xi}_u \right\} \\ \beta_k^f &= 0, \end{aligned}$$

where we used Eqs. (22) and (23) in the last display. We deduce the result by linearity. \square

4. Technical results

4.1. Probability computations

In this section we collect all elementary probability computations necessary for the computation of the α coefficients and the generalized α coefficients.

Lemma 3 (Activated only). *Let $p \geq 0$ be an integer. Then*

$$\mathbb{P}_s(z_1 = 1, \dots, z_p = 1) = \frac{(d-p)!}{d!} \cdot \frac{(d-s)!}{(d-s-p)!}.$$

Proof. Conditionally to $S = s$, the choice of S is uniform among all subsets of $\{1, \dots, d\}$. Therefore we recover the proof of Lemma 4 in Mardaoui and Garreau (2021). \square

The following lemma is a slight generalization, which coincides when $q = 0$.

Lemma 4 (Activated and deactivated). *Let p, q be integers. Then*

$$\mathbb{P}_s(z_1 = \dots = z_p = 1, z_{p+1} = \dots = z_{p+q} = 0) = \binom{d-p-q}{s-q} \binom{d}{s}^{-1}.$$

Proof. Conditionally to $S = s$, the deletions are uniformly distributed. Therefore, the total number of cases is $\binom{d}{s}$. Now, the favorable cases correspond to superpixels $p+1, \dots, p+q$ deleted: these are q fixed deletions. We also need to have superpixels $1, \dots, p$ activated, these are p indices that are not available to deletions. In total, we need to place $s-q$ deletions among $d-p-q$ possibilities. We deduce the result. \square

4.2. Algebraic identities

In this section we collect some identities used throughout the proofs.

Proposition 17 (Four letter identity). *Let A, B, C , and D be four finite sequences of real numbers. Then it holds that*

$$\sum_j A_j C_j \cdot \sum_j B_j D_j - \sum_j A_j B_j \cdot \sum_j C_j D_j = \sum_{j < k} (A_j D_k - A_k D_j)(C_j B_k - C_k B_j).$$

Proof. See the proof of Exercise 3.7 in Steele (2004). \square

Proposition 18 (A combinatorial identity). *Let $d \geq 1$ be an integer. Then*

$$V_d := \sum_{j < k} \binom{d}{j} \binom{d}{k} (j-k)^2 = d \cdot 4^{d-1}.$$

Proof. We first notice that

$$\begin{aligned} V_d &= \frac{1}{2} \sum_{j,k} \binom{d}{j} \binom{d}{k} (j-k)^2 && \text{(by symmetry)} \\ &= \sum_{j,k} \binom{d}{j} \binom{d}{k} k^2 - \sum_{j,k} \binom{d}{j} \binom{d}{k} jk && \text{(developing the square)} \\ &= \sum_j \binom{d}{j} \sum_k \binom{d}{k} k^2 - \left(\sum_j \binom{d}{j} j \right)^2. \end{aligned}$$

It is straightforward to show that

$$\sum_j \binom{d}{j} = 2^d, \sum_j \binom{d}{j} j = d \cdot 2^{d-1}, \text{ and } \sum_j \binom{d}{j} j^2 = d(d+1) \cdot 2^{d-2}.$$

We deduce that

$$c_d = 2^d \cdot d(d+1) \cdot 2^{d-2} - d^2 \cdot 2^{2d-2} = d \cdot 4^{d-1}.$$

□

5. Additional results

In this section, we present additional qualitative results on the three pre-trained models used in the paper: MobileNetV2 (Sandler et al., 2018), DenseNet121 (Huang et al., 2017), and InceptionV3 (Szegedy et al., 2016).

What does LIME really see in images?

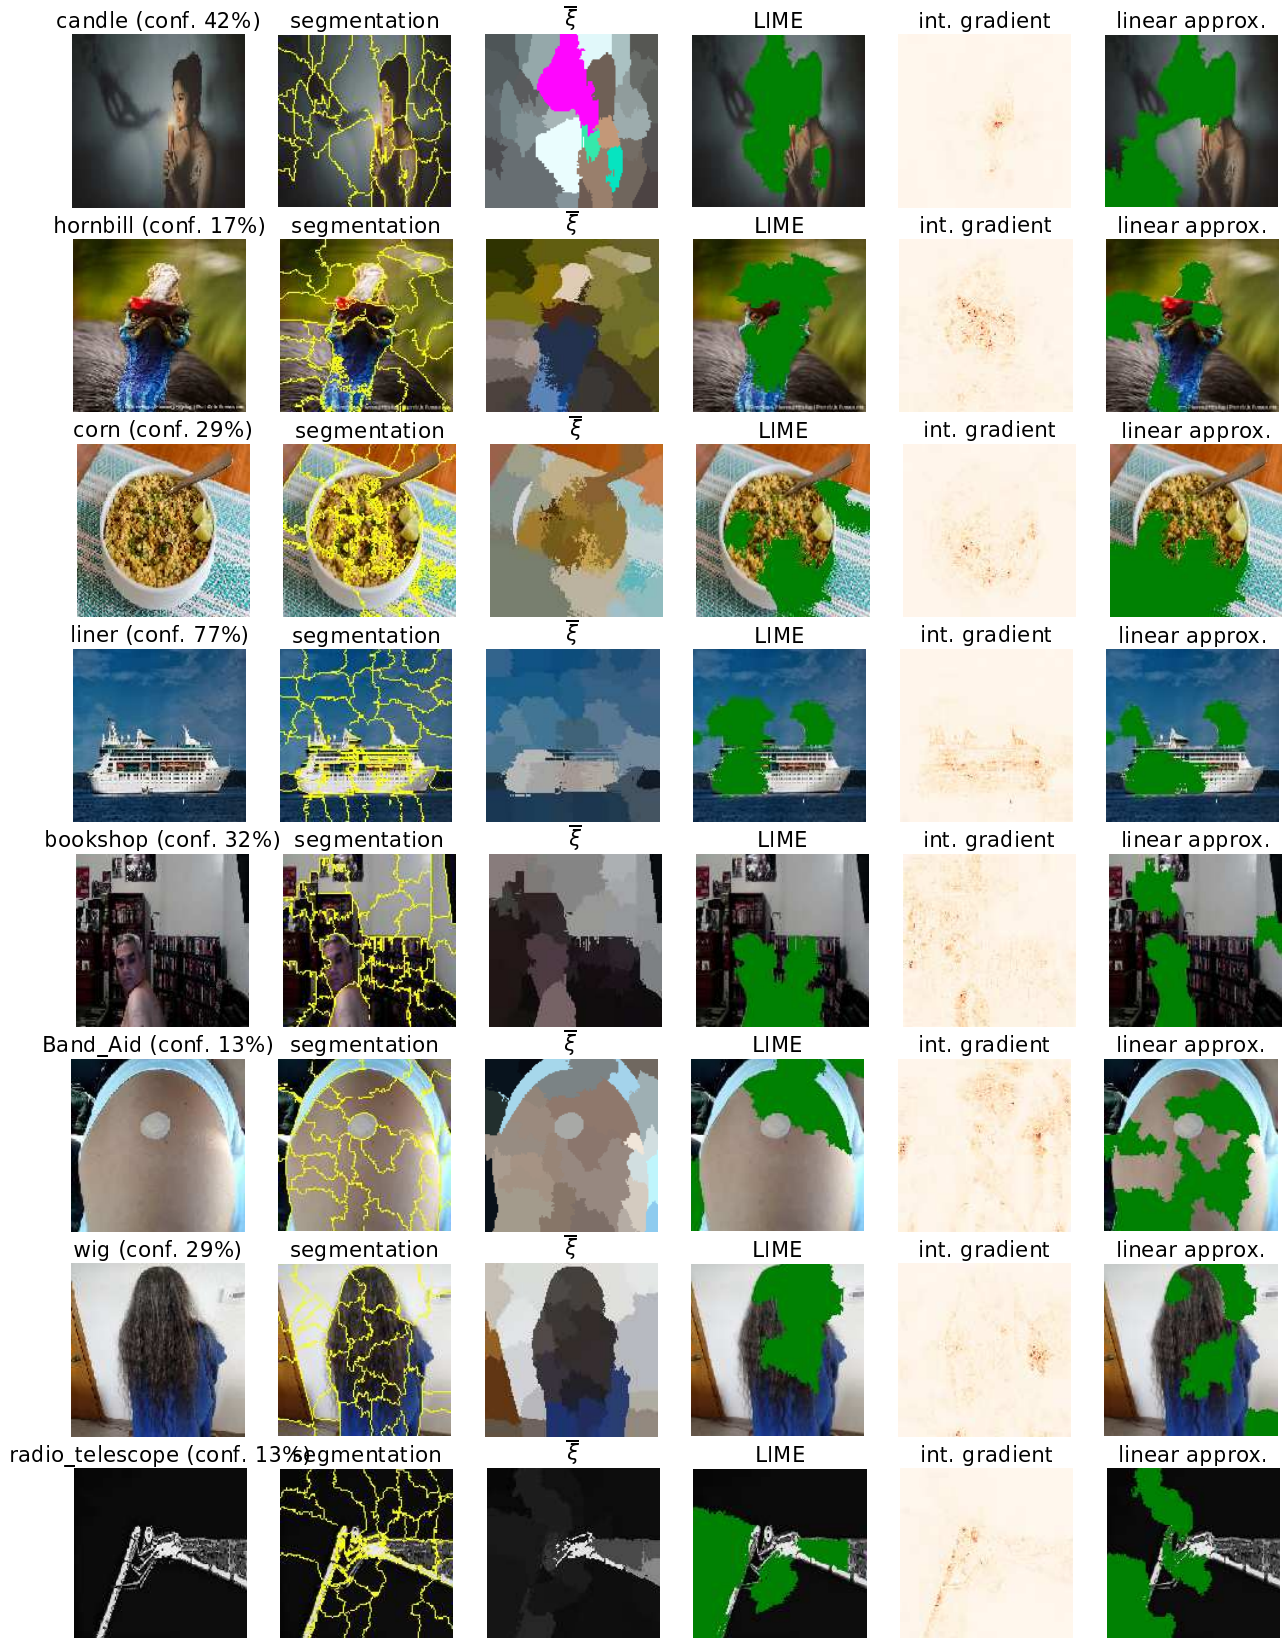


Figure 10: Empirical explanations, integrated gradient, and approximated explanations for images from the ILSVRC2017 dataset. The model explained is the likelihood function associated to the top class given by MobileNetV2.

What does LIME really see in images?




Figure 11: Empirical explanations, integrated gradient, and approximated explanations for images from the ILSVRC2017 dataset. The model explained is the likelihood function associated to the top class given by DenseNet121.

What does LIME really see in images?

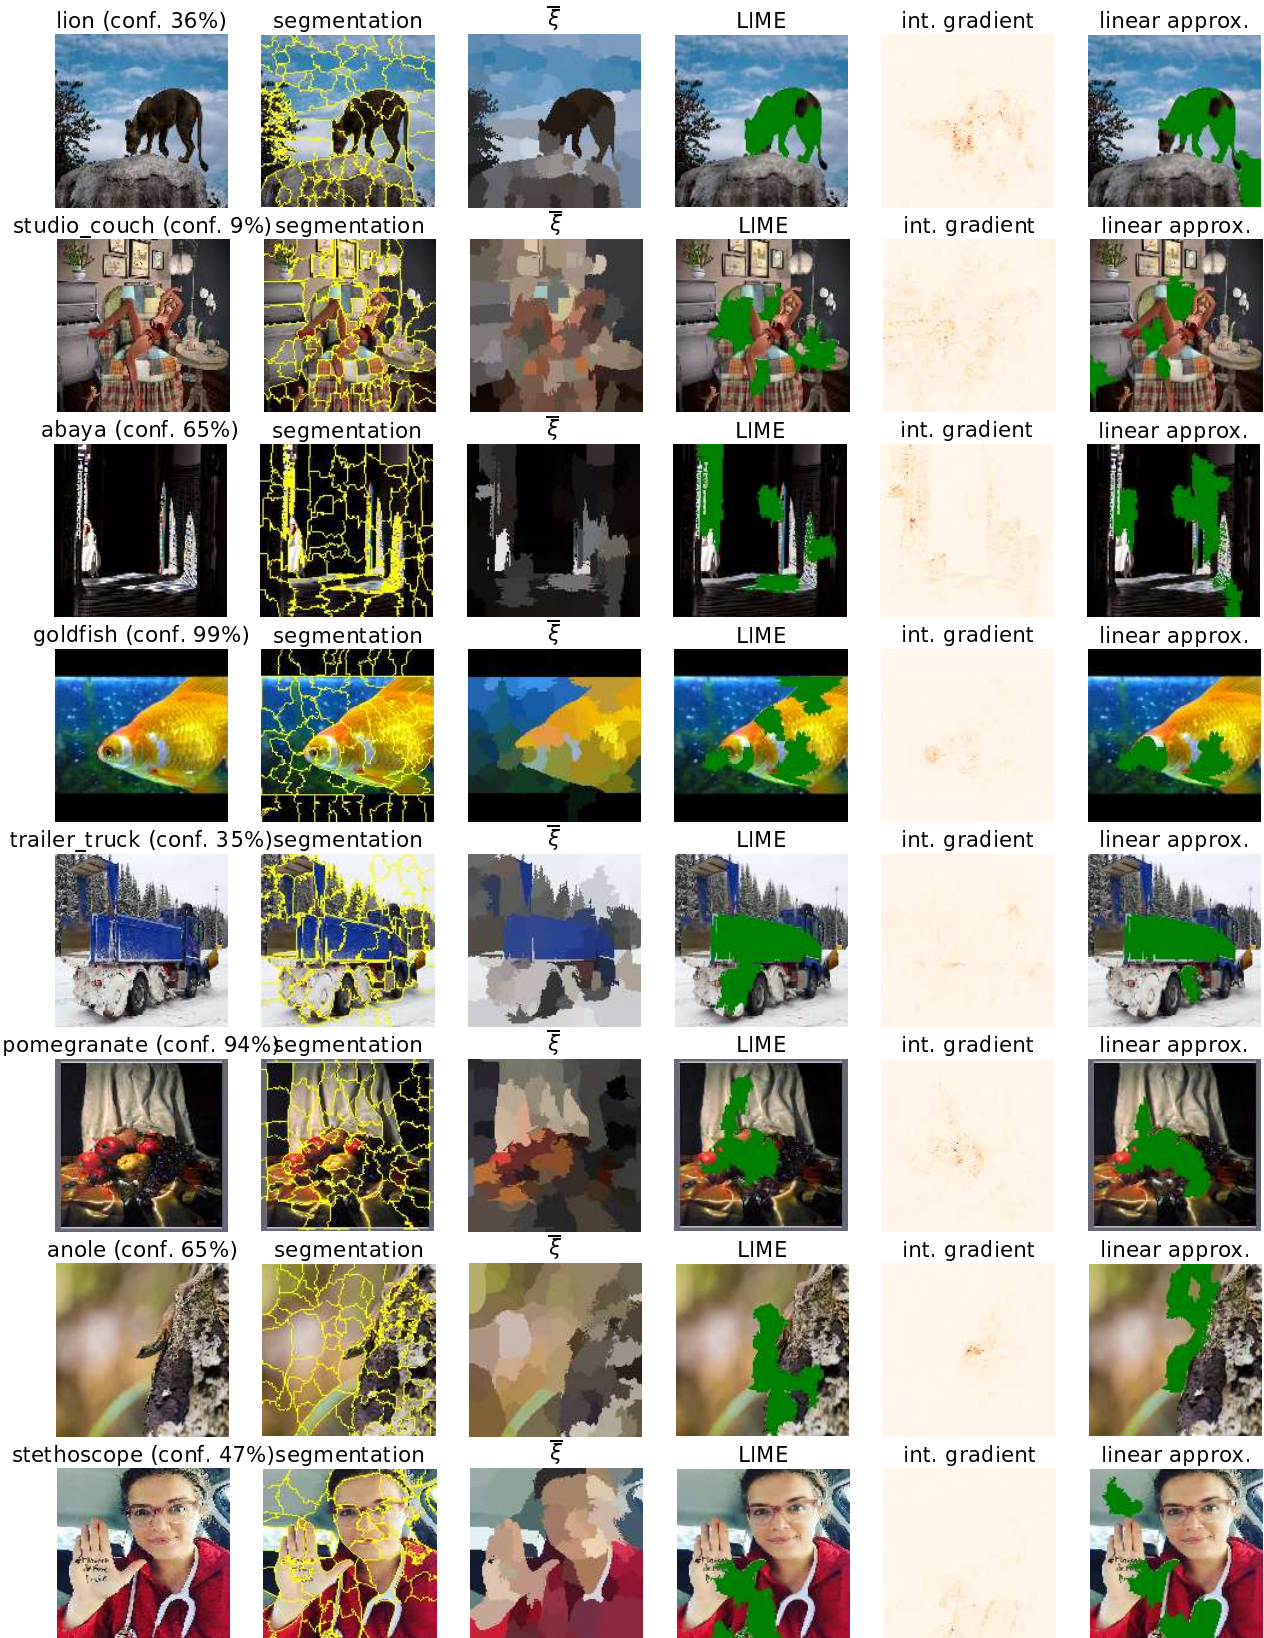


Figure 12: Empirical explanations, integrated gradient, and approximated explanations for images from the ILSVRC2017 dataset. The model explained is the likelihood function associated to the top class given by InceptionV3.