

Deep Prototypical-Parts Ease Morphological Kidney Stone Identification and are Competitively Robust to Photometric Perturbations

Daniel Flores-Araiza¹ Francisco Lopez-Tiro¹ Jonathan El-Beze² Jacques Hubert²
Miguel Gonzalez-Mendoza¹ Gilberto Ochoa-Ruiz¹ Christian Daul³

¹Tecnologico de Monterrey, School of Engineering and Sciences, Mexico

²CHU Nancy, Service d'urologie de Brabois, Nancy, France

³CRAN UMR 7039, Université de Lorraine and CNRS, Nancy, France

Abstract

Identifying the type of kidney stones can allow urologists to determine their cause of formation, improving the prescription of appropriate treatments to diminish future relapses. Currently, the associated ex-vivo diagnosis (known as Morpho-constitutional Analysis, MCA) is time-consuming, expensive and requires a great deal of experience, as it requires a visual analysis component that is highly operator dependant. Recently, machine learning methods have been developed for in-vivo endoscopic stone recognition. Deep Learning (DL) based methods outperform non-DL methods in terms of accuracy but lack explainability. Despite this trade-off, when it comes to making high-stakes decisions, it's important to prioritize understandable Computer-Aided Diagnosis (CADx) that suggests a course of action based on reasonable evidence, rather than a model prescribing a course of action. In this proposal, we learn Prototypical Parts (PPs) per kidney stone subtype, which are used by the DL model to generate an output classification. Using PPs in the classification task enables case-based reasoning explanations for such output, thus making the model interpretable. In addition, we modify global visual characteristics to describe their relevance to the PPs and the sensitivity of our model's performance. With this, we provide explanations with additional information at the sample, class and model levels in contrast to previous works. Although our implementation's average accuracy is lower than state-of-the-art (SOTA) non-interpretable DL models by 1.5%, our models perform 2.8% better on perturbed images with a lower standard deviation, without adversarial training. Thus, Learning PPs has the potential to create more robust DL models. Code at: https://github.com/DanielF29/Prototypical_Parts

1. Introduction

Urolithiasis is a disease in which kidney stones form somewhere in the urinary tract [26]. In developed countries, this disease has become a public health problem and presents a high incidence of kidney stones episodes, affecting up to 10% of their population [16]. The formation of kidney stones is strongly related to factors such as diet [12,23]. However, there are other factors that can lead to the production of a kidney stone such as age, chronic diseases, hereditary-family history, among others [4,6]. The timely identification of a kidney stone aids clinical specialists in determining the causes of its formation and to prescribe a personalized treatment, which can prevent relapses [12].

The morpho-constitutional analysis (MCA) is the standard technique for determining the different types of whole kidney stones removed during surgery (ex-vivo stones) [5,6]. Using this method it is possible to identify 21 different types of kidney stones, with both pure and mixed compositions [4]. Although the MCA efficiently establishes the type of ex-vivo kidney stones, it is very difficult to provide a reliable diagnosis during an endoscopic intervention (in-vivo), and the results of the MCA on extracted stones (ex-vivo) may take several days. Also, it can be difficult to carry out the inspection during surgery, as the extraction process can take up to one hour [4]. Furthermore, it is very difficult to train specialists on this technique, even considering a high incidence of kidney stone episodes.

Given the importance and difficulty of performing the visual inspection in a repeatable and highly reliable manner, recently, several Artificial Intelligence (AI) methods have been proposed for automating the kidney stone classification process. Using AI to automate the classification of kidney stones entails either replacing MCA or aiding the surgeon during the extraction process.

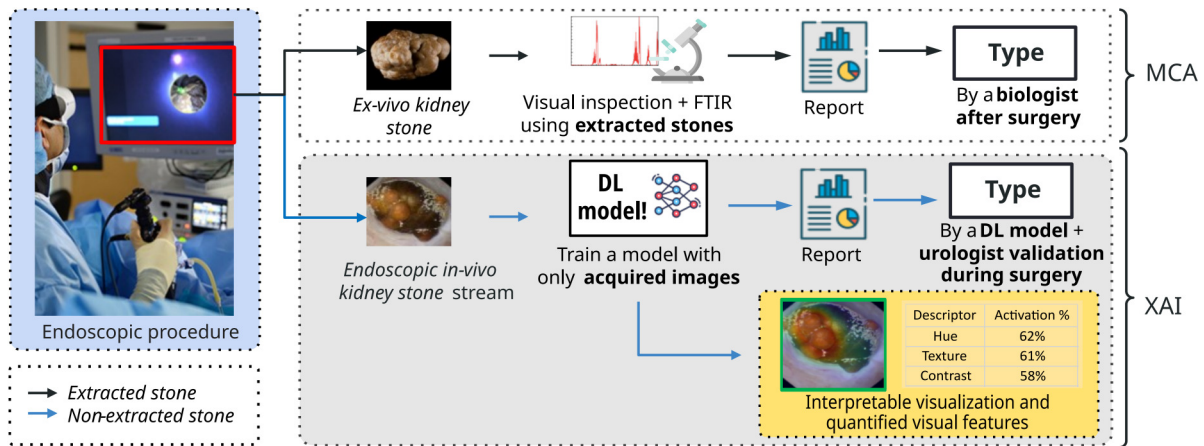


Figure 1. Traditional Morpho-constitutional Analysis (MCA) involves an endoscopic procedure to locate and extract a kidney stone (Ex-vivo), and subject it to visual and Fourier Transform Infrared Spectroscopy (FTIR) analysis by a specialist for classification, in the MCA section of the figure. Recent works have proposed using Deep Learning (DL) to classify kidney stones during an endoscopic procedure (In-vivo) automatically, represented in the XAI section of the figure. We argue and show, eXplainable AI (XAI) is needed for the adoption of DL as a Computer Assisted Diagnosis (CADx) tool. In the orange box is the extension to the current DL systems this proposal enables.

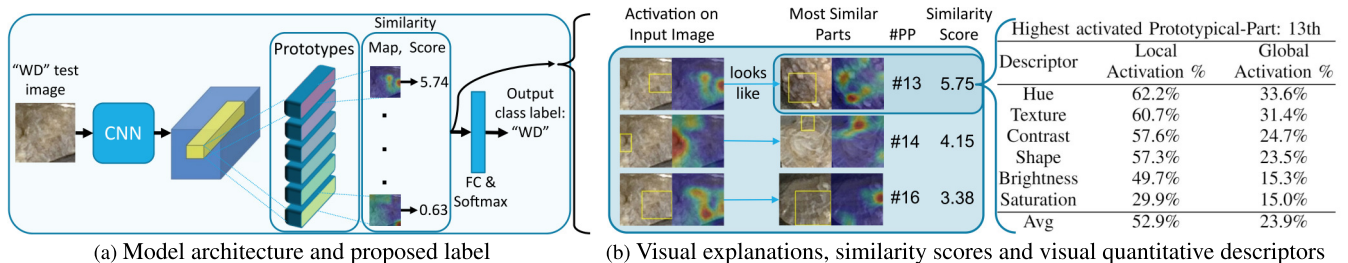


Figure 2. (a) The model classifies correctly a kidney stone patch, sub-type "WD", by measuring the similarity of learned Prototypical-Parts (PPs) to parts in the input image. As seen in (b), first, the most relevant area on the input image is indicated by a yellow bounding box and a heatmap representing levels of similarity on the input image for each Prototypical-Part (PP). Second, the image with the representative part from the training set that *looks like* the input indicated part is also shown on a yellow box, with its corresponding heatmap. Third, the level of saliency of visual characteristics and the descriptors are obtained. In this example descriptors for the most relevant PP are shown, indicating hue as the most relevant feature for the activation of this PP. Then, PPs identified work both as our explanations and for the classification, rendering this architecture interpretable, easing the urologist task of kidney stone classification, as they are provided with the reasons used by the model for the suggested classification.

Several DL-based techniques in particular have demonstrated that it is possible to replicate the results obtained by well-trained specialists [2, 10, 13, 18, 19]. However, DL models lack interpretation of the extracted features, making these models of limited help in clinical settings. To facilitate clinical acceptance, a DL model will need to provide the specialist with its reasoning process, this implies providing evidence on the *what*, *how* and *why* behind any suggested classification [7, 24]. Therefore, to be a useful aid for CADx purposes, any evidence should reflect the same key distinctions used for MCA by a specialist.

In order to pave the way for AI-based interpretable MCA using deep learning techniques, we leverage recent strides in explainability that seek to base image classification on case-based reasoning [1, 3, 20]. In this work, both visual ex-

planations and quantitative information about visual characteristics deemed important by the network are provided [17]. It must be emphasized that the explanations provided by our approach follow the reasoning processes of urologists during MCA by showing the detected morphological relevant features of each image as learned PPs.

Overall, this work presents an interpretable DL classifier for MCA. This proposal facilitates human-machine collaboration for morphological analysis whose explicit reasoning can be easily inspected, understood and verified by urologists and medical specialists. In contrast to previous works, we evaluate our interpretable models and their DL counterparts under perturbations of global visual characteristics to describe their relevance to the PPs and the sensitivity of our model.

The of the paper is organized as follows. In the following Section 2 the proposed approach reasoning and justification are presented. In Section 4 the experimental setup is described for our experiments. In Section 5 we share our results and present a discussion on their implications, sharing one explanation case comparing a correct and incorrect classification. Finally, in Section 6 the conclusion and the insights obtained for further research in the field are presented.

2. Motivation

MCA consists of two complementary analyses of the kidney stone ex-vivo in the laboratory. First, a specialist performs a visual inspection of the characteristics (such as color, texture and shapes) of the surface (external) and section (cross-sectional) views. Second, a biochemical study is performed to determine the main components in the stone through an Fourier Transform Infrared Spectroscopy (FTIR) study [9]. Both studies complement each other to provide a detailed report to determine the type of kidney stone. This traditional MCA process is represented in the MCA area in Fig. 1.

Recently proposed AI systems for kidney stone classification are presented as automated CADx tools. This approach could enable faster and more reliable diagnosis based solely on kidney stone images, reducing the need for invasive procedures, on in-vivo [10] or ex-vivo [2, 13, 18] conditions, as represented in *MCA-Based DL*, the grey area at Fig. 1. Also, for in-vivo scenarios, the operations would be quicker to perform and less traumatic, due to the classification of the kidney stones right before performing dusting, a process intended to fragment and destroy the kidney stones inside the urinary tract. Current DL classifiers of kidney stones are black boxes, as those have the limitation of being unable to describe the inner workings that led to a given prediction beyond the class label. Therefore, these types of models cannot provide useful information to the specialist to understand how features on the input image were used to output the classification proposed by the DL model.

The eXplainable AI (XAI) field seeks to remedy these limitations. XAI methods are used to provide an explanation of how DL models arrived at a particular output, typically after the classification was assigned. If DL models achieve high performance and can provide explanations for the classification based on relevant features for the specialists, then it would be easier to maintain accountability for recommended treatments based on those analyses. In the area indicated by the XAI bracket at Fig. 1, a holistic explanation is represented, by providing the description of what and why visual features are relevant from the input image, as quantitative evaluations of visual features, and a saliency map highlighting where the described features are located.

Explanations that could be considered complete do not yet exist for kidney stone identification. Nonetheless, improving explanations by indicating precise visual characteristics would facilitate the adoption of DL CADx tools, resulting in faster, more accountable, and repeatable diagnosis recommendations for the treatment of urolithiasis.

Taking into account the previously mentioned desired characteristics of explanations, in this work, we seek to improve the current methods by adding the required explainability to facilitate the morphological analysis of kidney stones, our proposal of how to achieve this is depicted in Fig. 2. For context, Fig. 3 presents a visual comparison of different explanations. Then, to improve DL-generated explanations for the classification of kidney stones we began generating saliency maps of learned representative parts, the learned PPs, for each class [11], after, we evaluate which global visual characteristics PPs are relating for such classification [17].

3. Proposed Approach

Having a single heatmap as a visual explanation, as most current XAI visualization methods do, is an oversimplification of all the characteristics expected and used to classify a kidney stone into its corresponding type. To alleviate the unmet needs of most of the current DL models and XAI visualization methods, this proposal leverages a case-based reasoning process. Our proposal pipeline does this by extracting semantic features, representing parts, from an input image with a Convolutional Neural Network (CNN). Then, the extracted features are measured against the learned prototypical parts, in the form of similarity scores, as indicated by equation 1, and represented in Fig. 2b. Finally, the output prediction of the classification task is a weighted combination of the similarity scores. Using these same learned PPs for the output of the model and the explanations, the method guarantees faithfulness of the explanations to the inner workings of the model. Current case-based reasoning methods, instead of using high-dimensional features from the convolutional layers of a CNN, use the identification of a few prototypical parts [1, 3, 20, 21]. This identification of a limited number of prototypical parts enables user understanding of the learned reasoning behind the DL model output, which we show, achieves competitive performance against its non-interpretable counterparts. The advantage of the case-based reasoning process is that the explanations generated follow the same reasoning process a specialist used to classify a kidney stone image.

However, for images of kidney stones, a great level of expertise is required, and thus, explanations generated from PPs for kidney stone patches still depend on non-obvious characteristics from the input image for non-experts, reason for us to quantify the sensitivity of PPs to a set of perturbations [17], which we dub herein as “Descriptors”.

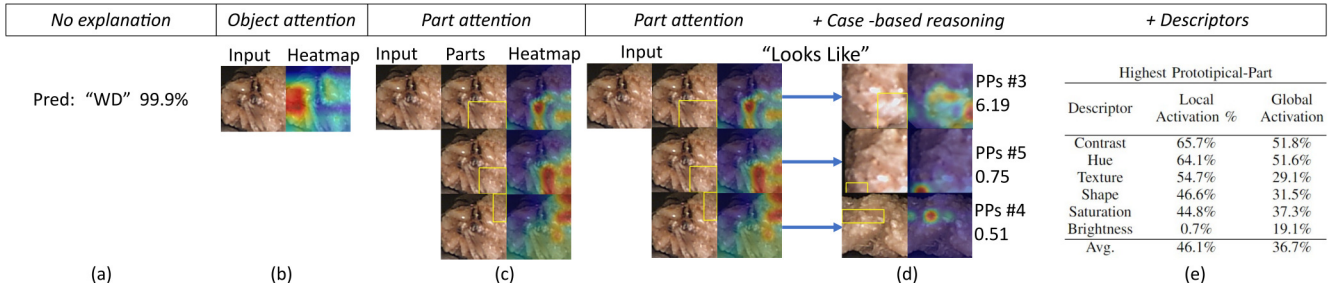


Figure 3. Visual comparison of different types of model interpretability on an example of a kidney stone image as an example. (a) Traditional non-interpretable deep learning output classification of a kidney stone [13], (b) object-level attention map (e.g., class activation map [22]), (c) part attention (provided by attention-based interpretable models [28]); and (d) part attention with similar prototypical parts [3], with additional (e) quantitative evaluations of visual features, the *descriptors* [11, 17].

Each perturbation is directed at a particular visual characteristic (Brightness, Saturation, Shape, Texture, Contrast, and Hue). These descriptors allow measuring individually the level of importance of visual features to each of the PPs, as shown in the table in Fig. 2b.

Our model shows the underlying decision-making process for each case, based on several explanations which correspond to the most similar cases identified and used by the model to give its suggested classification of the image. Fig. 2 shows an example of how our proposal works: **I)** Our model processes an input image through a feature extractor backbone (CNN in Fig. 2a). **II)** Similarity of the learned PPs to the extracted features from the CNN are calculated accordingly to equation 1 (Similarity score in Fig. 2a). **III)** From the Prototypes layer, in Fig. 2a, the scores of PPs calculated against the input image are processed through a linear weighted combination, a fully connected layer, to provide a final classification of a kidney stone image. **IV)** The main PPs that determined the classification are provided as the most similar cases to their respective parts from the input image, as seen in the blue box in Fig. 2b. **V)** Finally, the saliency of certain characteristics for each prototypical part is evaluated, the *Descriptors*, listed in the table in Fig. 2b and provide a description of the input image features that mainly triggered the activation of the corresponding learned Prototypical-Part (PP).

This capacity yields predictions that are easy to understand by specialists familiarized with the features related to each kidney stone type and subtype, rendering classifications interpretable for urologists. Thus, aiming to help in the morphological analysis of each kidney stone image, the trained model provides the specialist with the means to trust the AI, check its output for plausibility, and easily overrule it when necessary. Our approach, with its inherently interpretable reasoning process, contrasts directly with previous work that relied on post-hoc explanation techniques to explain a trained black-box model on particular classifications [10, 18] or with global explanations [8].

Table 1. Number of complete endoscopic images acquired.

Subtype	Main component	Key	Surface	Section	Mixed
Ia	Whewellite	WW	62	25	87
IIa	Weddellite	WD	13	12	25
IIIa	Acide Urique	AU	58	50	108
IVc	Struvite	STR	43	24	67
IVd	Brushite	BRU	23	4	27
Va	Cystine	CYS	47	48	95
TOTAL			246	163	409

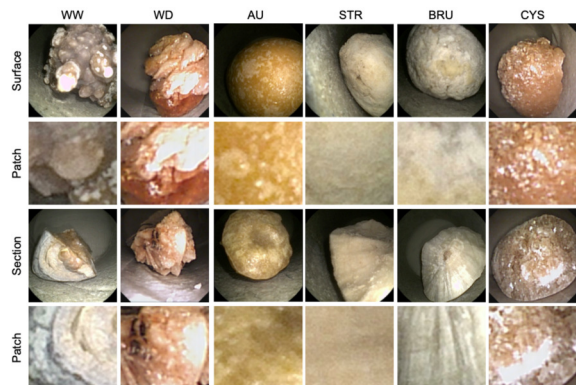


Figure 4. Examples of kidney stones patches extracted from the dataset used in [8]. This endoscopic dataset consists of six of the most common kidney stones namely Whewellite (WW), Weddellite (WD), Acide Urique (AU), Struvite (STR), Brushite (BRU), and Cystine (CYS).

Following the architecture proposed on [3, 20], in this study, we explore ProtoPNet architecture with some important differences, **I)** the analysis performed explores different backbones, with the intention to compare CNN with different levels of intricate connections. **II)** Different number of prototypical parts per class and **III)** The relevance of data augmentation for training. A relevant note is that the training of the models was performed without part annotations, only relying on the class label.

3.1. Kidney stone dataset

A simulated in-vivo dataset of endoscopic images was used for the experiments reported in this contribution (Table 1). This dataset consists of surface (246 images) and section (163 images) views. Each view contains six of the most common kidney stone sub-types: Ia (Whewellite, WW), IIa (Weddellite, WD), IIIa (Uric Acid anhydrous, UA), IVc (Struvite, STR), IVd (Brushite, BRU), and Va (Cystine, CYS). Two different reusable digital flexible ureteroscopes Karl Storz (Storz Image 1 Hub and Storz image1 S) captured images of real kidney stones in an environment that simulates real conditions (such as illumination or blur) during surgery. In addition, the acquired endoscopic images were labeled manually by the urologist Jonathan El-Beze (see [8] for more details).

In order to train DL-based models, 12,000 square-patches of 256×256 pixels were extracted from 409 endoscopic images of both views. This combined data is referred to as *mixed* dataset. Patches of the dataset are shown in Fig. 4. For more details, patches were extracted with the method described in [18] and [13].

3.2. Model architecture:

The architecture implemented uses a feature extractor $f(x)$, in this proposal a CNN, to perform the extraction of semantic features from the input images x and learn n prototypes per class, corresponding to parts in the input image. This first step of feature extraction from the input image implies the potential diversity and quality of possible prototypical parts to be learned from. Hence, three CNN backbones are explored (VGG16, ResNet50 and DenseNet201) in order to compare the different possible performances obtained with our approach and compared against pure non-interpretable CNNs present in the estate of the art for kidney stone identification.

Next, two layers of 1×1 convolutions are added, to adjust the depth of the feature activation maps to the depth selected of 128 for the layer of PPs, layer g . The variable l is used to index each of the $128 \times 1 \times 1$ patches $f(x)_l$ across the spatial dimensions. The layer of PPs g , contains n prototypes to be learned per class, with dimensions $128 \times 1 \times 1$.

Since a prototype has the same number of channels but a smaller spatial dimension than the convolutional feature maps $f(x)$, we can interpret the prototype as representing a prototypical activation pattern of its class and we can visualize the prototype as a patch of the training image it appears in. The distance between p_j and each of the $l \times 1$ spatial patches of the convolutional feature map $f(x)$ is measured by $d_{j,l} = \|p_j - f(x)_l\|_2^2$, and converts these distances to similarities $s_{j,l}$ by:

$$s_{j,l} = \log \frac{d_{j,l} + 1}{d_{j,l} + \epsilon} \quad (1)$$

Where ϵ is a small value to avoid division by zero. This similarity scores $s_{j,l}$ can be arranged spatially into a similarity map, upsampled and superimposed on the input image to generate a saliency map as the ones presented on the blue box in Fig. 2b. The overall similarity score of the corresponding PPs is the highest score s_j obtained from $s_{j,l}$. These similarity scores are then processed with a fully connected layer, as to combine them in a weighted combination to obtain an output per class, which is then normalized with a softmax function. The details for the initialization and training hyperparameters, except for the data augmentation used, follow [3].

3.3. Descriptors calculation:

Importance scores for image characteristics of the PPs are analyzed by measurement of the change in the similarity score s_j when perturbing the input image with certain modifications \hat{s}_j . Let $i \in \{\text{brightness, contrast, saturation, hue, shape, texture}\}$ denote the type of modification, as further detailed in [17]. Then the local importance score $\Phi_{local}^{i,j,k}$ of characteristic i for test image $j \in S_{test}$ with class k on the j -th prototype is the difference in similarity score:

$$\Phi_{local}^{i,j,k} = s_j - \hat{s}_j \quad (2)$$

And its global importance is a weighted arithmetic mean by weighing the local scores of all images in S_{train} by their respective original similarity score per PP:

$$\Phi_{global}^{i,j} = \frac{\sum_{k=1}^{|S_{train}|} \Phi_{local}^{i,j,k} \cdot s_j}{\sum_{k=1}^{|S_{train}|} s_j} \quad (3)$$

Hence, if unmodified image k gets a low similarity score with prototype j , it will provide a low weight during the global importance calculation. In contrast, if the prototype j is clearly present in image k , the architecture will assign a high similarity score and hence the local importance score for image k gets a high weight. As demonstrated in equation 4, these importance scores can be used to create global explanations that explain a prototype, and local explanations, equation 2, that explain the main visual features responsible for a high similarity score between a given image and a prototype. As appreciated from equation 2 and equation 4, with this *descriptors* the intention is to explain the identification of PPs on the input image from the models perspective, not from human perception, nonetheless, since the *descriptors* explored are humanly interpretable actually those do enrich the explanations obtained. The specifics for the different perturbations used to obtain the values of the PPs *descriptors* follow [17].

3.4. Evaluation metrics

To determine the different types of kidney stones in endoscopic images, MCA specialists perform a visual inspec-

tion of the surface (external part) and section (internal part) view, and provide a detailed report of each view of the kidney stone. In addition, in order to improve performance in this task, MCA also proposes combining information from surface and section views, since both views provide relevant information for classification. [4]

In this work, we perform different training of mixed views models (similar to how MCA specialists perform it). The accuracy and precision metrics are mainly used to quantitatively evaluate different implementations of DL models proposed in the SOTA mentioned to be intended as a CADx tool [2, 10, 13, 18, 25]. Therefore, our particular implementations were measured and compared on accuracy as well. It can be observed in Table 2 the comparison of the accuracy metric of recent works in the SOTA for the task of kidney stone identification against our implementations.

4. Experimental Setup

Three different CNN architectures are explored for our proposal: i) a VGG16 to examine the performance of a simple deep CNN, ii) a ResNet50 to consider a medium size CNN with residual connections and iii) a DenseNet201 for the evaluation of dense connections of a deep model to the SOTA for kidney stone identification. These architectures were pre-trained on ImageNet, to then train them with the kidney stone dataset. This enables having these CNNs as baselines, and ResNet 50 and VGG16 are also used on other implementations in the current SOTA for kidney stone identification [2, 10, 13–15, 27].

For the prototype layer g , the number of PPs per class explored are (1, 3, 10) PPs per each of the six classes of the three configurations of our dataset, described in section 3.1. The training of our implemented ProtoPNet models is a cycle consisting of three stages: (A_1) Updating the weights of the convolutional layers used as a feature extractor $f(x)$ per 10 epochs, using the loss function presented in [3]. (A_2) Projection of the prototype layer g to their corresponding closest parts after the convolutional layers $f(x)$ updates. Detailed description in supplementary material section S1.2 (A_3) Update of the fully connected layer h_1 weights for 20 epochs (FC & Softmax layer in Fig. 2a). We repeated this cycle three times and selected the best model on the mixed test data obtained during one of the (A_3) phases.

Preparation for the train and test sets were 80% (1,600 images per class) and 20% (400 images per class), respectively. The patches were also “whitened” using the mean m_i and standard deviation σ_i of the color values I_i in each channel ($I_i^w = (I_i - m_i)/\sigma_i$), with $i = R, G, B$). The different mentioned configurations for training were performed on the mixed data, described in Section 3.1, without data augmentation and with data augmentation. The data augmentation used consists in randomly applying one data augmentation at a time to an input image, from a set of visual

Table 2. Accuracy of state-of-the-art models and three of ours for identifying kidney stones. Models trained with data augmentation.

Method	Surface	Section	Mixed
Martinez, et al. [15]	56.2±23.3	46.6±12.8	52.7±18.9
Estrade, et al. [10]	73.7±17.9	78.8±10.6	70.1±22.3
Black, et al. [2]	73.5±19.0	88.8±02.8	80.1±13.8
Lopez, et al. [13]	81.0±03.0	88.0±02.3	85.0±03.0
Lopez, et al. [14]	83.2±01.2	90.4±04.8	85.6±00.1
Villalvazo-Avila, et al. [27]	88.8±02.8	84.4±06.0	96.6±00.5
Densenet201 3pps	80.2±00.6	87.6±06.2	85.2±01.5
Ours Resnet50 1pps	80.9±01.3	91.0±03.3	88.2±02.1
Vgg16 10pps	76.7±02.2	82.6±04.9	82.1±00.9

transformations. This set of transformations consists of horizontal flip, vertical flip, rotation between -180 and 180 degrees, a distortion scale of 0.4, translation up to 0.2, or symmetric padding of 50 pixels. After selecting the transformation it has a 50% chance of been applied.

After training the different model’s configurations, evaluation of the perturbations, mentioned in Section 3.3, is carried out as PPs *descriptors* and, additionally, we evaluate the performance of the models under these perturbations (graph in the supplementary material as Fig. S1). To preserve clarity of results, supplementary Figures S1 to S3, Fig. 6 and Tables 2 and 3, only show the results on the mixed test data. Additionally, the performance of the trained models under different levels of change on the most impactful perturbation, the hue channel perturbation, is evaluated. The original change in hue explored in [17] applies a color jitter modification with the torchvision library to the input images, fixed on a change of 0.1. We explore a range of values for the change in the hue channel to characterize the sensitivity of the performance of our models. Fig. S2 illustrates the performance of one model training configuration under different hue channel settings.

5. Results and Discussion

5.1. Implementation details

Experiments were carried out to explore different backbones for the feature extractor network $f(x)$, the numbers of PPs per class and the significance of data augmentation. This, is to evaluate the model’s performance and the obtained explanations. In contrast to CNNs, which are black-box classifiers, our proposal provides explanations for input images in three different ways: **i)** by showing the activation area, where each PP is identified. **ii)** The corresponding representative example image to each activated PPs. **iii)** Quantitative evaluations of how relevant six different visual descriptors are for the activation of each PP.

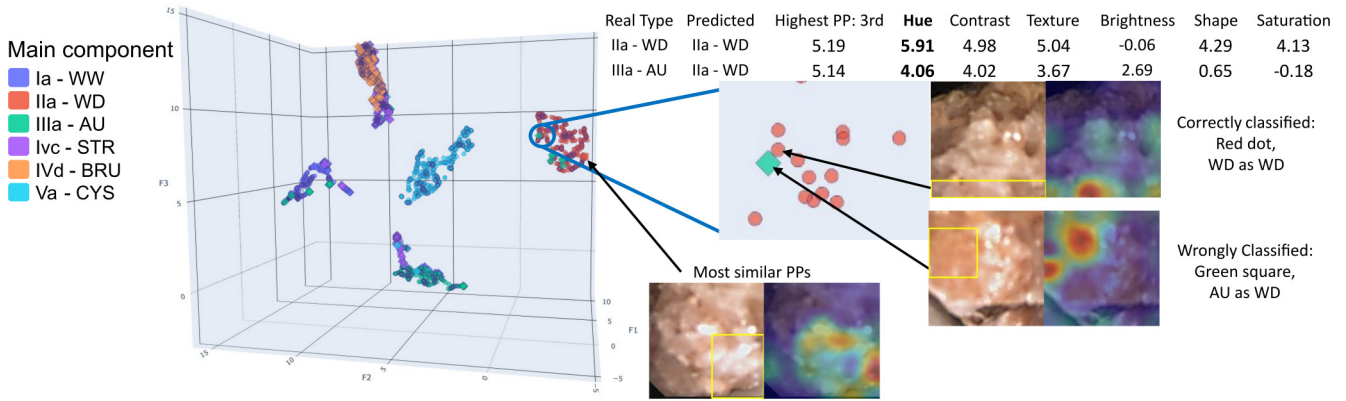


Figure 5. UMAP of the Part-Prototypes and descriptors activations on the test patches of kidney stone images, for the particular case of a ProtoPNet using DenseNet201 as the backbone and 3 Prototypical parts per class. Our approach allows obtaining separate clusters of the output classes.

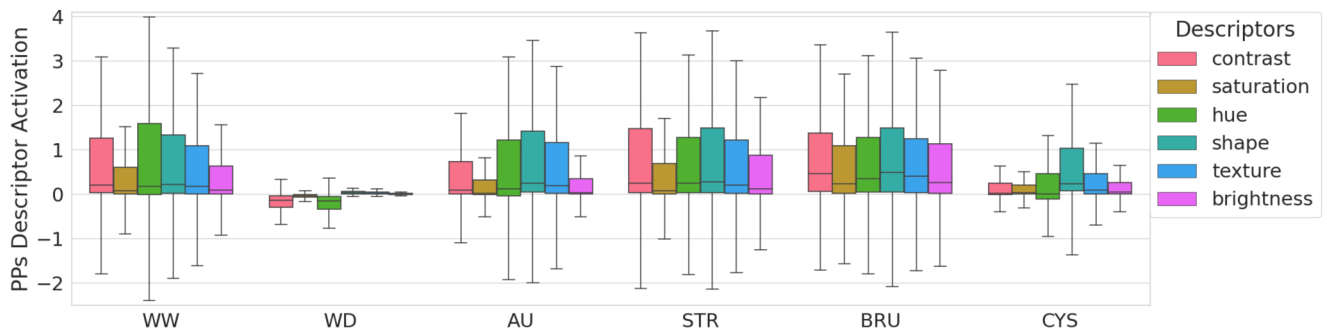


Figure 6. The six descriptor activations of the Prototypical Parts (PPs) per class, for the particular case of a ProtoPNet using DenseNet201 as the backbone and 3 Prototypical parts per class.

5.2. Visualization of a case example of explanations

In Fig. 5, an example of our model classifications is shown. We found that averaging the global scores of the PPs descriptors per class results in a hierarchy of descriptor significance, as depicted in Fig. 6. The WD class reveals a unique pattern of descriptor activation. An examination of several instances shows that the average activation of a class descriptors can hint at misclassifications. As shown in Fig. 5, a green diamond represents a misclassified AU image, which is misidentified as WD due to its activations corresponding to those of the WD class. This causes the image to appear within the WD class cluster on the UMAP plot. We use a UMAP visualization to present PPs activations for input images in the context of the top three discriminative dimensions. This UMAP highlights class separability for each output class of the ProtoPNet, as showcased in Fig. 5. The use of descriptors mitigates the cases for visually similar PPs by providing details on the characteristics most relevant for each PP [17]. Additionally, learning PPs from patches of kidney stone images allows observing PPs are learned proportionally to the architecture of the CNN used.

5.3. Performance results

A summary of the accuracy of the different configurations of models trained is displayed in Table 3. Tests carried out on the original mixed test dataset, corresponding to the training mixed dataset, are considered Independent Identically Distributed (IID) data, tests performed under perturbations applied to the original test dataset are considered Out Of Distribution (OOD) data and named accordingly in Table 3. Usage of data augmentation yielded an average 3% improvement accuracy compared to the non-augmented approach consistently across different evaluations. To ease presentation, we report results just on models trained with data augmentation.

The performance of our implemented ProtoPNet models is comparable with its corresponding non-interpretable CNN backbone models as also appreciated in Table 3. The average accuracy of our implementations keeps competitiveness with a difference not bigger than $\leq 0.18\%$, when compared with the average accuracy of the baseline CNN models. The mean values and standard deviations displayed in Table 3 are computed based on the configurations of the

Table 3. Performance metrics of our models and convolutional neural networks used as backbones. Evaluations are divided between tests performed with the same type of data the model was trained, considered Independent Identically Distributed (IID) data and tests performed under perturbations, which are considered Out Of Distribution (OOD) data. All models were pre-trained on ImageNet.

Model	IID			OOD		
	Accuracy (%)	Precision (%)	F1 (%)	Accuracy (%)	Precision (%)	F1 (%)
DenseNet201	89.67±3.60	90.51±3.60	89.47±3.46	58.44±28.06	60.22±29.62	54.54±32.73
ResNet50	83.40±5.32	85.58±4.30	83.11±5.21	45.12±20.82	49.16±24.08	38.25±23.47
Vgg16	81.47±2.72	84.71±2.64	80.62±3.84	58.35±22.62	63.73±19.93	54.38±26.30
CNNs Average	84.84±5.29	86.93±4.35	84.40±5.60	53.97±24.62	57.70±25.40	49.06±28.56
DenseNet201	1 PPs	86.29±1.91	87.54±1.34	86.25±1.78	60.31±20.32	70.29±14.38
	3 PPs	85.19±1.50	86.08±1.27	85.21±1.42	59.97±19.91	69.21±12.59
	10 PPs	87.29±0.92	87.99±0.94	87.23±0.92	61.20±19.33	71.26±11.92
ResNet50	1 PPs	88.21±2.07	88.61±1.88	88.21±2.05	58.26±24.23	55.16±27.50
	3 PPs	86.66±1.37	87.11±1.55	86.62±1.45	57.90±21.82	63.24±19.16
	10 PPs	85.44±1.44	86.16±1.15	85.43±1.42	56.31±22.66	64.81±16.55
Vgg16	1 PPs	81.78±1.60	83.38±1.94	81.82±1.62	51.78±19.02	53.82±19.40
	3 PPs	82.21±3.33	82.92±3.82	81.84±3.75	51.23±19.40	54.09±19.80
	10 PPs	82.08±0.90	83.39±0.92	82.13±0.87	54.69±18.05	56.48±17.56
ProtoPNets Average	85.02±2.83	85.91±2.66	84.97±2.89	56.85±20.63	63.30±17.90	53.30±23.83

models outlined therein. It is worth noticing that Table 3 shows our interpretable models to outperform traditional CNNs from the SOTA on kidney stone classification when evaluated on perturbed input images. Specifically, our models achieve 1.9% higher accuracy, 5.7% higher precision, and 3.2% higher F1 scores compared to traditional models.

6. Conclusions and Future Directions

We showed that by adapting and fine tuning CNN models into ProtoPNets is possible to convert considered black-box CNN models into interpretable ones, bringing detailed and faithful explanations. The main visual features in the input image that produce the activation of PPs in particular and in the average cases per class and per dataset are able to be extracted from the evaluated behavior of the network, as shown in Fig. 6. This case-based reasoning approach allows to generate explanations and additional activation details enabling urologists to use these models as assistance tools for the MCA. More importantly, we showed evidence for how to use the overall behavior of descriptors per class in combination with UMAP projection of the behavior of the model to identify the reasons behind misclassification, as shown in Fig. 5. Also, the UMAP visualization illustrates the blending of certain classes that share similar color and texture characteristics. To enhance cluster separability a metric learning approach could be explored. This to penalizes close PPs from different classes, encourages proximity for PPs of the same class but avoids their collapse.

Upon examining Table 3, on one hand, it is evident that despite a minor decrease in average performance for IID evaluations in PPs models compared to their corresponding CNNs. This could be due to restricting the classification to a few PPs per class, in comparison CNN models use a few hundred features. On the other hand, PPs models consistently exhibit reduced variance in performance metrics.

Additionally, these models consistently perform better in OOD evaluations, suggesting that PPs models tend to be more robust than their CNN counterparts. This characteristic could be attributed to a regularization effect resulting from measuring the similarity between the learned PPs and the image features. PPs models inherit most of the performance characteristics of CNNs, indicating that methods to enhance CNN performance should also improve PPs models, thus expanding potential improvements.

In future research, the model’s reasoning and reporting mechanisms could be refined to align more closely with the structured lexicon employed by urologists. This could be achieved by incorporating natural language descriptions for the PPs identified and their *descriptors* to make explanations more accessible to specialists and a broader range of users.

Compliance with Ethical Standards

The images were captured in medical procedures following the ethical principles outlined in the Helsinki Declaration of 1975, as revised in 2000, with the consent of the patients.

Acknowledgments

The authors extend their gratitude to the AI Hub and the Centro de Innovación de Internet de las Cosas at Tecnológico de Monterrey for generously providing access to an NVIDIA DGX computer, which facilitated the experiments conducted in this paper. Additionally, we would like to express our gratitude to CONACyT for the Ph.D. scholarships for Daniel Flores-Araiza and Francisco López-Tiro.

References

- [1] Alina Barnett, Fides Regina Schwartz, Chaofan Tao, Chaofan Chen, Yin hao Ren, Joseph Y. Lo, and Cynthia Rudin. Iaia-bl: A case-based interpretable deep learning model for classification of mass lesions in digital mammography. *Nat. Mach. Intell.*, 3:1061–1070, 2021. [2](#), [3](#)
- [2] Kristian M Black, Hei Law, Ali Aldoukhi, Jia Deng, and Khurshid R Ghani. Deep learning computer vision algorithm for detecting kidney stone composition. 2020. [2](#), [3](#), [6](#)
- [3] Chaofan Chen, Oscar Li, Chaofan Tao, Alina Jade Barnett, Jonathan Su, and Cynthia Rudin. This Looks Like That: Deep Learning for Interpretable Image Recognition. *CoRR*, jun 2018. [2](#), [3](#), [4](#), [5](#), [6](#)
- [4] Mariela Corrales, Steeve Doizi, Yazeed Barghouthy, Olivier Traxer, and Michel Daudon. Classification of stones according to michel daudon: a narrative review. *European Urology Focus*, 7(1):13–21, 2021. [1](#), [6](#)
- [5] Michel Daudon, Arnaud Dessombz, Vincent Frochot, Emmanuel Letavernier, Jean-Philippe Haymann, Paul Jungers, and Dominique Bazin. Comprehensive morpho-constitutional analysis of urinary stones improves etiological diagnosis and therapeutic strategy of nephrolithiasis. *Comptes Rendus Chimie*, 19(11-12):1470–1491, 2016. [1](#)
- [6] Michel Daudon, Paul Jungers, Dominique Bazin, and James C Williams. Recurrence rates of urinary calculi according to stone composition and morphology. *Urolithiasis*, 46(5):459–470, 2018. [1](#)
- [7] Brian Edwards. Fda guidance on clinical decision support: Peering inside the black box of algorithmic intelligence, Nov 2019. [2](#)
- [8] Jonathan El Beze, Charles Mazeaud, Christian Daul, Gilberto Ochoa-Ruiz, Michel Daudon, Pascal Eschwège, and Jacques Hubert. Evaluation and understanding of automated urolithiasis recognition methods. *BJU International*, 2022. [4](#), [5](#)
- [9] Laurence Estepa and Michel Daudon. Contribution of fourier transform infrared spectroscopy to the identification of urinary stones and kidney crystal deposits. *Biospectroscopy*, 3(5):347–369, 1997. [3](#)
- [10] Vincent Estrade, Michel Daudon, Emmanuel Richard, Jean-Christophe Bernhard, Franck Bladou, Gregoire Robert, and Baudouin Denis de Senneville. Towards automatic recognition of pure and mixed stones using intra-operative endoscopic digital images. *British Journal of Urology (BJU) International*, 2021. [2](#), [3](#), [4](#), [6](#)
- [11] Daniel Flores-Araiza, Francisco Lopez-Tiro, Elias Villalvazo-Avila, Jonathan El-Beze, Jacques Hubert, Gilberto Ochoa-Ruiz, and Cristian Daul. Interpretable deep learning classifier by detection of prototypical parts on kidney stones images. *arXiv preprint arXiv: Arxiv-2206.00252*, 2022. [3](#), [4](#)
- [12] Justin I Friedlander, Jodi A Antonelli, and Margaret S Pearle. Diet: from food to stone. *World journal of urology*, 33(2):179–185, 2015. [1](#)
- [13] Francisco Lopez, Andres Varelo, Oscar Hinojosa, Mauricio Mendez, Dinh-Hoan Trinh, Yonathan ElBeze, Jacques Hubert, Vincent Estrade, Miguel Gonzalez, Gilberto Ochoa, et al. Assessing deep learning methods for the identification of kidney stones in endoscopic images. In *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pages 2778–2781. IEEE, 2021. [2](#), [3](#), [4](#), [5](#), [6](#)
- [14] Francisco Lopez-Tiro, Juan Pablo Betancur-Rengifo, Arturo Ruiz-Sanchez, Ivan Reyes-Amezcuca, Jonathan El-Beze, Jacques Hubert, Michel Daudon, Gilberto Ochoa-Ruiz, and Christian Daul. Boosting kidney stone identification in endoscopic images using two-step transfer learning. *arXiv preprint arXiv: Arxiv-2210.13654*, 2022. [6](#)
- [15] Adriana Martínez, Dinh-Hoan Trinh, Jonathan El Beze, Jacques Hubert, Pascal Eschwege, Vincent Estrade, Lina Aguilar, Christian Daul, and Gilberto Ochoa. Towards an automated classification method for ureterosopic kidney stone images using ensemble learning. In *2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pages 1936–1939. IEEE, 2020. [6](#)
- [16] Anmar Nassir, Hesham Saada, Taghreed Alnajjar, Jomanah Nasser, Waed Jameel, Soha Elmorsy, and Hattan Badr. The impact of stone composition on renal function. *Urology Annals*, 10(2):215, 2018. [1](#)
- [17] Meike Nauta, Annemarie Jutte, Jesper C. Provoost, and Christin Seifert. This looks like that, because ... explaining prototypes for interpretable image recognition. *CoRR*, abs/2011.02863, 2020. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [11](#)
- [18] Gilberto Ochoa-Ruiz, Vincent Estrade, Francisco Lopez, Daniel Flores-Araiza, Jonathan El Beze, Dinh-Hoan Trinh, Miguel Gonzalez-Mendoza, Pascal Eschwège, Jacques Hubert, and Christian Daul. On the in vivo recognition of kidney stones using machine learning. *arXiv preprint arXiv:2201.08865*, 2022. [2](#), [3](#), [4](#), [5](#), [6](#)
- [19] Ege Gungor Onal and Hakan Tekgul. Assessing kidney stone composition using smartphone microscopy and deep neural networks. *BJUI Compass*, 3:310 – 315, 2022. [2](#)
- [20] Dawid Rymarczyk, Lukasz Struski, Jacek Tabor, and Bartosz Zielinski. Protopshare: Prototype sharing for interpretable image classification and similarity discovery. *CoRR*, abs/2011.14340, 2020. [2](#), [3](#), [4](#)
- [21] Dawid Rymarczyk, Łukasz Struski, Michał Górszczak, Koryna Lewandowska, Jacek Tabor, and Bartosz Zieliński. Interpretable image classification with differentiable prototypes assignment. *arXiv preprint arXiv: Arxiv-2112.02902*, 2021. [3](#)
- [22] Ramprasaath R. Selvaraju, Abhishek Das, Ramakrishna Vedantam, Michael Cogswell, Devi Parikh, and Dhruv Batra. Grad-cam: Why did you say that? visual explanations from deep networks via gradient-based localization. *CoRR*, abs/1610.02391, 2016. [4](#)
- [23] R Siener and A Hesse. Fluid intake and epidemiology of urolithiasis. *European journal of clinical nutrition*, 57(2):S47–S51, 2003. [1](#)
- [24] Shelly Soffer, Avi Ben-Cohen, Orit Shimon, Michal Marianne Amitai, Hayit Greenspan, and Eyal Klang. Convolutional neural networks for radiologic images: A radiologist’s guide. *Epub*, pages 590–606, 2019. [2](#)
- [25] Alejandro Torrell Amado. Metric learning for kidney stone classification. 2018. [6](#)

- [26] Adie Viljoen, Rabia Chaudhry, and John Bycroft. Renal stones. *Annals of clinical biochemistry*, 56(1):15–27, 2019. 1
- [27] Elias Villalvazo-Avila, Francisco Javier Lopez-Tiro, Jonathan El Beze, Jacques Hubert, Miguel González-Mendoza, Gilberto Ochoa-Ruiz, and Christian Daul. Improved kidney stone recognition through attention and multi-view feature fusion strategies. *ArXiv*, abs/2211.02967, 2022. 6
- [28] Heliang Zheng, Jianlong Fu, Tao Mei, and Jiebo Luo. Learning multi-attention convolutional neural network for fine-grained image recognition. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. 4