

SPACESeg: Automated Detection of Bed Junction Morphologies Indicating Signs of Life in Ediacaran Period

Padmaja Jonnalagedda^{†*}, Rachel Surprenant[§], Mary Droser[§], Bir Bhanu[†]
 Dept. of Electrical And Computer Engineering[†], University of California Riverside
 Dept. of Earth and Planetary Sciences[§], University of California Riverside
 {sjonn002, rsurp001, droser, bir.bhanu}@ucr.edu

Abstract

With *Perseverance* out looking for life on Mars, we identify the need to equip ourselves with automated techniques for remote assessment of geological information. The first step in this translational research is studying early signs of life on Earth. More specifically, we examine the Ediacaran sedimentological record of the Flinders region in Australia, whose unique bed ripple junction morphologies have been determined as the definite indicators of early life on Earth. We propose an automated technique, *SPACESeg*, that robustly detects the artifact-clouded, miniature ripple structures from cross-sectional views of the Ediacaran rocks. We demonstrate the efficacy of *SPACESeg* in precisely extracting the desired structures with high accuracy, outperforming many techniques. We also establish the robustness of this technique as it extracts desired biosignatures from drastically varying image conditions, even when the ripples comprise of $< 1\%$ of the image around significant artifacts. We provide quantitative and qualitative analysis and compare our method against many unsupervised rule-based and supervised deep learning methods, outperforming them all.

1. Introduction

The story of the origin and evolution of life on Earth is held within the natural archive of Earth's geologic record which contains the fossils and sedimentological signals necessary to develop our understanding of early evolution on Earth [36, 35]. In turn, study of this terrestrial record can provide valuable pointers for identifying life on other planets. The terminal Ediacaran Period is most well-known for its exceptional preservation of the first multicellular, community-forming organisms on Earth, known colloquially as the Ediacara Biota (574 - 539 Ma) [5, 19]. Not only do the Ediacara Biota hold broad significance as the first complex animals on Earth, the Ediacaran sedimentological

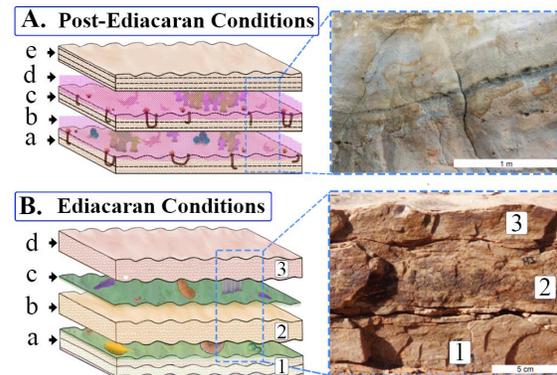


Figure 1: Illustration of cyclic depositional processes under modern and Ediacaran marine conditions (left) and subsequent sedimentary cross-sections (right). A) Modern conditions, no organic mat and the unstabilized seafloor; a. an established modern marine population; b. erosion of the population and top-most rippled surface during a storm event; c. deposition of a new sediment layer during the same storm; d. re-establishment of marine population on storm deposits followed by its erosion and top-most rippled surface during a second storm; e. deposition of new sediment during the second storm. B) Ediacaran conditions, the seafloor is covered and stabilized by organic mats; a: an established Ediacaran ecosystem; b. deposition of sediment during a storm event forming Double-Rippled Bedform (DRB); c. re-establishment of a new macrobiota community on the DRB; d. subsequent storm event resulting in stacked DRBs.

record is unique because the ripples on seafloor were stabilized by ubiquitous organic mats. Therefore, each storm event filled in the ripples instead of eroding them, subsequently creating "Double-Rippled" Bedforms (DRBs or "ripples") between consecutive storms. This resulted in a sedimentological record composed of stacked, discrete storm events with ripples on the top- and bottom-most surfaces (Fig. 1) [19]. This sedimentological record is unique to Ediacaran rocks that preserve organic surfaces and is, therefore, identified as a distinct and definitive biosigna-

*Corresponding author sjonn002@ucr.edu

ture. The detection of rippled biosignatures undoubtedly indicates the presence of past life. Thus we propose an automated technique, called "Scene-aware Perception Automation using Composite Embedding for Segmentation" (SPACESeg) which analyzes the terrain cross-sections to detect the DRBs.

To develop a remote methodology for identifying the presence of these organic mat biosignatures, we turn to the Ediacara Member of the Rawnsley Quartzite in the Flinders Ranges of South Australia. There are multiple localities in the Flinders Ranges that preserve extensive outcrops of the Ediacara Member wherein meters of stratigraphic section are characterized by double-rippled bedforms [8]. In order to tap into the astrobiological and paleobiological potential of DRBs, an objective method of recognizing and characterizing cross-sectional bed junction morphologies of mat-dominated bedforms is needed first. The automation of organic mat cross-sectional identification and subsequent analysis directly contributes to the study of the evolution of multicellular life on Earth. Modeling the parameters using the unequivocally biotically-mediated sedimentary structures from the Ediacaran Period will provide a useful tool for identifying less well-developed biologically important sedimentological signatures in rocks of the deeper Precambrian. Furthermore, our research can be translated to remotely recognize biologically-mediated sedimentary structures on other planets - thereby allowing for the detection and recognition of astrobiologically significant outcrops that have potential to identify the presence of past-life elsewhere in the Universe. This is in line with the main goal of the Mars 2020 Perseverance rover, which NASA states is to "Seek signs of ancient life and collect samples of rock and regolith for possible return to Earth" [7]. Our algorithm can be trained to analyze the images taken by Perseverance of geological features on Mars to recognize the presence or absence of the definitively biologically-mediated double-rippled bedforms, allowing for the rapid, remote and quantitative identification of astrobiologically significant outcrops on Mars that warrant further assessment.

As a computer vision task, the automated detection for DRBs poses many challenges. In addition to the cross-sectional morphology of DRBs being diverse, the data are very limited and the field of view is highly variable. The characteristic bed junction morphology of DRBs is very thin, which adds to the challenges. Furthermore, there are no set standards for image resolution, time of the day (location of sun, intensity variations, haze etc.), perspective frame of reference or external artifacts, among other variables. While these constraints pose a challenge for the development of automated methods, a robust methodology that can overcome these extraneous variables will be most well-suited for extended application to the study of bed junction morphologies in geological cross-sections in dif-

ferent regions on Earth, in other time periods, and on other planets. We propose a hybrid model (*SPACESeg*) based on rule-based optimization (Composite-image Optimized Region Activation or *CORA*) and Saliency Guided Encoder (SGE) to achieve a bed junction segmentation. *SPACE-Seg* adapts to and leverages varying local terrain properties while rejecting perceptually inconsequential features. The benefit of such a model is that it operates with limited domain information and a diverse dataset, while being adaptable to incrementally increasing data availability.

The design of *SPACESeg* tackles a lot of practical problems in computer vision tasks such as the variation in imaging conditions - both device and atmospheric, limited quantity and quality of data and diversity of feature set of real-world settings. In addition to these, *SPACESeg*'s detection algorithm places an importance on perceptual quality as opposed to statistical similarity in feature space. As an example, the bed ripple junctions are almost visually indistinguishable from shadows. However, the shadows provide no usable information in the characterization of the ripples as they have no biotic origin. Thus, we develop our method so it can distinguish between the information-rich ripples and the low-entropy shadow artifacts. With the increasingly practical applications of smart computational technologies, techniques such as *SPACESeg* can be invaluable in extracting and distinguishing usable information from visually similar and useless artifacts.

2. Related Work and Contributions

2.1. Related Work

In Astrobiology and Geology: The unique window into early life and its impact on the environment provided by the Ediacaran fossil record has garnered substantial attention, with extensive and active research being conducted on Ediacaran organic surfaces [9, 10, 34, 37, 25]. The development of a methodology to consistently and correctly identify and describe the sedimentological impact of organic mats has been brought to attention [22], though solutions have thus far focused on the microscopic characterization of microbial fabrics or the subjective identification of discrete organically-induced morphological features [3, 23, 29]. This process is tenuous and is not universally agreed to represent definitive biosignatures, thus limiting their paleobiological and astrobiological utility. The identification and quantitative characterization of the established double-rippled bedforms utilizing nothing more than a standard digital camera provides a rapid, low-cost, and easily replicated process for characterizing surfaces that otherwise evade objective definition.

In Computer Vision and Machine Learning: Semantic segmentation has been one of the most widely researched area of computer vision. There exists a wide array of

traditional and deep network methods proposed for segmentation [2, 21, 14, 31, 38, 39] and boundary detection [24, 33, 15, 1]. Specifically, there has been attention towards scene-aware segmentation [38], concrete crack detection [39] that could potentially translate to our research. However, little precedent exists for the computer-based identification of cross-sectional bed junction morphologies on Earth, let alone on another planet. With increasingly accessible technology, geologists have focused efforts on producing work-flows for the 3D visualization of outcrops, allowing for the field to be brought back to the lab, but still relying on human inference to interpret and gather data from these models [6, 16]. Secondary focus has been placed on utilizing computer vision and image analysis methods in the Geosciences, using these methods to characterize sediment particle movement [26, 28], to conduct areal shape analysis of carbonate reefs using satellite imagery [32], to characterize thin sections [18, 27], to characterize lithology in core samples [12], to process seismic data to reconstruct stenography at depth [4], and to characterize ripple features across a plane using photogrammetric models [17]. While computational methods are clearly widely applied in the Geosciences, we believe ours is the first of any programs developed to characterize cross-sectional geometries of bed junctions in outcrops formed under normal marine or biologically-mediated conditions (e.g., Fig. 1).

In light of the current interest in exobiological investigation as well as the increased accessibility of computer vision methodologies, this research aims to answer the fundamental questions of whether we could develop an approach to quantify the recognition of organically-mediated bedforms on Earth using images so that it could be deployed for use on remote imagery more broadly. In line with this question, we propose computer vision and deep learning techniques to study the impact of widespread organic substrates on cross-sectional geometries of sedimentary sections to develop a robust array of computational tools for translational research, paleontology, and exobiology.

2.2. Contributions

The contributions of this work are:

- Automated detection of biosignatures from sedimentological records, allowing for remote and translational astrobiological analyses
- SPACESeg: A hybrid semi-supervised model that can detect perceptually desired regions from complex scenes
- CORA: An unsupervised technique for isolating region boundaries by employing multi-constraint optimization
- SGE: Supervised segmentation network that is guided by region activations from saliency maps
- Successful detection of ripple bed junctions as thin as 1mm from a 10cm Field of View

3. Technical Approach

We propose a hybrid semi-supervised approach named "Scene-aware Perception Automation using Composite Embedding for Segmentation" (SPACESeg). The task of automating bed ripple junctions poses many challenges, some of which are - variations in time of day (hence shadows), variations in image resolution, field of view, artifacts (such as objects used for scale depiction), etc. The latent space representation to isolate DRBs is high-dimensional. Modeling a complex non-linear feature association is where deep segmentation networks excel. However, due to the lack of data, it is a challenging task. On the other hand, designing unsupervised rule-based techniques are beneficial when the data are limited. But the resultant feature association is either linear or low-level at best, because most techniques employ approximation (ignoring higher-order terms) and a limited set of features to achieve convergence. Our solution is a hybrid model, that combines rule-based (CORA) and neural encoding-based (SGE) models. The resultant SPACESeg is designed to tackle these issues by performing saliency analysis to achieve a low-level rendition of feature association as seed to be provided to the deep segmentation network. Therefore, it adapts to the topography (scene-aware) and rejects the seemingly indistinguishable artifacts to detect only DRBs (perception automation) while guiding a segmentation network towards the global minimum. SPACESeg does so by generating various composite representations of an image and embedding them hierarchically in the latent space to isolate the artifacts, shadows and the desired ripple beds (DRBs). The composite representations are generated using CORA, a Composite-image Optimized Region Activation technique, which are then provided as saliency maps to SGE, that embeds, quantifies and isolates the regions in the latent space. The descriptions of CORA, SGE and subsequent post-processing for DRB analysis are provided in the following subsections.

3.1. Problem Formulation

Let \mathbf{X} be a cross-sectional view of stacked DRBs and the ground-truth DRB annotation be $\hat{\mathbf{Y}}$. Our objective is to isolate the ripple junctions while rejecting the varying conditions of the image. As the image conditions vary and availability of data, it is challenging to isolate ripple junctions accurately without the information on atmospheric conditions, field of view, artifacts etc. To tackle this we utilize an unsupervised Composite-image Optimized Region Activation (CORA) module, that performs a rudimentary scene parsing to generate region activation maps. The region activation maps provide the contextual information for the region of interest, which is otherwise missing in the stacked DRBs images. These activation maps with contextual knowledge guides the Saliency Guided Encoder (SGE) to achieve the final fine-grained segmentation map $\hat{\mathbf{Y}}$.

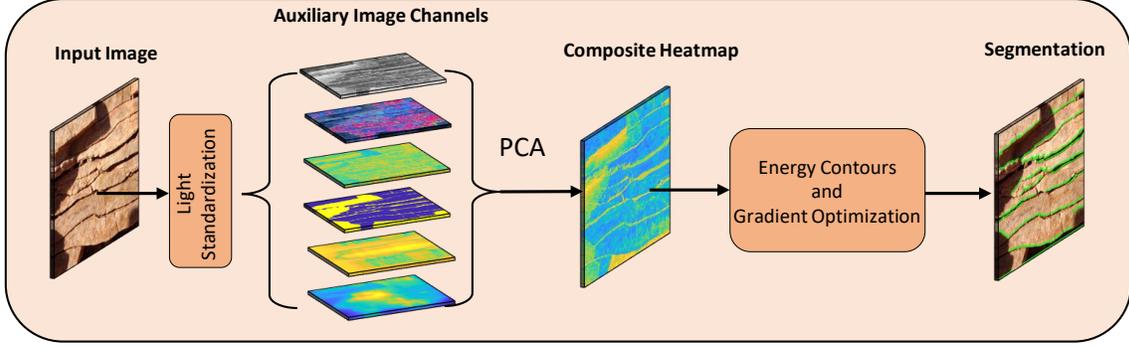


Figure 2: Overview of the Composite-image Optimization Region Activation (CORA) module

3.2. CORA

The overview of Composite-image Optimized Region Activation (CORA) is shown in Fig. 2. The goal of this module is to generate a region activation map via unsupervised multi-constraint optimization. The region activation map highlights the desired bed junctions over imperceptible obstructing artifacts so the subsequent deep segmentation network can train accurately. The algorithm formulation is detailed as follows:

Pre-processing: To address the issue of variations in daylight and brightness conditions in images we employ light standardization technique. While it is complicated to estimate the exact light conditions of the captured image, it is possible to standardize the intensity representation to a certain degree. To standardize the image decomposition and energy functions, we estimate the atmospheric light using a dark channel prior [11]. In this method, the radiance of an image can be estimated by estimating the transmission map (coefficient of light scattered at a given patch of the image). Scene radiance is estimated using inverse Koschmieder law [11] and the resultant image (\mathbf{J}) obtained by filtering using this estimate. We use the same assumptions as in [11] for atmospheric light estimate \mathbf{A} and transmission factors t . The pre-processed image \mathbf{J} is obtained as:

$$\mathbf{J}(\mathbf{x}) = \frac{\mathbf{I}(\mathbf{x}) - \mathbf{A}}{\max(t(x), t_0)} + \mathbf{A} \quad (1)$$

where, x is search space and t_0 is lower-bound of $t(x)$.

Auxiliary Image Representation and Decomposition: Exploratory Data Analysis (EDA) revealed the representations of data that trigger and highlight the features corresponding to ripple junctions, terrain, shadows and artifacts (collectively referred to as image attributes). A rudimentary analysis that included constrained Hough transform, Eigen decomposition, Gabor and Fourier descriptors revealed that the image attributes can be partially isolated in a larger search space of a Composite Energy Gradient (CEG). We, therefore, create a CEG of hybrid image channels. The

observations of EDA are cast on to a 6-channel Auxiliary Image Representation (\mathbf{J}_{aux}). \mathbf{J}_{aux} , as cast into a visible color scale. It is then subjected to Eigen Value Decomposition and Principal Component Analysis (PCA) to create a Composite Image Representation (\mathbf{J}_{comp}). The composite image is cast as a mock-RGB image by combining filtered auxiliary image channel aggregates and an image representation of the first principal component of the PCA. Any direct energy computation on auxiliary channels will reside in a limited search space, similar to employing fuzzy clustering and swarm optimization techniques on an image, which calls for the use of composite image instead of the auxiliary image. The composite image as shown in Fig. 2 is cast onto the perceptual 8-bit range. It can be noted in Fig. 2 that the composite image heat-map already begins to partially isolate shadow from ripple prior to the optimization step, thereby justifying the use of this composite representation.

Composite Energy and Gradient Optimization: This is composed of fuzzy contour optimization step and the gradient optimization step. To activate desired regions, we turned to Active Contour Models (ACMs) [13, 30]. General ACMs are affected by image inhomogeneity, initialization of level-set functions and non-convex properties of the energy function. With the resultant search space of the composite image, we suppress the impact of image inhomogeneity. For the initialization, a binarized image saturation prior acts as a psuedo-level set function. The region energy function is defined as the combination of edge and region energies.

To optimize the energy gradients, we employ the Fuzzy Energy Active Contour (FEAC) [21] method that incorporates fuzzy sets in the contour calculations. A combination of these steps helps us tackle the non-convexity in energy function. The output image is then optionally subjected to secondary gradient optimization that exerts constraints on the orientation, solidity, extent, convexity, saturation and luminescence. We employ FEAC-like energy functions $\mathcal{F}(\mathbf{J}, \mathbf{C}, \delta)$, where \mathbf{C} is the contour and δ is the initialization. We make use of the dark-channel priors from

Y channel of YCbCr space and S channel of HSV space to provide a pseudo-initialization level set δ . From the \mathbf{J}_{aux} , we extract the orientation information (ω) and contour property descriptors ($C_i \in \mathbf{C}$) to define a pseudo-initialized, multi-constrained problem as follows:

$$\min_{\alpha, \beta, \gamma} \min_{\omega} \min_{C_i} \sum_{\mathbf{S}} \mathcal{F}(\mathbf{J}, \mathbf{C}, \delta) \quad (2)$$

where α , β and γ channels are the 3 channels of \mathbf{J}_{comp} instead of R, G, B channel of the original image and \mathbf{S} is the search space based on delta.

The optimization process yields an output that contains partially isolated ripple boundaries. The output of CORA is interpreted as a seed activation and given to SGE as a saliency map (σ). The design of this guided segmentation network is described in the next subsection.

3.3. Saliency Guided Encoder (SGE)

Deep segmentation networks are difficult to train with limited data, and with the added complexity of a diverse and high dimensional feature space, the task difficulty is compounded. Furthermore, for our task, the dataset is practical and, thereby, imperfect. This adds a variety of artifacts that are difficult to characterize. For an automated method, artifact boundaries, regolith and shadows look similar to ripple junctions in the feature space. Testing segmentation deep networks such as U-Net [31], DeepCrack [39], PSPNet [38] etc. revealed the issues of over-segmentation, confusion between artifacts and ripple junctions, and/or sub-optimal convergence. We, therefore, propose Saliency Guided Encoding (SGE, Fig. 3), that leverages the activated maps σ and partially isolated ripple boundaries \mathbf{X}_c obtained from CORA to generate composite embedding. The generated embeddings are associated image attributes at varying levels of abstraction. Detailed technical description of SGE is given in the following subsections.

3.3.1 Composite embedding

Saliency Guided Encoder is composed of two encoders and one decoder module. The encoder \mathcal{E} is for the original input image \mathbf{X}_c and encoder \mathcal{S} is for the region activated map σ . The composite embedding is obtained by combining the output of each layer of encoder and passing it to the next layer of the image encoder (see Fig. 3) as presented below.

$$\sigma_i = \begin{cases} \mathcal{S}_0(\sigma), & \text{if } i == 0. \\ \mathcal{S}_i(\sigma_{i-1}), & \text{otherwise.} \end{cases} \quad (3)$$

$$\mathbf{X}_i = \begin{cases} \mathcal{E}_0(\mathbf{X}_c), & \text{if } i == 0. \\ \mathcal{E}_i(\mathbf{X}_{i-1} \oplus \sigma_{i-1}), & \text{otherwise.} \end{cases} \quad (4)$$

where, \mathcal{E}_i and \mathcal{S}_i represents the i^{th} layer of image and context encoder, respectively; σ_i is the contextual output and

\mathbf{X}_{i-1} is image encoded representation from the i^{th} layer of their respective encoders. \oplus is the concatenation operation.

3.3.2 Hierarchical Segmentation

The composite embedding obtained from each layer is then utilized to perform n-level hierarchical reconstruction. We also include the activation maps here to emphasize on the context information during segmentation. Segmentation of output at each level ($\hat{\mathbf{Y}}_i$) is defined as:

$$\hat{\mathbf{Y}}_i = \begin{cases} \mathcal{G}_i(\mathbf{X}_i \oplus \sigma_i), & \text{if } i == n. \\ \mathcal{G}_i(\hat{\mathbf{Y}}_{i+1} \oplus \sigma_{i-1}), & \text{otherwise.} \end{cases} \quad (5)$$

where, \mathcal{G}_i represents the i^{th} layer of decoder; σ_i is the contextual output and \mathbf{X}_n is image encoded representation from the n^{th} layer and $\hat{\mathbf{Y}}_i$ is the i^{th} -level segmentation output.

3.3.3 Objective function

The SGE does not merely employ an attention block to leverage the information obtained from CORA. Instead, it performs feature multiplexing at every level of abstraction of encoding and fusion-image segmentation at every level of decoding. The combination of feature multiplexing and fusion decoding encompasses the composite embedding part of the SGE. SGE gets its name not only because of saliency embedding, but also because every composite representation is included in a compounded loss function. The loss function for SGE is as follows:

$$\mathcal{L} = \sum_n \mathcal{L}_{rec}^i(\hat{\mathbf{Y}}_i, L_i) \quad (6)$$

where \mathcal{L}_{rec}^i is reconstruction loss at i^{th} level, L_i is the down-sampled label image for respective level and $\hat{\mathbf{Y}}_i$ is segmented output from \mathcal{G}_i .

4. Experimentation and Results

4.1. Dataset

The dataset used for this analysis was obtained from three sources. All of the images from these sources were taken in the Northern Flinders region in Australia. The most high-resolution images covering the largest mat substrate area were collected by MLD. The rest of the images were collected from the research published by [20, 35]. 86% of the data has an associated field of view. For this subset of the data, the field of view ranges from 2cm to 1m, indicating a large variation of the mat substrate associated. The resolution of the images varies from 3-265pixels/cm. The cumulative of the dataset spans > 4 m. There are 28 raw images in the dataset which are subsequently divided into smaller images for assessment. The smaller images are randomized in terms of extent of zooming and crop sizes and

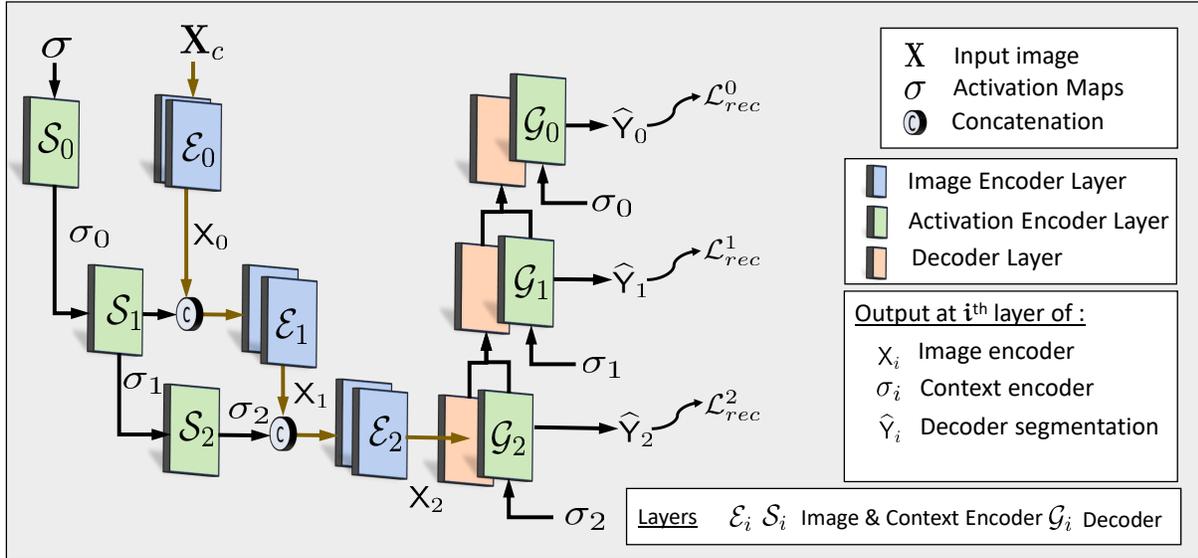


Figure 3: Overview of Saliency Guided Encoder (SGE)

locations. This helps the model to look at varying magnification levels and fields of view of input data. We obtain a total of 2304 sub-images, and this set of data are used for training and evaluation. During CORA processing, the images can be any size. For SGE training, we resize images to 256x256 and employ 80 percent of images for training. For the sake of consistency, all deep networks are trained using the same division of dataset. The qualitative and quantitative results are detailed in the following subsections.

4.2. Qualitative Results

In Fig. 4 and Fig. 5 for CORA and SGE analysis, we picked 3 images based on their varying characteristics. Image (a) was chosen as it has minimal shadow artifact. It does, however, present a difficult artifact - the rock itself. The second-bottom ripple of Image (a) is easy to miss by most methods. Image (b) was chosen to demonstrate the robustness of CORA and SGE against shadow artifacts. Image (c) does not have shadow but its ripple formation is complex and it can be seen that most methods had difficulty detecting the ripples in this image. Additionally, this image has 2 significant artifacts, that CORA and SGE are able to efficiently reject.

CORA: CORA is the first step of our segmentation model SPACESeg. We compare results of CORA against the standard unsupervised segmentation techniques. The number of clusters for Fuzzy c-means (FCM) [2] and Particle Swarm Optimization (PSO) [14] were chosen based on the performance achieved. We tested cluster numbers between 3 to 6 and it was noted that $n=4$ gave the best performance. The results are summarized in Fig. 4. For the active contours, we used Chan-Vese [13]. While CORA

optimization is based on FEAC [21], we test our results against FEAC itself to demonstrate the performance efficacy of CORA against various surfaces. For Image (a), we notice that all methods are unsuccessful in either isolating the shadow artifact or the rock artifact. We, however, note that CORA is successfully able to isolate all shadows and is also able to activate the general ripple region that even looks like the rock. Effect of shadow can also be noted in results for Image (b). For Image (c), while most methods are able to detect the general region of the ripples, they either also include the surrounding rock or in the case of Active Contours, get stuck at a local minima of the artifact due to sensitivity in initialization. The output of CORA precisely detects the ripple boundaries.

SPACESeg: SGE is guided by the output of CORA, which is the second step of our SPACESeg model. We compare our SPACESeg results against UNet [31], PSPNet [38], and DeepCrack [39] which is designed to detect cracks in concrete surfaces, which can be perceived as a similar task. However, we note that DeepCrack is unable to adapt to the varying thickness of the double rippled bedforms. While it does completely reject shadows, it gets penalized when the ripple boundaries are mistaken for shadows. On the other hand, UNet suffers from shadow confusion and ends up over-segmenting. PSPNet is designed for scene parsing and while it does show promise in detecting the right regions, compared to SPACESeg, it falls short in isolating the artifacts. It can be noted that in Image (a), the rock-like ripple is rejected while in Image (b), ripples and shadows are both partly detected and partly rejected. In Image (c), the top of the image reveals the over-segmentation problem because of the very closely spaced ripples. It can be seen that SPACESeg is robust to all these issues.

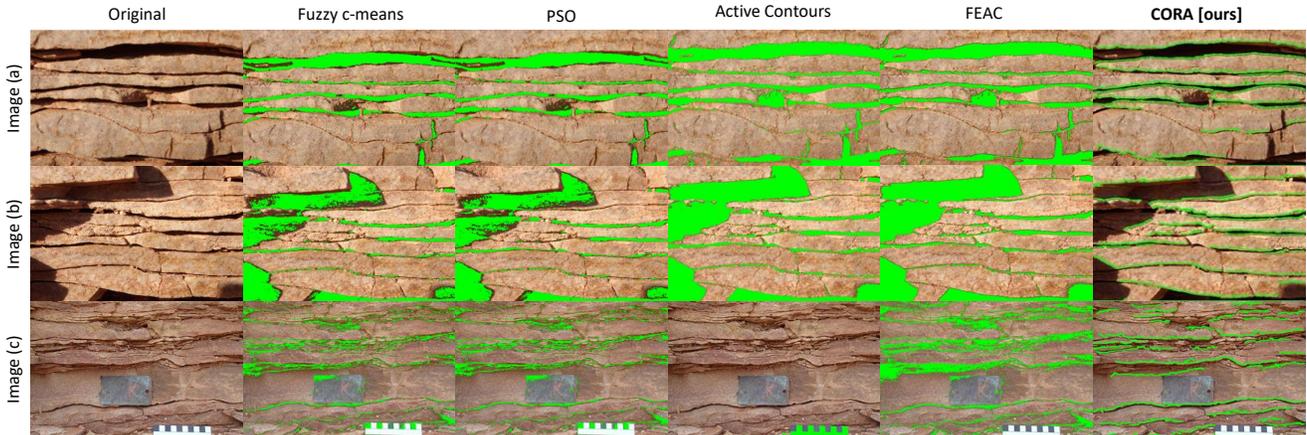


Figure 4: Visual comparison of performance: CORA vs other rule-based methods

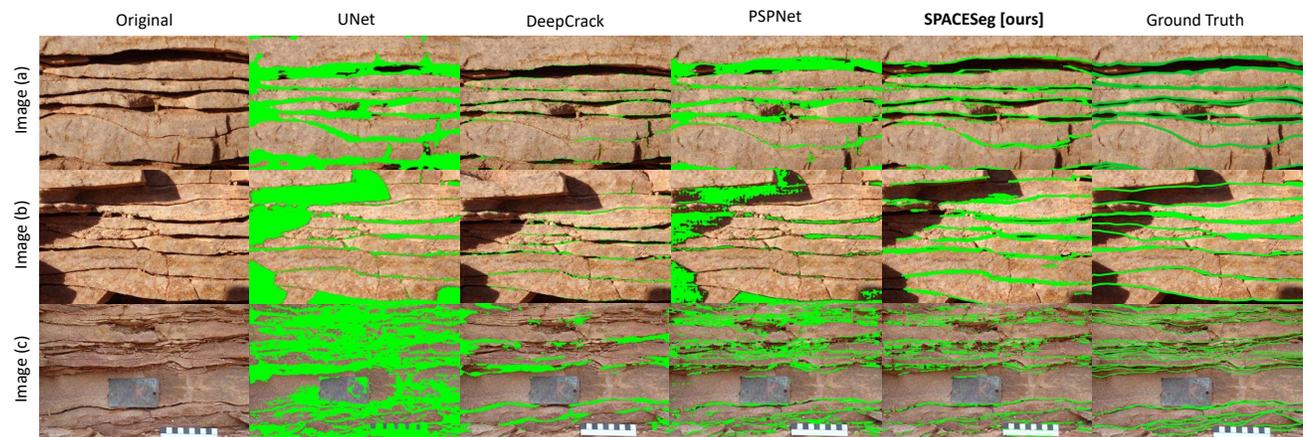


Figure 5: Visual comparison of performance: SPACESeg vs deep segmentation networks

Table 1: Quantitative performance of our proposed SPACESeg (Row 8) vs other rule-based segmentation (Rows 1-4) and deep segmentation networks (Rows 5-7). Metrics - ACC: Accuracy, PRE: Precision, REC: Recall, SSIM: Structural Similarity. Metrics are in %.

Algorithm	ACC	PRE	REC	SSIM
FCM [2]	64 (21)	46 (19)	78 (06)	62 (16)
PSO [14]	67 (19)	48 (16)	76 (12)	64 (16)
Chan-Vese [13]	63 (19)	57 (19)	60 (13)	67 (19)
FEAC [21]	70 (15)	64 (19)	74 (10)	69 (17)
UNet [31]	72 (13)	50 (13)	85 (04)	67 (16)
DeepCrack [39]	78 (12)	62 (03)	70 (12)	76 (14)
PSPNet [38]	80 (08)	58 (10)	82 (06)	70 (14)
SPACESeg (ours)	89 (03)	71 (05)	82 (08)	85 (03)

CORA vs SGE vs SPACESeg: It can be argued that fine-tuning CORA enough or training SGE long enough would yield the same results as the integrated SPACESeg. In Fig. 6, we demonstrate the conditions where each one of them fails. We note that in cases where rocks and artifacts completely overshadow the extremely thin ripples in the center, CORA outputs spurious activations, but SGE picks up the desired regions approximately. On the other hand, as in the case of the image in the second row, where the field of view is too large and not cross-sectional, CORA is able to adapt and learn the scene, thereby bootstrapping the performance of SGE. Thus, SPACESeg as an integrated system is robust to the individual weaknesses of rule-based and trainable transformations of CORA and SGE respectively. The average improvement in SSIM is 14% by using integrated SPACESeg as opposed to SGE or CORA alone.

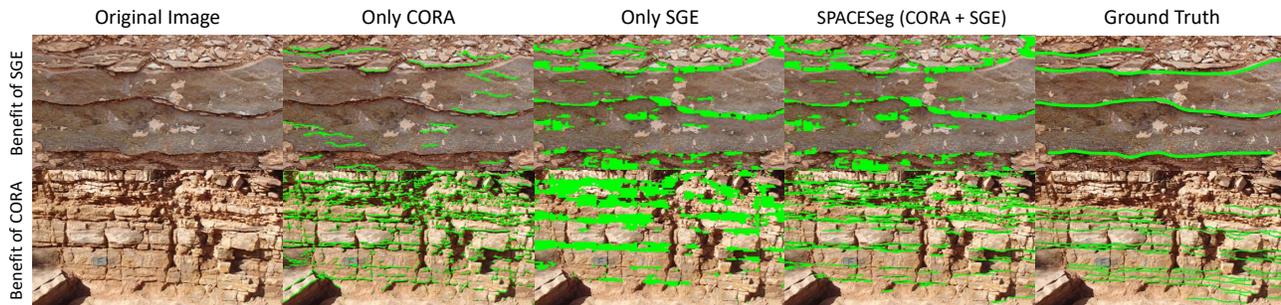


Figure 6: Isolating the effects of CORA and SGE and comparing with the integrated SPACESeg

4.3. Quantitative Results

For all tested algorithms, we report Accuracy (ACC), Precision (PRE), Recall (REC) and Structural Similarity Index Metric (SSIM). Table 1 reports all of the metrics for the unsupervised and supervised techniques as compared with SPACESeg. All metrics are reported as *mean value (std)*. For the deep segmentation models trained using the configurations in the papers were used and fine-tuned. We pretrained these networks on concrete crack images before fine-tuning on our dataset. We report these results instead of trained from scratch as pre-training on domain related dataset improved performance. Analyzing the numbers, it can be seen that for accuracy, the general performance is high, however, the standard deviation is also high. This can be attributed to the fact that ripples occupy a very small fraction of the image and over-segmentation is a common problem for networks training on complex and limited datasets. The precision for networks that get confused between shadow and ripple is low, because while they segment the ripple, it is included with the shadow, causing the precision to fall due to over-segmentation. On the other hand, recall is the extent of ripples being correctly detected, which shoots very high in cases where the shadow encompasses the ripples, but not so much when other artifacts (such as ripple looking like the rock surface) are introduced. Thus, the precision and recall alternate between highs and lows for most techniques, based on extent of artifact vs ripple presence, as indicated by the high standard deviation. Since the ripple structures have varying thickness and shapes, we also report SSIM. As opposed to the other metrics, SSIM will include the overall properties of the structures being segmented, emphasizing on whether perceptually meaningful regions have been picked up by the networks or not.

It can be seen that SPACESeg has consistently the best performance among all methods across all metrics. It has a slightly lower performance than UNet in recall because UNet does over-segmentation, which increases its sensitivity, which is an issue that SPACESeg efficiently tackles. The holistic assessment of the qualitative and quantitative

performance of SPACESeg demonstrates that not only is it efficient in learning under adverse data conditions, it also reports perceptually meaningful outcomes.

5. Conclusions

There is an increasing interest in astrobiological research, and resources are being invested in the search for life on other planets. In such a pivotal point in this field, data are a valuable and elusive resource. It is, therefore, important to develop tools that allow for remote analysis of data. The Ediacaran period holds massive significance in its unique capability to present us with definitive biosignatures that directly correlate to evolution of life. This study can open avenues for translational research in studying the Cambrian evolution of life and take us many steps closer to studying life signatures on other planets. To this end, we proposed SPACESeg, a robust detector of biosignatures from the Ediacaran organic substrates. SPACESeg’s hybrid approach of CORA and SGE allows it to tackle the biggest issues that real-world data faces. Practical dataset poses many challenges in its limited quantities, largely varying conditions and uncontrolled artifact manifestations. Moreover, it is able to do so while detecting miniature structures from the images. We report quantitative and qualitative results supporting this and note that SPACESeg outperforms other methods by at least 10%, while efficiently isolating human perceivable outcomes. The robustness of SPACESeg opens doors for scalable, translational, remote assessment tools that can aid in our search for signs of terrestrial and extraterrestrial life.

6. Acknowledgements

We thank Ross and Jane Fargher for access to the National Heritage Nilpena Ediacara fossil site on their property, acknowledging that this land lies within the Adnyamathanha Traditional Lands. The project was funded by NASA Exobiology Program (Grant 80NSSC19K0472) to MLD. Ian Hughes, Emmy Hughes, Scott Evans, Matt Dzaugis and Jim Gehling aided in field work.

References

- [1] Evgeny E Baraboshkin, Leyla S Ismailova, Denis M Orlov, Elena A Zhukovskaya, Georgy A Kalmykov, Oleg V Khotylev, Evgeny Yu Baraboshkin, and Dmitry A Koroteev. Deep convolutions for in-depth automated rock typing. *Computers & Geosciences*, 135:104330, 2020. [3](#)
- [2] James C Bezdek, Robert Ehrlich, and William Full. Fcm: The fuzzy c-means clustering algorithm. *Computers & geosciences*, 10(2-3):191–203, 1984. [3](#), [6](#), [7](#)
- [3] Sherry L. Cady, Jack D. Farmer, John P. Grotzinger, J. William Schopf, and Andrew Steele. Morphological biosignatures and the search for life on mars. *Astrobiology*, 3(2):351–368, 2003. [2](#)
- [4] Haibin Di, Leigh Truelove, Cen Li, and Aria Abubakar. Accelerating seismic fault and stratigraphy interpretation with deep cnns: A case study of the taranaki basin, new zealand. *The Leading Edge*, 39(10):727–733, 2020. [3](#)
- [5] Mary L. Droser, Lidya G. Tarhan, and James G. Gehling. The rise of animals in a changing environment: Global ecological innovation in the late ediacaran. *Annual Review of Earth and Planetary Sciences*, 45(1):593–617, 2017. [1](#)
- [6] Anette Eltner, Andreas Kaiser, Carlos Castillo, Gilles Rock, Fabian Neugirg, and Antonio Abellán. Image-based surface reconstruction in geomorphometry—merits, limits and developments. *Earth Surface Dynamics*, 4(2):359–389, 2016. [3](#)
- [7] Kenneth A Farley, Kenneth H Williford, Kathryn M Stack, Rohit Bhartia, Al Chen, Manuel de la Torre, Kevin Hand, Yulia Goreva, Christopher DK Herd, Ricardo Hueso, et al. Mars 2020 mission overview. *Space Science Reviews*, 216(8):1–41, 2020. [2](#)
- [8] J.g. Gehling. Environmental interpretation and a sequence stratigraphic framework for the terminal proterozoic ediacara member within the rawnsley quartzite, south australia. *Precambrian Research*, 100(1-3):65–95, 2000. [2](#)
- [9] James G. Gehling. Microbial mats in terminal proterozoic siliciclastics: Ediacaran death masks. *Palaios*, 14(1):40, 1999. [2](#)
- [10] Murray Gingras, James W. Hagadorn, Adolf Seilacher, Stefan V. Lalonde, Ernesto Pecoits, Daniel Petrash, and Kurt O. Konhauser. Possible evolution of mobile animals in association with microbial mats. *Nature Geoscience*, 4(6):372–375, 2011. [2](#)
- [11] Kaiming He, Jian Sun, and Xiaoou Tang. Single image haze removal using dark channel prior. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(12):2341–2353, 2010. [4](#)
- [12] Alexander V Ivchenko, Evgeny E Baraboshkin, Leyla S Ismailova, Denis M Orlov, Dmitry A Koroteev, and Evgeny Yu Baraboshkin. Conference paper, 2018. [3](#)
- [13] Michael Kass, Andrew Witkin, and Demetri Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1988. [4](#), [6](#), [7](#)
- [14] James Kennedy and Russell Eberhart. Particle swarm optimization. In *Proceedings of ICNN'95-international conference on neural networks*, volume 4, pages 1942–1948. IEEE, 1995. [3](#), [6](#), [7](#)
- [15] Iasonas Kokkinos. Pushing the boundaries of boundary detection using deep learning. *arXiv preprint arXiv:1511.07386*, 2015. [3](#)
- [16] Branko Kordić, Borna Lužar-Oberiter, Kristina Pikelj, Bojan Matoš, and Goran Vlastelica. Integration of terrestrial laser scanning and uas photogrammetry in geological studies: examples from croatia. *Periodica Polytechnica. Civil Engineering*, 63(4):989, 2019. [3](#)
- [17] Lu Li, Ran Wang, Jingyu Lin, Zhouxuan Xiao, and Yuanxiu Hui. A novel approach for extraction of ripple mark parameters based on sfm. *Sedimentary Geology*, 392:105523, 2019. [3](#)
- [18] Xiaokang Liu and Haijun Song. Automatic identification of fossils and abiotic grains during carbonate microfacies analysis using deep convolutional neural networks. *Sedimentary Geology*, 410:105790, 2020. [3](#)
- [19] Jack J. Matthews, Alex G. Liu, Chuan Yang, Duncan McIlroy, Bruce Levell, and Daniel Condon. A chronostratigraphic framework for the rise of the ediacaran macrobiota: New constraints from mistaken point ecological reserve, new-foundland. 2020. [1](#)
- [20] William McMahon, Alexander Liu, Benjamin Tindal, and Martin Kleinhans. Ediacaran life close to land: coastal and shoreface habitats of the ediacaran macrobiota in the central and southern flinders ranges, south australia. 2020. [5](#)
- [21] Ajoy Mondal, Susmita Ghosh, and Ashish Ghosh. Robust global and local fuzzy energy based active contour for image segmentation. *Applied Soft Computing*, 47:191–215, 2016. [3](#), [4](#), [6](#), [7](#)
- [22] Nora Noffke. The criteria for the biogenicity of microbially induced sedimentary structures (miss) in archean and younger, sandy deposits. *Earth-Science Reviews*, 96(3):173–180, 2009. [2](#)
- [23] Nora Noffke. Ancient sedimentary structures in the j3.7 ga gillespie lake member, mars, that resemble macroscopic morphology, spatial associations, and temporal succession in terrestrial microbialites. *Astrobiology*, 15(2):169–192, 2015. [2](#)
- [24] Robert J O’Callaghan and David R Bull. Combined morphological-spectral unsupervised image segmentation. *IEEE Transactions on Image Processing*, 14(1):49–62, 2004. [3](#)
- [25] Gretchen R. Oneil, Lydia S. Tackett, and Mike Meyer. Petrographic evidence for ediacaran microbial mat-targeted behaviors from the great basin, united states. *Precambrian Research*, 345:105768, 2020. [2](#)
- [26] Athanasios N. Papanicolaou, Panayiotis Diplas, Mahesh Balakrishnan, and Clinton L. Dancy. Computer vision technique for tracking bed load movement. *Journal of Computing in Civil Engineering*, 13(2):71–79, 1999. [3](#)
- [27] S. Polushkin, Y. Volokitin, I. Edelman, E. Sadykhov, O. Lokhanova, Y. Murzaev, and S. Pastushkov. Processing thin section photos with neural networks and computer vision. *First EAGE Digitalization Conference and Exhibition*, 2020. [3](#)
- [28] Alessio Radice, Sankar Sankar, and Francesco Ballio. Image-based lagrangian particle tracking in bed-load experiments. *Journal of Visualized Experiments*, (125), 2017. [3](#)

- [29] Vincenzo Rizzo. Why should geological criteria used on earth not be valid also for mars? evidence of possible microbialites and algae in extinct martian lakes. *International Journal of Astrobiology*, 19(3):283–294, 2020. [2](#)
- [30] Remi Ronfard. Region-based strategies for active contour models. *International Journal of Computer Vision*, 13(2):229–251, 1994. [4](#)
- [31] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation, 2015. [3](#), [5](#), [6](#), [7](#)
- [32] Grisel Jimenez Soto, Mirza Arshad Beg, Michael C. Poppelreiter, and Khaidi Rahmatsyah. Automatic extraction of modern reefs satellite images geometries using computer vision. 2020. [3](#)
- [33] Andrew N Stein, Thomas S Stepleton, and Martial Hebert. Towards unsupervised whole-object segmentation: Combining automated matting with boundary detection. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008. [3](#)
- [34] Lidya G. Tarhan, Mary L. Droser, James G. Gehling, and Matthew P. Dzaugis. Microbial mat sandwiches and other anactualistic sedimentary features of the ediacara member (rawnsley quartzite, south australia): Implications for interpretation of the ediacaran sedimentary record. *Palaios*, 32(3):181–194, 2017. [2](#)
- [35] Lidya G Tarhan, Mary L Droser, James G Gehling, and Matthew P Dzaugis. Microbial mat sandwiches and other anactualistic sedimentary features of the ediacara member (rawnsley quartzite, south australia): Implications for interpretation of the ediacaran sedimentary record. *Palaios*, 32(3):181–194, 2017. [1](#), [5](#)
- [36] Shuhai Xiao, Zhe Chen, Chuanming Zhou, and Xunlai Yuan. Surfing in and on microbial mats: Oxygen-related behavior of a terminal ediacaran bilaterian animal. *Geology*, 47(11):1054–1058, 2019. [1](#)
- [37] Shuhai Xiao, Zhe Chen, Chuanming Zhou, and Xunlai Yuan. Surfing in and on microbial mats: Oxygen-related behavior of a terminal ediacaran bilaterian animal. *Geology*, 47(11):1054–1058, 2019. [2](#)
- [38] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2881–2890, 2017. [3](#), [5](#), [6](#), [7](#)
- [39] Qin Zou, Zheng Zhang, Qingquan Li, Xianbiao Qi, Qian Wang, and Song Wang. Deepcrack: Learning hierarchical convolutional features for crack detection. *IEEE Transactions on Image Processing*, 28(3):1498–1512, 2018. [3](#), [5](#), [6](#), [7](#)