



ISTI Technical Reports

Using Spatial Logic and Model Checking for Nevus Segmentation

Gina Belmonte, ATNO Fisica Sanitaria Nord, Lucca, Italy

Giovanna Broccia, CNR-ISTI, Pisa, Italy

Vincenzo Ciancia, CNR-ISTI, Pisa, Italy

Diego Latella, CNR-ISTI, Pisa, Italy

Mieke Massink, CNR-ISTI, Pisa, Italy



Using Spatial Logic and Model Checking for Nevus Segmentation

Belmonte G.; Broccia G.; Ciancia V.; Latella D.; Massink M.

ISTI-TR-2020/017

Abstract

Spatial and spatio-temporal model checking techniques have a wide range of application domains, among which large scale distributed systems and signal and image analysis. In the latter domain, automatic and semi-automatic contouring in Medical Imaging has shown to be a very promising and versatile application that can greatly facilitate the work of professionals in this domain, while supporting explainability, easy replicability and exchange of medical image analysis methods. In recent work we have applied this model-checking technique to the (3D) contouring of tumours and related oedema in magnetic resonance images of the brain. In the current work we address the contouring of (2D) images of nevi. One of the challenges of treating nevi images is their considerable inhomogeneity in shape, colour, texture and size. To deal with this challenge we use a texture similarity operator, in combination with spatial logic operators. We apply our technique on images of a large public database and compare the results with associated ground truth segmentation provided by domain experts.

Spatial Logics, Spatial Model-checking, Nevus Segmentation

Citation

Belmonte G.; Broccia G.; Ciancia V.; Latella D.; Massink M. *Using Spatial Logic and Model Checking for Nevus Segmentation*. ISTI Technical Reports 2020/017. DOI: 10.32079/ISTI-TR-2020/017

Istituto di Scienza e Tecnologie dell'Informazione "A. Faedo"

Area della Ricerca CNR di Pisa

Via G. Moruzzi 1

56124 Pisa Italy

<http://www.isti.cnr.it>

Using Spatial Logic and Model Checking for Nevus Segmentation^{*}

Gina Belmonte¹, Giovanna Broccia², Vincenzo Ciancia²,
Diego Latella², and Mieke Massink²

¹ Azienda Toscana Nord Ovest S. C. Fisica Sanitaria Nord, Lucca, Italy

² Consiglio Nazionale delle Ricerche - Istituto di Scienza e Tecnologie
dell'Informazione 'A. Faedo', CNR, Pisa, Italy

Abstract. Spatial and spatio-temporal model checking techniques have a wide range of application domains, among which large scale distributed systems and signal and image analysis. In the latter domain, automatic and semi-automatic contouring in Medical Imaging has shown to be a very promising and versatile application that can greatly facilitate the work of professionals in this domain, while supporting *explainability*, *easy replicability* and *exchange* of medical image analysis methods. In recent work we have applied this model-checking technique to the (3D) contouring of tumours and related oedema in magnetic resonance images of the brain. In the current work we address the contouring of (2D) images of nevi. One of the challenges of treating nevi images is their considerable inhomogeneity in shape, colour, texture and size. To deal with this challenge we use a *texture similarity* operator, in combination with spatial logic operators. We apply our technique on images of a large public database and compare the results with associated ground truth segmentation provided by domain experts.

Keywords: Spatial logics; Model Checking; Medical Imaging;

1 Introduction and Related Work

A nevus is a visible, usually small and benign, circumscribed lesion of the skin. Unfortunately, in some cases these are hard to distinguish from their malignant counterpart known as Melanocytic nevus. Melanoma is a very serious form of skin cancer. It may be lethal if the disease is not recognised in a very early stage.

^{*} Part of this work has been developed in the context of the Italian MIUR-PRIN 2017 project IT MaTTerS: Methods and Tools for Trustworthy Smart Systems[†]. The names of the authors are listed in alphabetical order. The major contributions to the development of the specification in `ImgQL` are those of G. Broccia and M. Massink, who also have been taking care running it on the datasets and collecting the results. D. Latella and V. Ciancia also contributed to the development of the specification. V. Ciancia developed `VoxLogicA` and G. Belmonte proposed the particular texture analysis approach. All authors contributed to drafting and finalising the paper.

In Europe alone melanoma causes over 20,000 deaths each year [21]. One of the difficulties is that reliable early detection requires highly trained specialists but in many countries there is only a limited number of such specialists available. It is therefore no surprise that there is much interest in automated systems that can help recognising the disease reliably and at an early stage so that more lives could be saved and the number of unnecessary biopsies can be reduced [17].

The most popular and well-performing automated techniques for the diagnosis of melanoma at the moment rely on deep learning [17]. In this paper we take a different approach based on recently developed *spatial* model checking techniques, in particular for the contouring or segmentation of nevi, which is one of the sub-tasks involved in the diagnosis of melanoma. In our previous work on (semi-) automatic contouring using spatial model checking techniques for the contouring of various kinds of brain tissues and brain tumours [4,6,5,2] we have shown that such techniques can reach a segmentation quality that is competitive with state-of-the-art techniques, while supporting explainability, easy replicability and exchange of medical image analysis methods. The segmentation of nevi poses additional challenges because dermoscopic images of nevi tend to be very dis-homogeneous in size, colour, contrast, location and kind of nevus/lesion and the presence of additional objects such as coloured patches, hairs, shadows and other optical effects.

Spatial (and spatio-temporal) model checkers use high-level, often domain oriented, specifications written in a logical language to describe spatial properties in order to automatically and efficiently identify spatial patterns and structures of interest. The origins of spatial logic can be traced back to the forties of the previous century when McKinsey and Tarski recognised the possibility of reasoning on space using topology as a mathematical framework for the interpretation of modal logic (see [1] for a thorough introduction). In recent work by Ciancia et al. [13,14] *Closure Spaces*, a generalisation of topological spaces, were used as underlying model for discrete spatial logic inspired by recent work by Galton [22,23,34,24]. This resulted in the definition of the Spatial Logic for Closure Spaces (SLCS), and a related model checking algorithm. Furthermore, in [12], a spatio-temporal logic, combining Computation Tree Logic with the spatial operators of SLCS was introduced. Spatial and spatio-temporal model checking have recently been applied in a variety of domains, ranging from Collective Adaptive Systems [11,15,16] to signals [33] and medical images [2,3,4,5,6].

Several proposals of use of computational methods for the analysis of medical images are available in the literature. *Computer-Aided Diagnosis* (CAD) aims at the classification of areas in images, based on the presence of signs of specific diseases [19]. *Image Segmentation* [25] is focused on the identification of areas that have specific features or perform specific functions. *Automatic contouring of Organs at Risk* or target volumes [7] is specifically devoted to supporting radiotherapy applications. Finally, specific *indicators* can be computed from the acquired images that can enable early diagnosis—or the understanding of microscopic characteristics of specific diseases—or can help in the identification of prognostic factors to predict a treatment output [10,41]. In [26] spiral electric

waves—a precursor to atrial and ventricular fibrillation—are detected and specified using a spatial logic and model-checking tools. The formulas of the logic are learned from the spatial patterns under investigation and the onset of spiral waves is detected using bounded model checking. In [28] a logic called Spatial-Temporal Logic (SpaTeL) is defined that is a unification of signal temporal logic (STL) and tree spatial superposition logic (TSSL). The logic can be used for describing high-level spatial patterns that change over time.

In our previous work on the use of spatial model-checking for the analysis of medical images mentioned earlier, we focused on image segmentation, in particular for the identification of glioblastomas—which are the most common malignant intracranial tumours—but also of regions of interest in healthy organs [5].

In this paper, we apply spatial model-checking to images of nevi from a public dataset released by the International Skin Imaging Collaboration (ISIC) for the 2016³ International Symposium on Biomedical Imaging (ISBI 2016) challenge titled “Skin Lesion Analysis toward Melanoma Detection” [17]. This dataset contains 900 annotated dermoscopic images, obtained by specialised high-resolution imaging of the skin that reduces skin surface reflectance. Among this set are 173 images of melanomas. Each image in the dataset has been segmented manually by experts and their segmentation result is available as ground truth images, which makes comparison with results of other state-of-the-art segmentation techniques applied to the same dataset possible. The original challenge consisted of three parts: Lesion Segmentation, lesion Dermoscopic Feature Extraction, and Lesion Classification. In the present work we focus on lesion segmentation.

The outline of the paper is as follows. Section 2 provides some background on spatial model checking, the spatial model checker `VoxLogicA` and in particular its texture similarity operator. Section 3 presents the spatial logic specification for the segmentation of nevi and Section 4 presents the model checking results on (subsets) of the ISIC 2016 training dataset. Finally, Section 5 presents the conclusions of this work.

2 Background on Spatial Model Checking

ImgQL (*Image Query Language*), first proposed in [2,6], is a spatial logical language developed for the analysis of medical images. It is based on SLCS (*Spatial Logic for Closure Spaces*) [13,14]. ImgQL is also the input language for the spatial model checker `VoxLogicA` presented in [6]. In this section we first recall the definition of the logical kernel of ImgQL and the underlying basic notions and then we show its extension supported by the tool. We refer to our earlier work for further details on theoretical aspects and the spatial model checking algorithms [13,14,2,6].

³ We focus on the 2016 Challenge, that is the first of a series; We leave the subsequent challenges for future work.

2.1 The logical kernel of ImgQL

SLCS is interpreted over *closure spaces*. A closure space—CS for short—is a pair (X, \mathcal{C}) where X is a set (of points) and $\mathcal{C} : 2^X \rightarrow 2^X$ is a function satisfying the following three axioms: (i) $\mathcal{C}(\emptyset) = \emptyset$; (ii) $Y \subseteq \mathcal{C}(Y)$ for all $Y \subseteq X$; (iii) $\mathcal{C}(Y_1 \cup Y_2) = \mathcal{C}(Y_1) \cup \mathcal{C}(Y_2)$ for all $Y_1, Y_2 \subseteq X$. The *interior* of a set $Y \subseteq X$ is obtained by duality, i.e. $\mathcal{I}(Y) = \mathcal{C}(\overline{Y})$ where $\overline{Y} = X \setminus Y$ is the complement of Y . Given any relation $R \subseteq X \times X$, (X, \mathcal{C}_R) , with $\mathcal{C}_R(Y) = Y \cup \{x \mid \exists y \in Y. y R x\}$, is a CS. In particular, a digital image can be modeled as a finite CS where X is the set of voxels and R their (reflexive and symmetric) *adjacency relation*⁴.

$(\mathbb{N}, \mathcal{C}_{succ})$ is the CS of the natural numbers \mathbb{N} with the binary successor relation $succ = \{(m, n) \in \mathbb{N}^2 \mid n = m + 1\}$. A (discrete) *path* over (X, \mathcal{C}) is a continuous function⁵ from $(\mathbb{N}, \mathcal{C}_{succ})$ to (X, \mathcal{C}) .

It is often convenient to equip the elements of X with *attributes* in a given set A over a given set of values V ; an *attributed CS* is a structure $((X, \mathcal{C}), \mathcal{A})$ where (X, \mathcal{C}) is a CS and $\mathcal{A} : A \times X \rightarrow V$, is the attribute evaluation function, such that $\mathcal{A}(a, x)$ maps attribute (named) a of point x to its value in V . For instance, if x is a voxel, then $\mathcal{A}(red, x)$ may represent the intensity of red of x , and similarly for $\mathcal{A}(green, x)$ and $\mathcal{A}(blue, x)$. Attribute values can be used in expressions α over V ; consequently function \mathcal{A} is assumed lifted to such expressions in the standard way.

In this paper we will use *distance CS*, i.e. structures $((X, \mathcal{C}), d)$ where $d : X \times X \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ is a *distance function*⁶, i.e. it satisfies $d(x, y) = 0$ if and only if $x = y$; d is lifted to sets in the usual way: $d(x, \emptyset) = \infty$ and for $\emptyset \subset Y \subseteq X$ $d(x, Y) = \inf\{d(x, y) \mid y \in Y\}$.

ImgQL is interpreted over *attributed distance closure models*, i.e. structures $((X, \mathcal{C}), d, \mathcal{A}, \mathcal{V})$ where (X, \mathcal{C}) is a CS, d and \mathcal{A} are the distance and the attribute evaluation functions, respectively, and $\mathcal{V} : P \rightarrow 2^X$ is a valuation which maps the *atomic predicates* of a given set P to the points satisfying them. In the sequel we recall the formal definition of the logical kernel of ImgQL:

Definition 1. For given set P of atomic predicates p , and interval I of \mathbb{R} , the syntax of ImgQL is the following:

$$\Phi ::= p \mid \neg \Phi \mid \Phi_1 \wedge \Phi_2 \mid \overset{\rightarrow}{\rho} \Phi_1[\Phi_2] \mid \overset{\leftarrow}{\rho} \Phi_1[\Phi_2] \mid \mathcal{D}^I \Phi.$$

Defined predicates are elements p of P for which a defining equation $p := \alpha$ is given, where α is an expression.

⁴ All the theory and related model checkers work both for 2D and 3D even though we use only 2D in the current work. Similarly, in the current work we use the word ‘voxel’ both for 3D ‘pixels’ and for 2D pixels.

⁵ A *continuous* function from CS (X_1, \mathcal{C}_1) to CS (X_2, \mathcal{C}_2) is a function $f : X_1 \rightarrow X_2$ such that $f(\mathcal{C}_1(Y)) \subseteq \mathcal{C}_2(f(Y))$ for all $Y \subseteq X_1$.

⁶ Several distance functions are defined in the literature; the specific distance to be used depends on the application. The interested reader is referred to [2]. In this work we use the Manhattan distance where 1 voxel is the unit distance.

Satisfaction $\mathcal{M}, x \models \Phi$ of a formula Φ at point $x \in X$ in model $\mathcal{M} = ((X, \mathcal{C}), d), \mathcal{A}, \mathcal{V}$ is defined recursively on the structure of formulas, where $[[\Phi]]^{\mathcal{M}}$ is the set $\{x \in X \mid \mathcal{M}, x \models \Phi\}$ of points satisfying Φ in \mathcal{M} :

$$\begin{aligned}
\mathcal{M}, x \models p \in P &\Leftrightarrow x \in \mathcal{V}(p) \\
\mathcal{M}, x \models \neg \Phi &\Leftrightarrow \mathcal{M}, x \not\models \Phi \text{ does not hold} \\
\mathcal{M}, x \models \Phi_1 \wedge \Phi_2 &\Leftrightarrow \mathcal{M}, x \models \Phi_1 \text{ and } \mathcal{M}, x \models \Phi_2 \\
\mathcal{M}, x \models \overset{\rightarrow}{\rho} \Phi_1[\Phi_2] &\Leftrightarrow \text{there is path } \pi \text{ and index } \ell \text{ s.t. } \pi(0) = x \text{ and } \mathcal{M}, \pi(\ell) \models \Phi_1 \\
&\quad \text{and for all indexes } j : 0 < j < \ell \text{ implies } \mathcal{M}, \pi(j) \models \Phi_2 \\
\mathcal{M}, x \models \overset{\leftarrow}{\rho} \Phi_1[\Phi_2] &\Leftrightarrow \text{there is path } \pi \text{ and index } \ell \text{ s.t. } \pi(\ell) = x \text{ and } \mathcal{M}, \pi(0) \models \Phi_1 \\
&\quad \text{and for all indexes } j : 0 < j < \ell \text{ implies } \mathcal{M}, \pi(j) \models \Phi_2 \\
\mathcal{M}, x \models \mathcal{D}^I \Phi &\Leftrightarrow d(x, [[\Phi]]^{\mathcal{M}}) \in I
\end{aligned}$$

Whenever p is a defined predicate with defining equation $p := \alpha$, we extend the satisfaction relation by letting $x \in \mathcal{V}(p)$ if and only if $\mathcal{A}(x, \alpha)$ is true. •

Classical derived operators are defined as usual: $\perp \equiv p \wedge \neg p$, $\top \equiv \neg \perp$, $\Phi_1 \vee \Phi_2 \equiv \neg(\neg \Phi_1 \wedge \neg \Phi_2)$ etc. In addition, we have the following more specific derived operators:

$$\begin{aligned}
\mathcal{N}\Phi &\equiv \overset{\leftarrow}{\rho} \Phi[\perp] \\
\Phi_1 \mathcal{S} \Phi_2 &\equiv \Phi_1 \wedge \neg \overset{\rightarrow}{\rho} \neg(\Phi_1 \vee \Phi_2)[\neg \Phi_2] \\
\text{touch}(\Phi_1, \Phi_2) &\equiv \Phi_1 \wedge \overset{\rightarrow}{\rho} \Phi_2[\Phi_1] \\
\text{grow}(\Phi_1, \Phi_2) &\equiv \Phi_1 \vee \text{touch}(\Phi_2, \Phi_1) \\
\text{smoothen}(r, \Phi_1) &\equiv \mathcal{D}^{<r}(\mathcal{D}^{\geq r} \neg \Phi_1).
\end{aligned}$$

Intuitively, a point x satisfies $\mathcal{N}\Phi$ if it is *near* Φ , i.e. if it can be reached *in one step* from a point laying in $[[\Phi]]$; it is easy to see that $\mathcal{M}, x \models \mathcal{N}\Phi$ if and only if $x \in \mathcal{C}([[\Phi]])^{\mathcal{M}}$. A point x satisfies $\Phi_1 \mathcal{S} \Phi_2$ if it lays in an area, where all points satisfy Φ_1 , that is *surrounded* by points satisfying Φ_2 , i.e. it is impossible to find a path starting from x that can reach a point satisfying neither Φ_1 nor Φ_2 , without first passing through a point satisfying Φ_2 ⁷. The meaning of $\text{touch}(\Phi_1, \Phi_2)$ should be clear. A point satisfies $\text{grow}(\Phi_1, \Phi_2)$ if it satisfies Φ_1 or it lays in a path of points all satisfying Φ_2 and leading to a point satisfying Φ_1 . A formula $\text{smoothen}(r, \Phi_1)$ is satisfied by points that are at a distance of less than r from a point that is at least at distance r from points that do not satisfy Φ_1 . This operator works as a filter; only contiguous areas satisfying Φ_1 that have a minimal diameter of at least $2r$ are preserved; these are also smoothed if they have an irregular shape (e.g. protrusions with a width that is less than the indicated distance).

We close this section with the description of an additional logical operator of ImgQL, namely the *Texture Analysis* operator Δ introduced in [6]. Texture

⁷ Note that in [13,14,2,6] \mathcal{N} and \mathcal{S} (denoted by \mathcal{U} in [13]) have been presented as basic operators while reachability operators have been defined as derived from the formers.

Analysis (TA) is an approach used for finding patterns in (medical) images. The approach has proved promising in a large number of applications in the field of medical imaging [31,32,8,18]; in particular it has been used in *Computer Aided Diagnosis* [42,29,30] and for classification or segmentation of tissues or organs [9,37,35]. The ImgQL TA operator Δ uses *first order* statistical methods⁸ and differs from those in the classical setting, e.g. [39,40], where the various moments (*mean*, *variance* etc.) of distributions of the two pictures to be compared are analysed. In ImgQL, instead, the statistical distributions—actually the histograms—of the two pictures are compared directly, using, as similarity measure, their *cross-correlation* (also called *Pearson's correlation coefficient*).

The intuitive semantics of Δ is presented schematically in Figure 1. Suppose we have an ImgQL formula Φ that specifies a *sample area of interest* $[\Phi]^M$ in a given image; suppose, furthermore, that the feature that makes the points in $[\Phi]^M$ ‘interesting’ is represented by the values of a certain attribute a of theirs, for instance voxel luminosity. A common representation of such values is their *histogram* H_Φ , where the range of values is split into adjacent intervals of equal width—called *bins*—and for each bin, say k , $H_\Phi(k)$ is the total number of points that have a value of a falling in k . A point x is considered *similar* to the sample area if the histogram H_x (of an attribute of interest) of the points laying in a small area around x *correlates* with H_Φ . Of course, the two histograms must have the same number of bins.




Fig. 1: Illustration of the ImgQL TA operator Δ

More precisely, with reference to model $\mathcal{M} = ((X, \mathcal{C}), \mathcal{A}, \mathcal{V})$, the histogram $\mathcal{H}(a, Y, m, M, k)$ of the distribution of the values of attribute a of the points in $Y \subseteq X$, in the interval $[m, M]$ with step size Δ and k bins, is the function $\mathcal{H} : A \times 2^X \times \mathbb{R} \times \mathbb{R} \times \mathbb{N} \rightarrow (\mathbb{N} \rightarrow \mathbb{N})$ such that, for all values $m, M \in \mathbb{R}$, with $m < M$, and $k \in \mathbb{N} \setminus \{0\}$, and $i \in \{1, \dots, k\}$, $\mathcal{H}(a, Y, m, M, k)(i) = |\{y \in Y \mid (i-1)\Delta \leq \mathcal{A}(y, a) - m < i\Delta\}|$, where $\Delta = \frac{M-m}{k}$. The *mean* \bar{h} of histogram h is $\frac{1}{k} \sum_{i=1}^k h(i)$. Given histograms $h_1, h_2 : \{1, \dots, k\} \rightarrow \mathbb{N}$, their *cross correlation* $\mathbf{r}(h_1, h_2)$ is given by $\mathbf{r}(h_1, h_2) = \frac{\sum_{i=1}^k (h_1(i) - \bar{h}_1)(h_2(i) - \bar{h}_2)}{\sqrt{\sum_{i=1}^k (h_1(i) - \bar{h}_1)^2} \sqrt{\sum_{i=1}^k (h_2(i) - \bar{h}_2)^2}}$.

⁸ First order statistical methods are statistics based on the probability distribution function of the intensity values of the pixels of parts, or the whole, of an image.

The value of \mathbf{r} is *normalised* so that $-1 \leq \mathbf{r}(h_1, h_2) \leq 1$; $\mathbf{r}(h_1, h_2) = 1$ indicates that h_1 and h_2 are *perfectly correlated* (that is, $h_1 = \alpha h_2 + \beta$, with $\alpha > 0$); $\mathbf{r}(h_1, h_2) = -1$ indicates *perfect anti-correlation* (that is, $h_1 = \alpha h_2 + \beta$, with $\alpha < 0$). On the other hand, $\mathbf{r}(h_1, h_2) = 0$ indicates no correlation⁹.

We now have all the ingredients for completing the definition of the logical kernel of `ImgQL` by extending the syntax given in Def. 1 with $\Delta_{\bowtie c} \left[\begin{smallmatrix} m & M & k \\ r & a & b \end{smallmatrix} \right]$, and adding the following clause to the definition of the satisfaction relation, where $S(x, r) = \{y \in X \mid d(x, y) \leq r\}$ is the *sphere* of radius r centred in x , $h_a(i) = \mathcal{H}(a, S(x, r), m, M, k)(i)$, $h_b(i) = \mathcal{H}(b, \llbracket \Phi \rrbracket^{\mathcal{M}}, m, M, k)(i)$, and $\bowtie \in \{=, <, >, \leq, \geq\}$:

$$\mathcal{M}, x \models \Delta_{\bowtie c} \left[\begin{smallmatrix} m & M & k \\ r & a & b \end{smallmatrix} \right] \Phi \Leftrightarrow \mathbf{r}(h_a, h_b) \bowtie c.$$

So $\Delta_{\bowtie c} \left[\begin{smallmatrix} m & M & k \\ r & a & b \end{smallmatrix} \right] \Phi$ compares the region of the space constituted by the sphere of radius r centred in x against the region characterised by Φ . The comparison is based on the cross correlation of the histograms of the values of the chosen attributes of (the points of) the two regions, namely attribute a for the points around x and attribute b for the points that satisfy Φ . Of course, these attributes may also be chosen to be the same. Both histograms share the same domain $([m, M])$ and the same (number of) bins $(\{1, \dots, k\})$.

2.2 VoxLogicA

`VoxLogicA`¹⁰ is a model-checker for `ImgQL` that is specialised for digital image analysis. It is a *global* spatial model-checker in the sense that, given a model \mathcal{M} (i.e. a digital image) and a formula Φ , it computes the set $\llbracket \Phi \rrbracket^{\mathcal{M}}$ of all voxels in the image that satisfy Φ . Such a set can be, and usually is, represented by a boolean image—i.e. a closure model of the same dimension and size of \mathcal{M} , where each point is assigned the value *true* if the corresponding voxel in \mathcal{M} satisfies Φ , and *false* otherwise.

Actually, this feature is pushed forward in `VoxLogicA` so that one can obtain a (result) *quantitative* image where each point has a numerical value that is, for instance, the cross-correlation score computed for the verification of a Δ -formula on the corresponding voxel of \mathcal{M} . This is precisely what is done in the following example:

```
let scores = crossCorrelation(5,inty,inty,sample,min(inty),max(inty),15)
```

where `sample` is a formula characterising the sample portion of the image at hand and every point of `scores` will be associated with the score of the correlation

⁹ Note that normalisation makes the value of \mathbf{r} undefined for constant histograms, having therefore standard deviation of 0; in terms of statistics, a variable with such standard deviation is only (perfectly) correlated to itself. This special case is handled by letting $\mathbf{r}(h_1, h_2) = 1$ when both histograms are constant, and $\mathbf{r}(h_1, h_2) = 0$ when only one of the h_1 or h_2 is constant.

¹⁰ `VoxLogicA` is available at <https://github.com/vincenzoml/VoxLogicA>. In the present paper we used version 0.5.99.1-experimental.

between the intensity (`inty`) histogram of the sphere of radius 5 centred in the corresponding voxel of the image and intensity histogram of the `sample` area in the image, both histograms having 15 bins.

Functions and predicates can be defined in `VoxLogicA` in the usual way. For instance

```
let strongCorr(r,a,b,F,m,M,k,c) = crossCorrelation(r,a,b,F,m,M,k) > c
```

is the `VoxLogicA` equivalent of $\Delta_{>c} \begin{bmatrix} m & M & k \\ r & a & b \end{bmatrix} F$ so that

```
let interesting = strongCorr(5,inty,inty,sample,min(inty),max(inty),15,9.8)
```

returns in `interesting` a boolean image where the value *true* is associated to each voxel corresponding to a point in the current image which is the centre of a sphere of radius 5 the intensity of which has a high—higher than 9.8—correlation with the `sample` portion of the current image, and *false* to any other point.

The following additional commands are available in `VoxLogicA` (more details can be found in [6]):

- `load x = "s"` loads an image from file `"s"` and binds it to `x` for subsequent usage;
- `save "s" e` stores the image resulting from evaluation of expression `e` to file `"s"`;
- `print "s" e` prints to the log the string `s` followed by the numeric, or boolean, result of computing `e`;
- `import "s"` imports a library of declarations from file `"s"`;

3 Segmentation of Nevus with VoxLogicA

As mentioned briefly in the introduction, one of the complications of the segmentation of nevi is their great variability in appearance and the dishomogeneity of the dermoscopic images themselves. Nevi may show very different colour ranges, also within the same nevus, have different sizes, can be more or less regular, appear on more or less regular skin where hairs or sebaceous follicles may be present as well. Furthermore, the images themselves also show quite a variety and may be of different size, showing black corners, rings, or shadows due to the lenses used, showing more or less contrast and intensity or the presence of patches near the nevus. The images in Fig. 2 show a few examples of this dishomogeneity as encountered in the ISIC 2016 training dataset¹¹.

In the present paper we therefore make the following simplifying assumptions on the nevus images:

1. The nevus is surrounded by skin texture and is not touching or overlapping with the (black) border of the image;

¹¹ Datasets can be found at <https://challenge.isic-archive.com/data>




Fig. 2: Example images from the ISIC 2016 dataset illustrating the dishomogeneity of nevi. They also differ greatly in resolution, e.g. the size of (a) is 486 KB and that of (f) 11,3 MB (compared in .png format).

2. The nevus is connected and does not consist of multiple parts that are separated by skin texture;
3. The skin texture type surrounding the nevus is of type I or II (fair skin colour) or at most type III (uniform light tan) [20].

In future work we will study to what extent these assumptions can be relaxed.

3.1 Nevus Segmentation using Texture Analysis

Since there is very little one can take for granted in the images in the ISIC 2016 dataset, in the following we show some very coarse heuristics to get the process of segmentation started and illustrate the core steps of the segmentation approach. Some intermediate results are shown in Figure 3, illustrating the role of the texture analysis operator in this process on the image of a nevus shown in Fig. 3a (Fig. 3e shows the associated intensity distribution). The main aim is to distinguish voxels that are part of the background (skin) from those that are likely part of the nevus. First we assume that our task is to find all voxels that are likely to be part of the background, so the healthy skin surrounding the nevus. We assume furthermore that at least part of the nevus is somewhere in the middle of the image so that we can take an area relatively close to the border as a sample of the background. Let Φ be the ImgQL formula that specifies such an area, shown in Fig 3b as a semi-transparent overlay in cyan on the original image—later in the paper we will show Φ in detail. Let us, for now, also assume that we work with the intensity of the voxels rather than their colour or other attributes. As described in Section 2, we construct the histogram H_Φ of the distribution of the intensity values of all the voxels that satisfy Φ . Assume that H_Φ has k bins, and a minimum and maximum value that correspond to the minimum and maximum pixel intensity in the whole image.

Next we construct the local histogram H_x for each pixel x in the image by taking the intensity of all the pixels that are present in a radius rad around pixel x . This second histogram has the same number of bins and minimal and maximal values as those of histogram H_ϕ . We can now compute the similarity score for each point x in the image by computing the Pearson’s correlation coefficient of the histograms H_ϕ and H_x . Recall that this provides normalised values between -1 and +1. A value equal to 1 indicates perfect correlation between the histograms, a value equal to -1 indicates perfect anti-correlation. A score of value 0 indicates that there is no correlation between the histograms. The result for Fig. 3a is shown in Fig. 3c as a semi-transparent yellow overlay where higher values of the score correspond to a brighter yellow hue. The associated histogram of these cross-correlation scores are shown in Fig. 3f. Finally, in Fig. 3d, those pixels with a cross-correlation score above 0.05 are shown as an overlay in pink.

Thus, this particular use of the texture operator can provide a rather good first approximation of the area covered by the nevus. Clearly, it is not perfect yet, as also some other areas remain that are not identified as part of the background, whereas they should be. But these areas can in principle be identified by other means, such as their relative position with respect to the border of the image and other aspects that distinguish them from the nevus itself, as done in Specification 1, shown later on, where we show in more detail how the above described procedure can be specified in `ImgQL`. This specification uses a predicate, `patch` and a derived operator `relDist`. The former is a predicate specifying voxels that are part of a patch. The latter is a derived operator that defines the *relative distance* in an image, depending on its size. The definition of both the operators are provided and explained after Specification 1.

Moreover, in Section 4 we will use common similarity indexes to compare the quality of the segmentation. These indexes are defined directly in terms of `ImgQL` and shown in Specification 5. They provide numeric support in the form of values of several commonly used similarity indexes that allow for an objective comparison with expert ground truth and with the performance of other state-of-the-art approaches. Note that the `VoxLogicA` procedure provided in the sequel does not require any particular preprocessing of the images as provided by the ISIC 2016 dataset, except for format conversion for the test data (from `jpg` format to `png` format) and colour conversion for the ground truth data (from grayscale to `RGB`). So we do not use explicit image pre-processing transformations that for example remove hairs or black corners or borders of the image or that increase contrast or normalise the size of the images.

Brief description of the VoxLogicA segmentation Specification 1. After importing the standard library (`stdlib.imgql`), containing derived `VoxLogicA` operator definitions, and loading the image with ground truth and the related nevus image (lines 1-3), two quantitative images are defined: `nevusImgIntens` for the nevus image and `groundIntens` for the ground truth image (lines 4-5). These quantitative images associate to each voxel its intensity (luminosity); intuitively, each of them can be thought of as a (2D) matrix.




Fig. 3: Illustration texture analysis operator *VoxLogicA*. Fig. 3a Original image of nevus (image ISIC-0000010 of the 2016 Challenge) and surrounding skin and related histogram of voxel intensity in Fig. 3e. Sample Φ of background voxels shown in cyan (Fig. 3b). Cross-correlation score for surrounding of each voxel w.r.t. histogram of background sample Φ . The score is shown in yellow. The higher the score the higher the intensity of the yellow colour of voxels in Fig. 3c. The distribution of the cross-correlation scores is shown in Fig. 3f. Voxels with a cross-correlation score of more than 0.05 are shown in pink in Fig. 3d.

ImgQL Specification 1: Nevus segmentation specification

```

1 import "stdlib.imgql"
2 load groundTruth = "$INPUTDIR/$NAME_seg_RGB.png"
3 load nevus = "$INPUTDIR/$NAME.png"
4 let nevusImgIntens = intensity(nevus)
5 let groundIntens = intensity(groundTruth)
6 let similarTo(a,rad) = crossCorrelation(rad,nevusImgIntens,
    nevusImgIntens,a,min(nevusImgIntens),max(nevusImgIntens),15)
7 let almostBlack = nevusImgIntens <. 40.0
8 let blackBorder = grow(distleq(relDist(5),border), almostBlack)
9 let bgSampleWidth = relDist(200)
10 let bgSample = distleq(bgSampleWidth, blackBorder) & (!blackBorder) &
    (!patch)
11 let bgSimScore = similarTo(bgSample, relDist(5))
12 let bgSim = (bgSimScore >. 0.05) & (! patch) & (!blackBorder)
13 let preSeg = ((!border) S (bgSimScore <. 0.11)) & (! patch) &
    (!blackBorder)
14 let nevSeg = smoothen(maxvol(preSeg & (! (distleq(relDist(50),
    blackBorder))))), relDist(3))
15 let nevSegSmooth = smoothen(maxvol(preSeg & (! (distleq(relDist(50),
    blackBorder))))), relDist(10))
16 let nevSeg1 = maxvol(nevSeg & nevSegSmooth & !patch)
17 let nevSegV0 = nevSeg1

```




Fig. 4: Illustration of segmentation procedure in Specification 1 of image ISIC_0008294 (also shown in Fig 2 (e)). Figure (a) shows the nevus intensities (greyscale); Figures (b) to (h) are each associated with a specific formula, as indicated below, with the exception of (d) where the score is shown as a varying intensity of yellow: (b) `blackBorder` (red); (c) `bgSample` (blue); (d) `bgSimScore` (yellow); (e) `bgSim` (green); (f) `preSeg` (magenta); (g) `nevSegV0` (cyan); (h) `nevSegV0` (cyan) and ground truth (blue).

In line 6 a similarity operator is defined with parameters `a` and `rad`, that is based on the cross correlation operator; `a` defines the sample area (denoted by Φ in Section 2) and `rad` defines the radius around each voxel x for the construction of the local histogram of x . We recall that the parameters of the cross correlation operator are, from left to right, the radius, the attribute of the voxels around x , the attribute of voxels of the sample area, the minimum and maximum value of the intensity of the whole nevus image, and finally the number of bins in both histograms.

The (intermediate) results of the segmentation procedure defined in the rest of the specification are illustrated in Fig. 4. Lines 7-8 specify the characteristics of voxels that are part of the black corners that can be observed in many images (in a similar way as shown in Fig. 2a and Fig 2b). Such voxels should not be considered in the sample of the skin texture. In line 7 voxels are specified that are almost black, i.e. having an intensity below 40. Then (line 8) only those almost black voxels are considered from which the border can be reached exclusively ‘passing by’ further almost black voxels, exploiting the `grow` operator.

In lines 9-10 a sample (`bgSample`) of the skin around the nevus is specified, namely a sample of voxels that are most likely part of the healthy skin without (or with very few) voxels that are part of the nevus. This sample consists of voxels that are at most at relative distance 200 (`bgSampleWidth`) from the black border (lines 8-9). In line 11 the similarity score of each voxel in the image w.r.t. the sample is computed using the `similarTo` operator and saved as a quantitative

image. Line 12 characterises all voxels that have a cross correlation score larger than 0.05. This line has been inserted only for illustration purposes here to highlight the voxels with skin texture; in Figure 4e we show this set of points (bgSim).

A preliminary segmentation is specified in line 13, where we look for voxels that are not part of the border and that are surrounded by voxels with a cross correlation score of less than 0.11, a small correlation score. The idea is that, at the border of the nevus and the healthy skin, the histograms of the area around those voxels represent in part the skin and in part the nevus, which have in general rather different intensity distributions. The cross correlation of such histograms with the sample area of the skin can therefore be expected to be quite small. The exact value of the threshold has been established in an empirical way; it is the value that gives on average good results for the investigated datasets. For optimal results on individual images this threshold value may differ slightly.

Of course, this pre-segmentation should exclude areas close to the black border and in patches. The latter are used in some images to indicate the position of nevus with little contrast (see for example Fig. 2e and Fig. 2f). There may also be other small darker areas on the skin that are not part of the nevus. Therefore, in line 14, the preliminary segmentation is refined by taking only the largest volume (maxvol) far enough from the black border and smoothing the specified area removing small noise and irregularities at the edge.

In line 15 the same procedure of line 14 is repeated with a larger smoothing factor. This is used to exclude possible thin protrusions attached to the segmented nevus that are originating from thin hairs or shadows. In line 16 the intersection of these intermediate results is taken to preserve the more detailed edge of the nevus and at the same time to exclude some larger protrusions (i.e. several hairs grouped together). Also voxels that are part of possible patches are excluded from the segmentation and we obtain a first nevus segmentation in line 17 (nevSegV0).

ImgQL Specification 2: Generating model checking results and similarity scores

```

1 import let groundTruth = groundIntens >. 0
2 save "$OUTPUTDIR/$NAME_nevSegV0.png" nevSegV0
3 save "$OUTPUTDIR/nevSegV0.nii.gz" nevSegV0
4 print "DICE V0" dice(nevSegV0,groundTruth)

```

In Specification 2 the manual segmentation performed by domain experts (the ‘ground truth’) is defined as a predicate that is satisfied by voxels in the image of the ground truth where the intensity of the voxel is positive (line 1). In fact, the ground truth is a black and white image of the same size as the image of the nevus where the area indicated by the expert is white (intensity 255) and the rest black. The resulting segmentation (but also other intermediate results as those shown in Fig 4) can be saved in .png format or in the NIfTI (.nii) format.

ImgQL Specification 3: Relative distances

```

1 let refImgPerimeter = 2 .* (1022 .+. 767)
2 let imgSizeFactor = (volume(border) ./ refImgPerimeter)
3 let relDist(x) = (imgSizeFactor .* x)

```

The latter format is used by various viewers used in medical imaging. We used the free viewer MRICron¹². The operator `dice` compares (line 4) the segmentation defined by `nevSegV0` with the `groundTruth` giving as result a similarity score as defined in Specification 5. Further details on these scores are provided in the next section.

Specification 1 uses the `relDist` operator defined in Specification 3. The ISIC 2016 datasets contain images of very different sizes. The `relDist` operator is introduced to scale the distance appropriately, with respect to a reference image. The size of the reference image is defined as the length of its perimeter, i.e. the number of voxels on its border. The reference image has a width of 1022 voxels and a height of 767 voxels. The perimeter of the image being analysed can be found as the volume (i.e. number of voxels) that form the border (i.e. one voxel wide edge) of the image. The property `border` is a built-in operator of `ImgQL`. The scaling of the distance is the fraction between the length of the perimeter of the image under analysis and that of the reference image.

ImgQL Specification 4: Patches

```

1 let bNev = blue(nevus)
2 let rNev = red(nevus)
3 let gNev = green(nevus)
4 let patchBlue = distleq(relDist(5),(bNev > (rNev +. 30)) &
5 (bNev > (gNev)) & (bNev >. 150))
6 let patchRed = distleq(relDist(5),(rNev > (bNev +. 100)) &
7 (rNev > (gNev +. 20))) & (rNev >. 130)
8 let patchGreen = distleq(relDist(5),(gNev > (rNev +. 20)) &
9 (gNev > bNev) & (gNev >. 100))
10 let patchPart(x,y) = ifB(volume(x) <. (y .* volume(tt)),x,ff)
11 let patchSample = patchPart(patchBlue,0.4) | patchPart(patchRed,0.4) |
    patchPart(patchGreen,0.4)
12 let patchAtBorder = touch(smoothen(patchSample,relDist(10)),
    distleq(relDist(20),border))
13 let patch = ifB(ppM(patchAtBorder) >. 0.5, patchAtBorder, ff)

```

Specification 1 also uses the predicate `patch` that is satisfied by voxels that are part of a patch. Patches are defined in Specification 4. Lines 1-3 define three quantitative images (matrixes) projecting the intensity of the blue, red and green

¹² <https://www.nitrc.org/projects/mricron>

part of the `rgb`-vector for each voxel of the image. Lines 4-9 define blue, red and green patches, respectively. These also cover intermediate hues such as yellow and orange. However, it is not enough to define the colour ranges of patches because nevi or skin may have occasionally colours in those ranges too (see for example Fig. 2a and Fig. 2b). Using further knowledge about the relative spatial position of patches (they are at the border of the image), their relative size (covering not more than 40 percent of an image) and their compactness (their Polsby-Popper measure, of compactness of a shape¹³, `ppM`, is at least 0.5), the specification `patch` is given in line 13. `ifB` is the boolean if-then-else construct of `VoxLogicA`. The definition of `ppM` is shown in Specification 5 (lines 6-8) in the next section.

4 Results

In the literature on medical imaging several indexes are used to compare the similarity between two segmentations of the same image, in particular the similarity between the manual and automatic segmentation. Commonly used similarity measures are the Dice index, the Jaccard index and the accuracy index. These coefficients give a result between 0 (no similarity) and 1 (perfect similarity). Further measures are the sensitivity (fraction of true positives) and specificity (fraction of true negatives). For example a Dice index of around 0.9 is considered as indicating very good similarity. The Dice index (D) and the Jaccard index (J) are related: $J = D / (2 - D)$. In Specification 5 these common similarity indexes are defined in `ImgQL` so that they can be calculated during the analysis. Their definitions should be self-explanatory recalling that the operator `volume(x)` gives the number of voxels that satisfy property x . It must be noted though, that no unique ‘gold standard’ for comparison exists because also manual expert markings have a considerable level of variability. In [17] it is reported that the average Jaccard index of agreement between 3 pairs of clinicians that each segmented a subset of 100 images was 0.786.

A large set of dermoscopic images is available from the ISIC gallery¹⁴. Furthermore, two datasets of dermoscopic images were made available for the ISIC 2016 Challenge¹⁵. One set of 900 images for *training* purposes and a *test* set of 379 images for the final evaluation of the competing methods.

In Table 1 we show the mean scores for `nevSegV0` of Specification 1 and compare them with the best mean scores¹⁶ of the teams that participated in the ISIC 2016 Challenge Part 1 on nevus segmentation [27]. The prevalent techniques used by the teams in the ISIC 2016 Challenge were based on deep learning [17].

For the comparison in Table 1 we have used two restricted subsets of images from the ISIC gallery that satisfy the criteria that were mentioned in Section 3. *Mean first 10* concerns the first 10 images of the ISIC gallery (images

¹³ Also known as “Isoperimetric quotient”.

¹⁴ See: <https://www.isic-archive.com/#!/topWithHeader/onlyHeaderTop/gallery>

¹⁵ These datasets can be found at <https://challenge.isic-archive.com/data>

¹⁶ These scores were obtained on the ISIC 2016 Challenge *test dataset*.

	Accuracy	Dice	Jaccard	Sensitivity	Specificity
Best ISIC 2016	0,953	0,91	0,843	0,91	0,965
V0: Mean first 10	0,969	0,92	0,855	0,96	0,971
V0: Mean first 50	0,945	0,90	0,828	0,89	0,980

Table 1: Similarity scores of the nevus segmentation method `nevSegV0` shown in specification 1 compared with the best score on the test set of the ISIC 2016 Challenge [17].

ISIC_0000000 to ISIC_0000010) with the exception of image ISIC_0000004 in which the nevus is touching the border. These 10 images have been used to develop the specification and to manually calibrate some thresholds. They are also part of the datasets for the ISIC 2016 Challenge, with the exception of image ISIC_0000005. Furthermore, some other typical example images from the ISIC 2016 training set have been used, such as those shown in Fig. 2, to develop the part of the specification dealing with patches and some other aspects.

Subsequently we have analysed a larger set of images. *Mean first 50* concerns the first 50 images of the ISIC 2016 gallery (ISIC_0000000 to ISIC_0000050) omitting¹⁷ images 4, 11, 24, 26, 31, 33 and 50. Also these images are part of the datasets of the ISIC 2016 Challenge. Table 1 shows the various scores for `nevSegV0` for these restricted subsets. In the first line of the table the best scores for the ISIC 2016 Challenge are reported for the complete ISIC 2016 test set (see [36,38]). Although the latter set is different and much larger, the scores we obtained for `nevSegV0` for the small sets are in line with these other scores. This means that it is possible to obtain segmentations of good quality with our method.

Next we investigate to what extent this is the case considering the much larger ISIC 2016 training set. In Table 2 we consider the whole ISIC 2016 training set except those images where the nevus is touching the black border, so in total we consider 750 of the 900 training set images. The table shows the number and fraction of images for which `nevSegV0` gives a Dice score above (or below) a certain threshold. This table shows that this rather simple specification gives an excellent score (Dice > 0.9) for more than 30% of the images. Moreover, in more than 65% of the cases it gives a reasonably good correspondence (Dice > 0.7). These results are rather surprising because of the enormous variability of the images and the relative simplicity of this version of the segmentation procedure. Fig. 5 shows some more examples of segmentation via `nevSegV0`.

The illustration of the intermediate results of the segmentation procedure in Fig. 4 shows that with this spatial model-checking approach each step is ‘explainable’, meaning that one can understand why a particular result is produced.

¹⁷ Images 4, 26 and 33 violate criterium (1), images 31 and 50 violate criterium (2).

For image 11, due to a technical issue, we had the wrong ground truth (namely that of image 00), image 24 has too low contrast.

ImgQL Specification 5: Similarity indexes and the Polsby-Popper measure

```

1 let dice(x,y) = (2 .* volume(x & y)) ./ (volume(x) .+. volume(y))
2 let jaccard(x,y) = dice(x,y) ./ (2 .-. dice(x,y))
3 let sensitivity(x,y) = volume(x & y) ./ (volume(x & y) .+.
  volume(!x & (y)))
4 let specificity(x,y) = volume(!x & (!y)) ./ (volume(!x & (!y))
  .+. volume(x & (!y)))
5 let accuracy(x,y) = (volume(x & y) .+. volume(!x & (!y))) ./
  (volume(x & y) .+. volume(!x & (!y)) .+. volume(x & (!y)) .+.
  volume(!x & (y)))
6 let square(x) = x .* x
7 let iboundary(x) = near(interior(x)) & !(interior(x))
8 let ppm(x) = (volume(x) .* 4 .* 3.14) ./
  (square(volume(iboundary(x))))

```

This is one of the advantages of our spatial model-checking approach. It makes it much easier to further improve the segmentation procedure, to discuss a particular method with experts and to compare different methods or use them to document a particular analysis. Furthermore, the resulting segmentations are amenable to further analysis of the nevus area itself. For example one may extract further features such as regularity of its shape or texture, size and other features that may be of interest for the diagnosis of nevi. Note also that the segmentation method can be calibrated to specific images to obtain more precise results by tuning the thresholds. We have not applied our method to the ISIC 2016 *test* set; we keep that set for testing future versions of the specification. The results in this section have been obtained with **VoxLogicA** version 0.6.0.osx-x64 on a MacBook Pro with 3.1 GHz Intel core i7 and 16GB of memory. For the large




Fig. 5: Images and their segmentation (cyan) and ground truth (blue): ISIC_0000002 (a) resp. (d), ISIC_0000043 (b) resp. (e) and ISIC_0004309 (c) resp. (f).

nevSegV0	number	fraction (of 750)
Dice > 0.9	250	0.33
Dice > 0.8	396	0.53
Dice > 0.7	491	0.65
Dice < 0.5	140	0.19
Dice = 0	45	0.06

Table 2: Dice score distribution of the nevus segmentation method `nevSegV0` for 750 images of the ISIC 2016 training set not touching the black border.

data sets an AMD Ryzen 7 2700 Eight-Core Processor with 32GB of memory has been used. `nevSegV0` takes approximately 2 seconds for a typical image of 1325 KB. The results together with the selected datasets discussed above are available in a git-repository¹⁸.

5 Conclusions and Future Work

We have shown how spatial model checking techniques and the related tool `VoxLogicA` can be used for the segmentation of `nevi`. Nevus segmentation based on dermoscopic images is an important part of many automatic procedures to diagnose malign skin tumours such as Melanoma. To the best of our knowledge, this is the first time that spatial model checking is applied to this specific domain. Spatial model-checkers use high-level, often domain oriented, logical languages to specify spatial properties. In this paper we have presented a segmentation method combining spatial operators inspired by the notion of closure spaces with more domain oriented operators such as a texture similarity operator. This first, rather simple method shows that an accuracy can be obtained that is in line with the state-of-the-art in nevus segmentation, i.e. a Dice score above 0.9. It has also been shown that it obtains this high accuracy in more than 30% of 750 images of one of the largest available training sets of dermoscopic images that is publicly available. An advantage of this spatial model-checking method is that the segmentation procedure is explainable and high-level. This makes the method amenable to further improvements by inspection of the intermediate results, exchange and discussion of the method specifications between domain experts, conservation of the method for the purpose of documentation of the analysis and independent replication by other experts.

The results we obtained so far are very promising and future work is envisioned to increase the class of images for which accurate segmentation can be obtained in a similar spirit as in which we have shown how one can deal with the presence of patches or the presence of other artifacts in the images that are due to the way the images have been produced. The enormous dishomogeneity in this type of images, both for what concerns the `nevi` and the images themselves, remains a great challenge.

¹⁸ <https://github.com/brocciagi/Spatial-Model-Checking-for-Nevus-Segmentation>

References

1. Aiello, M., Pratt-Hartmann, I., Benthem, van, J.: Handbook of Spatial Logics. Springer (2007)
2. Banci Buonamici, F., Belmonte, G., Ciancia, V., Latella, D., Massink, M.: Spatial logics and model checking for medical imaging. *Int. J. Softw. Tools Technol. Transf.* 22(2), 195–217 (2020), <https://doi.org/10.1007/s10009-019-00511-9>
3. Belmonte, G., Ciancia, V., Latella, D., Massink, M.: VoxLogicA: a Spatial Model Checker for Declarative Image Analysis (Extended Version). *ArXiv e-prints* (Nov 2018), <https://arxiv.org/abs/1811.05677>
4. Belmonte, G., Ciancia, V., Latella, D., Massink, M., Biondi, M., De Otto, G., Nardone, V., Rubino, G., Vanzi, E., Banci Buonamici, F.: A topological method for automatic segmentation of glioblastoma in MR FLAIR for radiotherapy - ESMRMB 2017, 34th annual scientific meeting. *Magnetic Resonance Materials in Physics, Biology and Medicine* 30(S1), 437 (oct 2017), <https://doi.org/10.1007/s10334-017-0634-z>
5. Belmonte, G., Ciancia, V., Latella, D., Massink, M.: Innovating medical image analysis via spatial logics. In: ter Beek, M.H., Fantechi, A., Semini, L. (eds.) *From Software Engineering to Formal Methods and Tools, and Back - Essays Dedicated to Stefania Gnesi on the Occasion of Her 65th Birthday*. *Lecture Notes in Computer Science*, vol. 11865, pp. 85–109. Springer (2019), https://doi.org/10.1007/978-3-030-30985-5_7
6. Belmonte, G., Ciancia, V., Latella, D., Massink, M.: Voxlogica: A spatial model checker for declarative image analysis. In: Vojnar, T., Zhang, L. (eds.) *Tools and Algorithms for the Construction and Analysis of Systems - 25th International Conference, TACAS 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings, Part I*. *Lecture Notes in Computer Science*, vol. 11427, pp. 281–298. Springer (2019), https://doi.org/10.1007/978-3-030-17462-0_16
7. Brock, K.: *Image processing in radiation therapy*. CRC Press (2014)
8. Castellano, G., Bonilha, L., Li, L., Cendes, F.: Texture analysis of medical images. *Clinical Radiology* 59(12), 1061–1069 (dec 2004)
9. Chen, C., Da Ponte, J., Fox, M.: Fractal feature analysis and classification in medical imaging. *IEEE Transactions on Medical Imaging* 8(2), 133–142 (jun 1989)
10. Chetelat, G., Baron, J.: Early diagnosis of alzheimer’s disease: contribution of structural neuroimaging. *NeuroImage* 18(2), 525–541 (2003)
11. Ciancia, V., Gilmore, S., Latella, D., Loreti, M., Massink, M.: Data verification for collective adaptive systems: Spatial model-checking of vehicle location data. In: *Eighth IEEE International Conference on Self-Adaptive and Self-Organizing Systems Workshops, SASOW*. pp. 32–37. IEEE Computer Society (2014)
12. Ciancia, V., Grilletti, G., Latella, D., Loreti, M., Massink, M.: An experimental spatio-temporal model checker. In: *Software Engineering and Formal Methods - SEFM 2015 Collocated Workshops*. *Lecture Notes in Computer Science*, vol. 9509, pp. 297–311. Springer (2015)
13. Ciancia, V., Latella, D., Loreti, M., Massink, M.: Specifying and verifying properties of space. In: *Theoretical Computer Science - 8th IFIP TC 1/WG 2.2 International Conference, TCS 2014, Rome, Italy, September 1-3, 2014*. *Proceedings*. *Lecture Notes in Computer Science*, vol. 8705, pp. 222–235. Springer (2014)
14. Ciancia, V., Latella, D., Loreti, M., Massink, M.: Model Checking Spatial Logics for Closure Spaces. *Logical Methods in Computer Science Volume 12, Issue 4* (Oct 2016), <http://lmcs.episciences.org/2067>

15. Ciancia, V., Latella, D., Massink, M., Paškauskas, R.: Exploring spatio-temporal properties of bike-sharing systems. In: 2015 IEEE International Conference on Self-Adaptive and Self-Organizing Systems Workshops, SASO Workshops 2015, Cambridge, MA, USA, September 21-25, 2015. pp. 74–79. IEEE Computer Society (2015), <https://doi.org/10.1109/SASOW.2015.17>
16. Ciancia, V., Latella, D., Massink, M., Paškauskas, R., Vandin, A.: A tool-chain for statistical spatio-temporal model checking of bike sharing systems. In: Margaria, T., Steffen, B. (eds.) Leveraging Applications of Formal Methods, Verification and Validation: Foundational Techniques - 7th International Symposium, ISoLA 2016, Imperial, Corfu, Greece, October 10-14, 2016, Proceedings, Part I. Lecture Notes in Computer Science, vol. 9952, pp. 657–673 (2016), https://doi.org/10.1007/978-3-319-47166-2_46
17. Codella, N.C.F., Nguyen, Q.B., Pankanti, S., Gutman, D., Helba, B., Halpern, A., Smith, J.R.: Deep learning ensembles for melanoma recognition in dermoscopy images. *IBM Journal of Research and Development* 61(45) (2017), <http://www.research.ibm.com/journal/>, special Issue on Deep Learning
18. Davnall, F., Yip, C., Ljungqvist, G., Selmi, M., Ng, F., Sanghera, B., Ganeshan, B., Miles, K.A., Cook, G.J., Goh, V.: Assessment of tumor heterogeneity: an emerging imaging tool for clinical practice? *Insights into Imaging* 3(6), 573–589 (oct 2012)
19. Doi, K.: Computer-aided diagnosis in medical imaging: Historical review, current status and future potential. *Comput. Med. Imaging Graph.* 31(4-5), 198–211 (2007)
20. Fitzpatrick, T.B.: The Validity and Practicality of Sun-Reactive Skin Types I Through VI. *Archives of Dermatology* 124(6), 869–871 (06 1988), <https://doi.org/10.1001/archderm.1988.01670060015008>
21. Forsea, A., del Marmol, V., de Vries, E., Bailey, E., Geller, A.: Melanoma incidence and mortality in europe: new estimates, persistent disparities. *British Journal of Dermatology* 167(5), 1124–1130 (2012), <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1365-2133.2012.11125.x>
22. Galton, A.: The mereotopology of discrete space. In: Freksa, C., David, M. (eds.) *Spatial Information Theory. Cognitive and Computational Foundations of Geographic Information Science*, Lecture Notes in Computer Science, vol. 1661, pp. 251–266. Springer Berlin Heidelberg (1999), http://dx.doi.org/10.1007/3-540-48384-5_17
23. Galton, A.: A generalized topological view of motion in discrete space. *Theor. Comput. Sci.* 305(1-3), 111–134 (2003), [https://doi.org/10.1016/S0304-3975\(02\)00701-6](https://doi.org/10.1016/S0304-3975(02)00701-6)
24. Galton, A.: Discrete mereotopology. In: Calosi, C., Graziani, P. (eds.) *Mereology and the Sciences: Parts and Wholes in the Contemporary Scientific Context*, pp. 293–321. Springer International Publishing (2014), https://doi.org/10.1007/978-3-319-05356-1_11
25. Gordillo, N., Montseny, E., Sobrevilla, E.: State of the art survey on MRI brain tumor segmentation. *Magn. Reson. Imaging.* 31(8), 1426–1438 (2013)
26. Grosu, R., Smolka, S., Corradini, F., Wasilewska, A., Entcheva, E., Bartocci, E.: Learning and detecting emergent behavior in networks of cardiac myocytes. *Commun. ACM* 52(3), 97–105 (2009)
27. Gutman, D., Codella, N.C.F., Celebi, M.E., Helba, B., Marchetti, M.A., Mishra, N.K., Halpern, A.: Skin lesion analysis toward melanoma detection: A challenge at the international symposium on biomedical imaging (ISBI) 2016, hosted by the international skin imaging collaboration (ISIC). *CoRR abs/1605.01397* (2016), <http://arxiv.org/abs/1605.01397>

28. Haghghi, L., Jones, A., Kong, Z., Bartocci, E., Grosu, R., Belta, C.: Spatel: A novel spatial-temporal logic and its applications to networked systems. In: Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control. pp. 189–198. HSCC '15, ACM, New York, NY, USA (2015)
29. Han, F., Wang, H., Zhang, G., Han, H., Song, B., Li, L., Moore, W., Lu, H., Zhao, H., Liang, Z.: Texture feature analysis for computer-aided diagnosis on pulmonary nodules. *Journal of Digital Imaging* 28(1), 99–115 (aug 2014)
30. Heinonen, T., Arola, T., Kalliokoski, A., Dastidar, P., Rossi, M., Soimakallio, S., Hyttinen, J., Eskola, H.: Computer aided diagnosis tool for the segmentation and texture analysis of medical images. In: IFMBE Proceedings, pp. 274–276. Springer Science (2009)
31. Kassner, A., Thornhill, R.: Texture analysis: A review of neurologic MR imaging applications. *Am. J. Neuroradiol.* 31(5), 809–816 (2010)
32. Lopes, R., Ayache, A., Makni, N., Puech, P., Villers, A., Mordon, S., Betrouni, N.: Prostate cancer characterization on MR images using fractal features. *Med. Phys.* 38(1), 83 (2011)
33. Nenzi, L., Bortolussi, L., Ciancia, V., Loreti, M., Massink, M.: Qualitative and Quantitative Monitoring of Spatio-Temporal Properties with SSTL. *Logical Methods in Computer Science* 14(4), 1–38 (2018), DOI 10.23638/LMCS-14(4:2)2018. Published on line: 23 Oct. 2018. ISSN: 1860-5974
34. Randell, D.A., Landini, G., Galton, A.: Discrete mereotopology for spatial reasoning in automated histological image analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* 35(3), 568–581 (2013), <https://doi.org/10.1109/TPAMI.2012.128>
35. Rodriguez Gutierrez, D., Awwad, A., Meijer, L., Manita, M., Jaspan, T., Dineen, R., Grundy, R., Auer, D.: Metrics and textural features of MRI diffusion to improve classification of pediatric posterior fossa tumors. *American Journal of Neuroradiology* 35(5), 1009–1015 (dec 2013)
36. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. *CoRR abs/1505.04597* (2015), <http://arxiv.org/abs/1505.04597>
37. Sharma, N., Ray, A., Sharma, S., Shukla, K., Pradhan, S., Aggarwal, L.: Segmentation and classification of medical images using texture-primitive features: Application of BAM-type artificial neural network. *J Med Phys* 33(3), 119 (2008)
38. Shelhamer, E., Long, J., Darrell, T.: Fully convolutional networks for semantic segmentation. *IEEE Transactions on Pattern Analysis & Machine Intelligence* 39(04), 640–651 (apr 2017)
39. Srinivasan, G., Shobha, G.: Statistical texture analysis. In: Proceedings of World Academy of Science, Engineering and Technology. vol. 36, pp. 1264–1269 (dec 2012)
40. Tijms, B., Series, P., Willshaw, D., Lawrie, S.: Similarity-based extraction of individual networks from gray matter MRI scans. *Cerebral Cortex* 22(7), 1530–1541 (aug 2011)
41. Toosy, A.: Diffusion tensor imaging detects corticospinal tract involvement at multiple levels in amyotrophic lateral sclerosis. *J. Neurol. Neurosurg. Psychiatry* 74(9), 1250–1257 (2003)
42. Woods, B., Clymer, B., Kurc, T., Heverhagen, J., Stevens, R., A., O., Bulan, O., Knopp, M.: Malignant-lesion segmentation using 4d co-occurrence texture analysis applied to dynamic contrast-enhanced magnetic resonance breast image data. *J. Magn. Reson. Imaging* 25(3), 495–501 (2007)