



Skeleton Sparsification and Densification Scale-Spaces

Julia Gierke¹ · Pascal Peter²

Received: 9 September 2025 / Accepted: 8 June 2026
© The Author(s) 2026

Abstract

The Hamilton–Jacobi skeleton, also known as the medial axis, is a powerful shape descriptor that represents binary objects in terms of the centres of maximal inscribed discs. Despite its broad applicability, the medial axis suffers from sensitivity to noise: Minor boundary variations can lead to disproportionately large and undesirable expansions of the skeleton. Classical pruning methods mitigate this shortcoming by systematically removing extraneous skeletal branches. This sequential simplification of skeletons resembles the principle of sparsification scale-spaces that embed images into a family of reconstructions from increasingly sparse pixel representations. We combine both worlds by introducing skeletonisation scale-spaces: They leverage sparsification of the medial axis to achieve hierarchical simplification of shapes. Unlike conventional pruning, our framework inherently satisfies key scale-space properties, such as hierarchical architecture, controllable simplification, and equivariance to geometric transformations. We provide a rigorous theoretical foundation in both continuous and discrete formulations and extend the concept further with densification. By growing the skeleton successively instead of shrinking it, we allow inverse progression from coarse to fine scales. Densification scale-spaces can even reach beyond the original skeleton to produce overcomplete shape representations with relevancy for practical applications. Through proof-of-concept experiments, we demonstrate the effectiveness of our framework for practical tasks including robust skeletonisation, shape compression, and stiffness enhancement for additive manufacturing.

Keywords Skeleton · Medial axis · Scale-space · Sparsification · Densification

1 Introduction

Classical scale-spaces embed images into a family of hierarchically simplified versions of the original. Such image evolutions are typically based on partial differential equations (PDEs) [1–5] which remove information by smoothing, thus simplifying the image. Cárdenas et al. [6] have introduced sparsification as an alternative: They create coarser scales by explicitly removing image pixels. At each scale, they interpolate the missing image parts to create a sequence of complete images. On one hand, these scale-spaces share

many important properties with classical scale-spaces. On the other hand, they also open up new applications for scale-space theory such as image compression.

In the field of shape analysis [7], notions of hierarchical simplification that resemble sparsification are well-known in a different context. Shapes can be equivalently described by the medial axis introduced by Blum [8] as a symmetry-based, central shape descriptor intended to mimic biological sensory processes. Due to its appearance as a thin set of connected lines and arcs centred in the shape, the medial axis is also referred to as the skeleton. Unfortunately, it is highly sensitive to perturbations on the shape boundary, which lead to unwanted branches in the skeleton. Therefore, so-called pruning techniques [9–14] have been proposed. They sequentially remove parts of the skeleton. The similarity to the sparsification processes of Cárdenas et al. [6] suggests that there might be hitherto unexplored connections to scale-space theory.

This is surprising, since shape analysis has been known to be closely connected to scale-space theory. Significant research efforts have been dedicated to the exploration of

Julia Gierke and Pascal Peter have equally contributed.

✉ Pascal Peter
peter@math.uni-sb.de

Julia Gierke
gierke@mia.uni-saarland.de

¹ Mathematical Image Analysis Group, Saarland University, Campus E1.7, 66041 Saarbrücken, Germany

² Department of Mathematics and Computer Science, Saarland University, Campus E2.4, 66041 Saarbrücken, Germany

morphological scale-spaces [1, 15–19]. Here, simplifying shape evolutions are derived from the basic erosion and dilation operators [20]. Thus, in contrast to the skeleton view point, they operate on the shape outline instead of a central shape descriptor. While multi-scale ideas have been previously connected to the medial axis transform [21–26], these focus on counteracting noise or on generalisations of the shape descriptor. In contrast, we aim to build a new scale-space theory on the combined concepts of skeletonisation and sparsification.

1.1 Our Contribution

We propose a scale-space framework that combines the ideas of image sparsification [6] and multi-scale shape analysis: Our *shape sparsification* transfers classical scale-space properties [1] such as hierarchical architecture, quantifiable simplification, and equivariances to the medial axis setting [8]. To this end, we consider sparsified skeletons instead of sparsified images and derive theoretical guarantees that apply to shape descriptors and their associated binary images. Furthermore, we complement the sparsification idea by a densification counterpart: By extending the original skeleton, we can reach finer skeleton scales that go beyond the initial scale corresponding to the original shape. With proof-of-concept examples, we illustrate how our theory supports a variety of different practical applications ranging from compression over skeleton robustification to pre-processing for 3-D printing.

Beyond a more rigorous discussion of theoretical background and a more detailed review of related work, we extend our previous conference contribution [27] on skeletonisation scale-spaces in the following ways:

1. We propose a fully continuous scale-space theory and discuss differences between continuous, discrete, and semi-discrete skeletonisation scale-spaces in more detail.
2. With extended experiments for compression and branch pruning, we provide more in-depth quantitative results for potential practical applications.
3. Densification adds a whole new paradigm to our framework. A generalised bottom-up scale-space can not only grow a skeleton from a single point, but also allows to add finer scales beyond the original skeleton.
4. We illustrate the usefulness of densification scale-spaces with a practical application for stiffness enhancement [28].

1.2 Related Work

Research on the medial axis transform and scale-space theory provide the foundation for our work. We rely on classical definitions and fundamental properties from both fields. In

the following we provide a short overview of the literature on skeleton pruning and thinning, as well as its image counter-part of sparsification and densification. These two fields inspire concrete implementations and practical applications of our scale-space framework.

1.2.1 Medial Axis Transform

As noted by Hajłasz [29], the underlying concepts of the medial axis were already studied by Erdős [30] before the term itself was coined. Blum [8] rediscovered the medial axis as an alternative to Euclidean geometry, which he deemed too specific, and topology, which he saw as too general. The idea of a central shape descriptor was motivated by the so-called grass-fire analogy, where prairie fires spread from the boundaries. The skeleton is formed by shock locations where these propagating fronts meet inside of the object. From these points, the object can be reconstructed [31]. The evolution of the propagating front in the grass-fire model is closely related to the Eikonal equation, which can be solved by fast marching methods [32] or in terms of viscosity solutions [33]. However, solving the Eikonal equation alone does not suffice to compute the skeleton since the locations of the shocks need to be explicitly identified [34].

Relations between the shape boundary and the medial axis have also been explored by Kuijper and Olsen [35] as well as Damon [36]. The medial axis is a subset of the more general symmetry sets [37], which are closely related, but not part of our theoretical framework. Therefore, we refer to [21, 23, 38] for a more detailed review of this literature. For our purpose, classification of skeleton points into categories like end or branching points is crucial for our concrete task-specific scale-spaces. Here, we rely on the work of Giblin and Kimia [39].

Our work is not the first to connect skeletons to multi-scale ideas. In Sect. 1.2.3 we discuss pruning and thinning algorithms separately as prime examples that come closest to our own notion of skeleton sparsification. Additionally, Pizer et al. [21] and Saha et al. [23] also provide an overview of more broadly related methods that use PDEs for shock computation [40] or use the concept of cores [26]. The latter uses the notion of medialness, a relaxation of the strict binary conditions of the medial axis by observing evidence for points being approximately centred in a shape in a disc-shaped neighbourhood. The radius of this neighbourhood defines a scale-parameter and cores are defined via traces in such a scale-space. Like similar concepts such as the multi-scale medial axis of Pizer et al. [25] or the learning-based approach of Kokkinos [24], the focus differs from our own work: We develop a scale-space theory by explicitly applying sparsification techniques to the skeleton.

The medial axis transform has a wide range of practical applications that use either 2-D or 3-D [38] skeletons. Clas-

sically, the medial axis has been used in computer vision for tasks like object recognition and shape matching [41], but there are also applications in physics simulations [42], compression [43], CAD modelling [44] or 3-D printing [28].

1.2.2 Scale-Space Theory

Classical scale-spaces as introduced by Iijima [2] rely on Gaussian blur to gradually remove image features with increasing scale. Such evolutions can be expressed by a linear partial differential equation (PDE). This theory has been studied in great detail [1, 2, 45–48] and has been extended to nonlinear [4, 5] or pseudodifferential evolutions [49–52]. Beyond linear and nonlinear diffusion processes, variational models have been studied in a scale-space context as well. The total variation (TV) flow is a particularly well-known alternative nonlinear evolution that has been studied extensively [53, 54]. More general variational approaches to image and shape evolution have been considered for instance in shape space [55]. While we do not rely on PDEs or variational methods in this work and restrict ourselves to operations on skeletons, the core principles of classical scale-space theory are highly relevant. We build our theory on the architectural, invariance, and information reduction concepts introduced by Alvarez et al. [1].

Regarding our goal of constructing hierarchical shape evolutions, the class of morphological scale-spaces [1, 15–19] is closer to our work than the aforementioned Gaussian category. The basic building blocks of morphological scale-spaces are dilation and erosion [20] operations. In a binary image, these supremum and infimum operations affect the boundary of the shape and thus naturally lead to boundary evolutions. Our framework shares similarities with some of these scale-spaces, for instance the hierarchical shrinking of objects for increasingly coarser scales. However, we define our evolution entirely on the skeleton as an equivalent shape descriptor instead of modifying the boundary.

1.2.3 Skeleton Pruning and Thinning

The sequential removal of skeleton parts has been studied for two major purposes: *Thinning* addresses the fact that discrete skeletons obtained by practical methods of computation do not necessarily have the desired one-pixel width. *Pruning* approaches focus instead of removing spurious skeleton branches that arise from noisy shape boundaries. Both can be used in tandem, typically thinning the skeleton first and applying pruning afterwards. Some algorithms attempt to address both goals simultaneously by integrating thinning and pruning criteria into a single framework [56–59].

Thinning: At first glance, thinning methods [56–60] seem closely related to our sparsification since they remove individual pixels step-by-step. However, they are a method

of discrete skeleton computation. Thus, they successively remove points from the original shape based on some criterion of medialness. This ensures that the final skeleton is thin (one pixel width). In contrast, we remove points from the already computed skeleton. However, for the computation of our initial skeleton in our experiments on discrete data, we use maximal disc thinning [59]. It is a homotopic thinning method that preserves the connectivity of the initial shape and removes shape pixels ordered by their distance from the boundary. Skeleton end points define a stopping criterion for the algorithm and rely on the maximal disc identification method of Remy and Thiel [61].

Pruning: In contrast to thinning, branch pruning methods [9–14] are not limited to the discrete setting. Small bumps or kinks in the object boundary can have a large impact on the skeleton, adding substantial branches both in the continuous and the discrete setting. Since pruning approaches remove parts of the skeleton, they are more closely related to our sparsification scale-spaces. We generalise the removal of skeleton points, but also show that classical branch pruning is one of the practical applications that is covered by our scale-space theory Sect. 4. In particular, our experiments are inspired by the hierarchic branch pruning of Ogniewicz [11].

1.2.4 Sparsification and Densification

Our scale-space framework is inspired by two previous publications that successively remove image content to define a family of simplified images. Cárdenas et al. [6] use spatial sparsification: They remove image pixels to transition from one scale to the next. These techniques originate from so-called spatial optimisation problems in image compression and sparse data representation [62–64]. Here, a set of optimal known data points has to be chosen for image reconstruction with interpolation methods such as PDE-based inpainting. Similar concepts have also been used for a sparsification of the co-domain. Quantisation scale-spaces [65] reduce the number of different grey values to construct a multi-scale family of images.

The opposite notion of densification [66–68] is well-known in spatial optimisation as well. Such approaches typically start with an empty image and greedily add pixels that are most important for the image reconstruction. Compared to sparsification, these approaches are often more robust to noise, since the initially chosen pixels have a global impact instead of the highly localised influence of individual pixels in a dense known data set. For our skeletonisation scale-spaces, we leverage both of the aforementioned paradigms.

1.3 Organisation

Section 2 provides the theoretical background required for our new scale-spaces in Sect. 3, which we adapt to specific

applications in Sect. 4. We conclude with a discussion and outlook on future work in Sect. 5.

2 Review: Skeletonisation

For our shape sparsification scale-spaces, we consider binary images that arise for instance from segmentation or object detection algorithms. They separate the image into the background represented by value 0 and the object marked by value 1. In the continuous setting with an image domain $\Omega \subset \mathbb{R}^2$, we write images as functions $f : \Omega \rightarrow \{0, 1\}$. We denote the set of object points by $O = \{x \in \Omega \mid f(x) = 1\}$. Thus, the binary image f can be interpreted as the indicator function of the object set O , i.e. $f = \chi_O$.

In particular, we assume that O is bounded and closed. Thus the set is compact and the boundary ∂O is included in O . While we can describe the object either by the set O or its boundary ∂O , we require an alternative shape descriptor for our scale-space framework.

The medial axis transform (MAT) of Blum [8] is the foundation on which we build our shape sparsification scale-spaces. In order to define it, we require a distance $d(x, y)$ for $x, y \in \Omega$, which we assume to be Euclidean. The MAT is a symmetric shape descriptor with respect to the boundary. Therefore, we first define the distance map, which maps each spatial location in the image domain to its distance to the object boundary ∂O .

Definition 1 (*Distance Map D*) The distance map denotes the minimum distance to the boundary ∂O according to

$$D : O \rightarrow \mathbb{R}, \quad x \mapsto \min_{y \in \partial O} d(x, y) \quad (1)$$

The distance map values for each location $x \in \Omega$ coincide with the radii of the largest discs that fit into the object if centred at x .

Definition 2 (*Inscribed Disc $B_{D(x)}$*) We denote the largest disc with centre x that is fully contained in O by

$$B_{D(x)}(x) := \{y \in \Omega \mid d(x, y) \leq D(x)\} \quad (2)$$

These discs are useful for defining the medial axis $\Sigma \subset O$. It contains all points that are centred in the object O , which means that they have equal distance from at least two boundary points. Thus, all medial axis points x are centres of inscribed discs that touch the image boundary at least at two points. This corresponds exactly to the subset of largest inscribed discs which are not contained in any other inscribed disc with a different centre y . This allows us to define the medial axis or skeleton.

Definition 3 (*Skeleton Σ*) The skeleton or medial axis (MA) Σ is defined as the set of centres of inscribed discs of maximal radius [70]:

$$\Sigma := \{x \in O \mid \forall y \in O \setminus \{x\} : B_{D(x)}(x) \not\subset B_{D(y)}(y)\}.$$

In case the object O is not connected, each component of O will have its own separate skeleton. Note that the skeleton definition also does not require smoothness assumptions on the boundary of O . In our scale-space framework, we will not only consider a single static object and corresponding skeleton, but whole families of these pairs. Therefore we explicitly introduce notations for mapping images to skeletons and vice versa.

Definition 4 (*Skeleton, Reconstruction, and Object Transform*)

- The skeleton transform $\mathcal{S}(f) = \Sigma$ maps the image f to its skeleton Σ as defined in Def. 3.
- The object transform $\mathcal{O}(\Sigma)$ maps a skeleton back to the corresponding object according to

$$\mathcal{O}(\Sigma) := \bigcup_{x \in \Sigma} B_{D(x)}(x). \quad (3)$$

- The reconstruction transform $\mathcal{R}(\Sigma)$ maps the reconstructed skeleton back to a binary image $\mathcal{R}(\Sigma) : \Omega \rightarrow \{0, 1\}$ with

$$\mathcal{R}(\Sigma)(x) := \begin{cases} 1 & \text{for } x \in \mathcal{O}(\Sigma) \\ 0 & \text{else} \end{cases}. \quad (4)$$

According to the classical definitions of Blum [8] and Pfaltz and Rosenfeld [70], all three transforms are well-posed. However, in the discrete setting, uniqueness cannot always be guaranteed. We discuss these aspects in the following Sects. 2.1 and 2.2 and consider the respective impact on our scale-space frameworks in Sect. 3.

Note that the transforms also imply that we can interpret operations on the skeleton in terms of their effect on the object and vice versa. Removing or inserting skeleton points modifies the set Σ and therefore changes the associated object $O = \mathcal{O}(\Sigma)$. Conversely, modifying the object and recomputing its skeleton via the skeleton transform $\mathcal{S}$ leads to a modified skeleton structure. Thus both representations are coupled through the transforms $\mathcal{S}$ and $\mathcal{O}$. Operations performed in one domain induce corresponding changes in the other.

Typically, a single skeleton point x carries more information than an object point since it represents the whole section of the shape covered by the corresponding disc $B_{D(x)}$. However, its influence on the overall shape on removal also




Fig. 1 Skeleton examples for images apple-8 and apple-9 of the CE-Shape-1 database [69]. In (a), the discrete skeleton of the black shape is marked in red. The corresponding distance map (b) indicates large distances to the boundary by bright grey values. The skeleton corresponds to ridges of the distance map inside of the object. The false colour representation of the overlap (c) illustrates how many skeleton points

reconstruct the same object point: blue indicates the lowest amount of reconstruction overlap and red the highest. a, d together show the impact of boundary perturbations on the skeleton. The leaf and stem at the top of the apple create large additional branches by adding bumps to the boundary of the elliptic shape, while apple-9 introduces new cavities to the boundary that also create additional skeleton branches

depends heavily on the rest of the skeleton. We address in more detail with the unique impact criterion in Definition 13. Conversely, a very small impact on the overall shape, e.g. removing single pixels at the boundary, can have a significant impact on the skeleton as illustrated in Fig. 1a, d.

2.1 Central Skeleton Properties

The following four central properties of skeletons are all relevant for our scale-space considerations. All of them can be already found directly or indirectly in the original publication of Blum [8], even though formal proofs were provided later in some cases.

Property S1 (Equivalent Shape Descriptor) In the continuous setting, the skeleton Σ , together with the radii of the maximal inscribed discs provides an equivalent representation of the object, that is $\mathcal{O}(\Sigma) = \mathcal{O}$. Consequentially, in this setting, skeleton and reconstruction transforms are truly inverse transforms with $\mathcal{R}(\mathcal{S}(f)) = f$ [70]. Furthermore, $\mathcal{O}(X)$ defines a continuous mapping of any set X . This implies that if X is compact, then the reconstruction $\mathcal{O}(X)$ is compact as well. This is relevant when we build multiple reconstructions on sparsified skeletons $X \subset \Sigma$.

Property S2 (Thin Set) The medial axis of any closed set consists of thin curves. Mathematically speaking, it has a Hausdorff measure of zero [29, 30]: $\mathcal{H}(\Sigma) = 0$.

Property S3 (Homotopy) The skeleton transform preserves the connectivity of the original shape [39, 71].

Property S4 (Equivariance under Euclidean Motion) Consider Euclidean transformations $\mathcal{T}$ that consist of translation, rotation, reflection, or uniform scaling. Then, the skeleton transform $\mathcal{S}$ commutes with the Euclidean transformation: $\mathcal{S}(\mathcal{T}(f)) = \mathcal{T}(\mathcal{S}(f))$.

2.2 Discrete Setting

In Sect. 3 we discuss not only continuous scale-spaces, but also their discrete or semi-discrete variants. In a *spatially discrete setting*, we transition from the continuous domain Ω of image coordinates to an image with a finite resolution $n_x \times n_y$. In this case, we use a vectorial notation $f \in \{0, 1\}^n$ where the indices $\Omega_D = \{1, \dots, n\}$ correspond to a row-by-row sequence of the $n = n_x \cdot n_y$ pixels. We transfer Definitions 1 and 2 to the discrete setting by replacing Ω by Ω_D and using the Euclidean distance between pixel centres. Then, the skeleton from Definition 3 and the transforms from Definition 4 carry over as well. However, not all of the properties are preserved.

On a discrete pixel grid, neither the exact positions of the medial axis points nor the shape of discs can be reproduced exactly. Thus, enforcing a thin skeleton in the sense of property S2 with one pixel width might also imply that not all points of the original shape are covered by discrete discs. Therefore, reconstructions according to Definition 4 do not necessarily reconstruct the full object, i.e. the case $\mathcal{O}(\Sigma) \neq \mathcal{O}$ and $\mathcal{R}(\mathcal{S}(f)) \neq f$ can occur.

The equivariance Property S2 still holds for on-grid translations, reflections, and rotations in 90° degree steps. However, rescaling operations and rotations by other angles affect the discrete boundary of the object. For these cases there is no guarantee that computing the skeleton from the transformed shape is identical to transforming the discrete skeleton. This may even vary depending on the algorithm used to compute the discrete skeleton. The choice of this algorithm also affects Property S2. In our work, we always use homotopy preserving skeleton computations, namely the maximal disc thinning algorithm [59]. For such homotopy preservation, the following definitions are vital:

Definition 5 (Endpoints, Branching Points, Simple Points, and Arcs) Let a 4-neighbourhood contain left, right, upper,

and lower adjacent pixels, while an 8-neighbourhood additionally contains the diagonally adjacent pixels.

- An **endpoint** x of Σ has either: 1.) zero or one skeleton points in their 8-neighbourhood; 2.) two skeleton points in its 8-neighbourhood which are adjacent to one another; 3.) three skeleton points in its 8-neighbourhood that are either all above, below, left, or right of x . We denote the set of endpoints by $E(\Sigma)$.
- A **branching point** of Σ has at least three skeleton points in its 8-neighbourhood which are not horizontally or vertically adjacent. We denote the set of branching points by $B(\Sigma)$.
- A **simple point** is a skeleton point which is neither an endpoint nor a branching point. We write $S(\Sigma) := \Sigma \setminus (E(\Sigma) \cup B(\Sigma))$.
- We define $\mathcal{A}(\Sigma)$ as the set of all arcs/branches. Each of its elements $A = \{a_1, \dots, a_k\} \subset \Sigma$ fulfils two properties: 1.) The arc is *connected*, i.e. for each i the points a_i and a_{i+1} are 8-neighbours. 2.) Exactly a_1 and a_k are *end- or branching points*, the rest of the a_i are simple.

In some cases the skeleton may consist solely of closed loops without endpoints or branching points, for instance for shapes with holes such as an annulus. In this case we have $E(\Sigma) = \emptyset$ and $B(\Sigma) = \emptyset$, so Definition 5 does not yield any arcs. We treat such special cases appropriately in our framework (e.g. in Algorithm 1).

These definitions are also relevant for practical applications such as shape matching. Moreover, we require them for our task-specific scale-spaces and quality measures in Sect. 4.

3 Skeletonisation Scale-Spaces

In the following, we transfer the concept of sparsification scale-spaces to the medial axis and propose a new class of scale-spaces. Intentionally, we keep our new class of skeletonisation scale-spaces very general. In particular, we keep the amount of requirements for our frameworks very low. As a drawback of this approach, in this section we are only able to provide generic simplification guarantees. This is well worth the cost however, since we also enjoy the liberty of designing sparsification or densification paths with very small amounts of restrictions. In Sect. 4, we make use of these opportunities with a diverse set of application-specific scale-spaces that illustrate how we can equip our foundational frameworks with additional requirements that yield meaningful simplification guarantees.

For sparsification scale-spaces, Cárdenas et al. [6] have so far only provided a fully discrete theory where both the spatial domain of image pixels and the scale are discrete. In

contrast, we cover the full spectrum of continuous, discrete, and semi-discrete theory. We discuss the fully continuous and discrete setting in detail. To avoid redundancies, we treat space-continuous and scale-continuous semi-discrete scale-spaces jointly in Sect. 3.4.

3.1 Fully Continuous Skeleton Sparsification Scale-Spaces

We construct a scale-space by sequentially removing skeleton points from the initial skeleton $\Sigma(0) := \mathcal{S}(f)$ obtained from the binary image $f : \Omega \rightarrow \{0, 1\}$ containing the discrete object $O = \{x \in \Omega \mid f(x) = 1\}$. With the artificial time parameter $t > 0$ we represent the evolution of the skeleton $\Sigma(t)$ starting with initial time $t = 0$. For increasing t , we transition to coarser scales.

Definition 6 (Skeleton Evolution) Given the initial skeleton $\Sigma(0)$, a *continuous sparsification path* $\Sigma : [0, \infty) \rightarrow \mathcal{P}(\Sigma(0))$ is a function that maps any positive time t to an element of the power set $\mathcal{P}(\Sigma(0))$ of the initial skeleton. This powerset contains all possible subsets of the skeleton that can result from subsequent removal of skeleton points. The evolution needs to fulfil the following requirements:

1. **CR1 (nested skeletons):** The skeletons $\Sigma(t)$ are nested over time, i.e. $\Sigma(t_2) \subset \Sigma(t_1)$ for $t_2 > t_1$.
2. **CR2 (enveloping disc):** There exists an enveloping closed disc: $D(t)$ with linearly decreasing radius $r(t) := \max\{-t \cdot s + r_0, 0\}$ with $r_0 > 0$ and $s > 0$, and centre $c_D \notin \Sigma(0)$ such that at any time $t > 0$, the skeleton $\Sigma(t)$ is included in the enveloping disc: $\Sigma(t) \subset D(t)$.

Since we want our scale-space framework to cover a broad range of relevant evolutions, we keep the requirements very general. The nested skeleton criterion CR1 ensure that once a skeleton point has been removed, it is never added back. The enveloping disc criterion CR2 is less intuitive, but we discuss why we took this specific design choice in Sect. 3.1.2. CR2 is an auxiliary construction designed to guarantee a minimal shrinkage speed s and makes the discrete and continuous setting more consistent. As such, we can choose the disc centre freely outside of the skeleton. Based on the skeleton evolution, we define the sparsification path that indicates, which parts of the skeleton are removed.

Definition 7 (Continuous Sparsification Path) Given a skeleton Σ , a *continuous sparsification path* $P : [0, \infty) \times [0, \infty) \rightarrow \mathcal{P}(\Sigma)$ is a function that maps a time frame with start t_1 and end $t_2 > t_1$ to the subset $P(t_1, t_2)$ of the skeleton Σ that is removed in this time frame:

$$P(t_1, t_2) := \Sigma(t_1) \setminus \Sigma(t_2)$$

In particular, for all $t \geq 0$ it holds that $P(t, t) = \emptyset$. We can now define our scale-space.

Definition 8 (*Continuous Skeleton Sparsification Scale-Space*)

Consider the binary image $f : \Omega \rightarrow \{0, 1\}$ with domain Ω , skeleton $\Sigma = \mathcal{S}(f)$, reconstruction $\mathcal{R}(\Sigma)$ and sparsification path $P(\cdot, \cdot)$. The *continuous skeleton sparsification scale-space* is the family $(u(t), \Sigma(t))$ of images $u(t, \mathbf{x}) : [0, \infty) \times \Omega \rightarrow \{0, 1\}$ and evolving skeletons $\Sigma : [0, \infty) \rightarrow \mathcal{P}(\Sigma)$ obeying

1. $\Sigma(0) = \Sigma = \mathcal{S}(f)$,
2. $\Sigma(t) = \Sigma(0) \setminus P(0, t)$ for $t > 0$,
3. $u(t) := \mathcal{R}(\Sigma(t))$ and $O(t) := \mathcal{O}(\Sigma(t))$ for $t \geq 0$.

3.1.1 Continuous Scale-Space Properties

Most of the properties C1–C6 discussed in the following verify classical scale-space concepts of Alvarez et al. [1]. The presence of corresponding architectural, simplification, and equivariance features identifies our newly constructed framework as a scale-space. Only Property C4 is entirely specific to our shape analysis case since it verifies that sparsification does not destroy important skeleton characteristics. All other properties have classical scale-space analogues.

Property C1 (Original Skeleton and Image as Initial State)

By Definition 8, we have $\Sigma(0) = \mathcal{S}(f) = \Sigma$. Due to Property S1 of the medial axis transform, $u(0) = \mathcal{R}(\Sigma(0)) = \mathcal{R}(\mathcal{S}(f)) = f$.

Property C2 (Causality) We can equivalently reach the skeleton-image-pair $(u(t), \Sigma(t))$ in a step of size t from time 0 or in a step of size $t - \hat{t}$ from time $\hat{t} < t$. This holds according to Definitions 7 and 8 since

$$\Sigma(0) \setminus P(0, t) = \Sigma(t) = \Sigma(\hat{t}) \setminus P(\hat{t}, t). \quad (5)$$

In fact, due to the nested set property CR1, we always have $P(t_2, t_3) \subset P(t_1, t_3)$ for $t_1 < t_2 < t_3$.

Property C3 (Lyapunov Sequences) Scale-spaces need quantifiable simplification over time. Due to the generality of our framework, we only prove generic Lyapunov sequences here. For concrete applications in Sect. 4, we can find Lyapunov sequences that reflect task-specific measures of simplification. A simple Lyapunov sequence can be defined via the object area.

Proposition 1 (Decreasing Object Area) *The object area $a(t) := \mathcal{H}(O(t))$ in terms of its Hausdorff measure decreases over time t , i.e. $a(t_2) \leq a(t_1)$ for $t_2 > t_1$.*

Proof For a time $t_2 > t_1$, we have $(*) : \Sigma(t_2) \subset \Sigma(t_1)$ due to requirement CR1. Therefore, we can infer

$$\mathcal{O}(\Sigma(t_2)) = \bigcup_{\mathbf{x} \in \Sigma(t_2)} B_{D(\mathbf{x})}(\mathbf{x}) \stackrel{(*)}{\subset} \bigcup_{\mathbf{x} \in \Sigma(t_1)} B_{D(\mathbf{x})}(\mathbf{x}) = \mathcal{O}(\Sigma(t_1)) \quad (6)$$

From $\mathcal{O}(\Sigma(t_2)) \subset \mathcal{O}(\Sigma(t_1))$ we can conclude $a(t_2) \leq a(t_1)$. $\square$

Instead of the area, we can also use the diameter $\text{dia}(\mathcal{O}(\Sigma(t)))$ of the reconstructed object to define a Lyapunov sequence. For a closed set $S \subset \Omega$ we define

$$\text{dia}(S) := \max\{d(\mathbf{x}, \mathbf{y}) \mid \mathbf{x}, \mathbf{y} \in S\},$$

where $d(i, j)$ is the Euclidean distance.

Proposition 2 (Decreasing Object Diameter) *The object diameter decreases with increasing time t , i.e. for $t_2 > t_1 \geq 0$, we have*

$$\text{dia}(\mathcal{O}(\Sigma(t_2))) \leq \text{dia}(\mathcal{O}(\Sigma(t_1))). \quad (7)$$

Proof Assume that $\text{dia}(\mathcal{O}(\Sigma(t_2))) > \text{dia}(\mathcal{O}(\Sigma(t_1)))$. Then, there are $\mathbf{x}, \mathbf{y} \in \mathcal{O}(\Sigma(t_2))$ with $d(\mathbf{x}, \mathbf{y}) > \text{dia}(\mathcal{O}(\Sigma(t_1)))$. However, we know that for $t_2 > t_1$, we have $\mathcal{O}(\Sigma(t_2)) \subset \mathcal{O}(\Sigma(t_1))$ according to our previous proof for Proposition 1. Therefore, we also have $\mathbf{x}, \mathbf{y} \in \mathcal{O}(\Sigma(t_1))$ which contradicts our assumption. Thus, the object diameter at t_2 cannot be larger than the diameter at time t_1 . $\square$

Property C4 (Skeleton Properties) The properties S1 (equivalent shape descriptor) and S2 (thin set) remain fulfilled under sparsification. We cover the equivariances S2 separately.

S1 : By definition, $u(t) = \mathcal{R}(\Sigma(t))$.

S2 : For all $t > 0$, we have $\Sigma(t) \subset \Sigma$ and $\mathcal{H}(\Sigma) = 0$. Therefore, also $\mathcal{H}(\Sigma(t)) = 0$ and thus all evolving skeletons are thin sets.

Note that we intentionally do not require the homotopy preservation Property S2 to be fulfilled for all of our scale-spaces. This means that due to sparsification, a shape can potentially get disconnected. In Sect. 4.2, we show how additional requirements can add this property if desired.

Property C5 (Equivariance) In Sect. 2.1 we have reviewed the inherent equivariance property S2 of the medial axis: The MAT commutes with any Euclidean transformation $\mathcal{T}$, including translation, rotation, reflection, and scaling. In Property C4, we have also established that the reconstruction criterion $u(t) = \mathcal{R}(\Sigma(t))$ holds for all t . Thus, for all $t > 0$, we have skeleton/image pairs for which the equivariance

$\mathcal{S}(\mathcal{T}(u(t))) = \mathcal{T}(\mathcal{S}(u(t))) = \mathcal{T}(\Sigma(t))$ holds. Consequently, the equivariances transfer directly from the MAT to the whole scale-space evolution: if we apply the transform $\mathcal{T}$ to f at time $t = 0$, the image skeleton pairs are transformed accordingly by $\mathcal{T}$.

Property C6 (Empty Image as Steady State) We know that the enveloping disc D is centred in $c_D \notin \Sigma$. Let r_m denote the minimal distance of c_D to the set Σ . At time $T > \frac{r_0 - r_m}{s}$, we thus get $D(T) \cap \Sigma(0) = \emptyset$. Since $\Sigma(T) \subset \Sigma(0)$, the requirement $\Sigma(T) \subset D(T)$ can only be fulfilled for $\Sigma(T) = \emptyset$. Thus, the reconstruction is an empty image as well with object $O(t) = \emptyset$ and we reach an empty steady state in finite time.

3.1.2 Impact of the Enveloping disc Requirement

If we consider the proofs of our six scale-space properties, the enveloping disc requirement of Definition 7 has only an impact on Property 6, the steady state. It ensures that the skeleton is removed completely in finite time and thus the final reconstruction is an empty image. Alternatively, we could require the nesting property CR1 to enforce proper subsets $\Sigma(t_2) \subsetneq \Sigma(t_1)$ for $t_2 > t_1$.

However, the intersection of all $\Sigma(t)$ would still be non-empty and (given a shrinking diameter) would converge to a single skeleton point. While this would constitute a valid steady state evolution in line with typical scale-spaces, it would deviate from the discrete case which has an empty steady state as we discuss in Sect. 3.2. Here, for the sake of consistency, we have chosen to enforce empty steady states for both. Alternatively, we could also modify the requirements for discrete scale-spaces such that their steady state becomes a single skeleton point.

Note that the linear shrinkage of the enveloping disc is not an overly restrictive requirement for possible sparsification paths. It is an auxiliary construction that only influences the implicit parametrisation of the artificial time variable. Nonlinear and asymmetric shrinkage of the skeleton is not ruled out by the definition, as long as this shrinkage is overall fast enough relative to the reference speed s provided by the shrinkage of the enveloping disc. Thus, with an appropriate choice of s and c_D , all skeleton sparsification paths of interest are possible.

3.2 Fully Discrete Skeleton Sparsification Scale-Spaces

The fully discrete setting is relevant for our practical implementations in Sect. 4. Here we consider an $n_x \times n_y$ discrete binary image $f \in \{0, 1\}^{n_x n_y}$ with discrete image domain $\Omega_D = \{1, \dots, n_x n_y\}$. In contrast to the fully continuous setting, it is more convenient to explicitly define the sparsi-

fication path first, especially since we are going to use it as a design criterion for the scale-space later on.

Definition 9 (Discrete Sparsification Path) Given a skeleton $\Sigma \subset \Omega_D$, a *sparsification path* $P = (P_1, \dots, P_m)$ with $m \in \mathbb{N}^+$ is an ordered collection of non-empty sets that form a partition of Σ , i.e.:

1. $P_k \neq \emptyset$,
2. $P_k \cap P_\ell = \emptyset, k \neq \ell$
3. $\bigcup_{\ell=1}^m P_\ell = \Sigma$.

Based on this finite number of sparsification steps, we can define the scale-space. It is not only discrete in space, but also has a discrete scale parameter ℓ .

Definition 10 (Discrete Skeleton Sparsification Scale-Space)

Consider the discrete binary image $f \in \{0, 1\}^{n_x n_y}$ with domain Ω_D , skeleton $\Sigma = \mathcal{S}(f)$, reconstruction $\mathcal{R}(\Sigma)$ and sparsification path $P = (P_1, \dots, P_m)$ partitioning Σ , $m \in \mathbb{N}^+$. The *discrete skeletonisation scale-space* is the family $(u_\ell, \Sigma_\ell)_{\ell=0}^m$ of images u_ℓ and skeletons Σ_ℓ obeying

1. $\Sigma_0 := \Sigma = \mathcal{S}(f)$,
2. $\Sigma_\ell := \Sigma_0 \setminus \bigcup_{i=1}^\ell P_i$ for $\ell \in \{1, \dots, m\}$,
3. $u_\ell := \mathcal{R}(\Sigma_\ell)$ and $O_\ell = \mathcal{O}(\Sigma_\ell)$ for $\ell \in \{0, \dots, m\}$.

As in the continuous case, the sparsification path defines which points to prune in the transition from scale ℓ to scale $\ell + 1$. However, the number m of path sets uniquely determines the finite number $m + 1$ of discrete scales. Both the number of steps and the composition of the sets can be freely chosen as long as the partition condition is fulfilled. Note that we also require $P_\ell \neq \emptyset$ and therefore at least one skeleton point is removed per sparsification steps. This implies a maximum of $m \leq |\Sigma|$ sparsification steps.

3.3 Discrete Scale-Space Properties

The following properties D1-D6 are discrete counterparts to C1-C6 from the continuous setting. Many of the properties carry over in very similar fashion, but there are a few notable exceptions.

Property D1 (Original Skeleton as Initial State) As in the continuous setting, $\Sigma_0 = \mathcal{S}(f)$ at scale $\ell = 0$ is the skeleton of the original object due to Definition 10. However, as we have pointed out in Sect. 2.2, there is no guarantee that $\mathcal{R}(\Sigma) = f$ in the discrete setting. Thus, the initial image u_0 is typically only an approximation to f .

Property D2 (Causality) By following the sparsification path $P = (P_1, \dots, P_m)$, a scale $\ell \in \{0, \dots, m\}$ can be

reached either from the initial scale 0 or from any intermediate scale $k \in \{0, \dots, \ell - 1\}$. Similarly to the continuous case, we use Definition 10 to verify

$$\Sigma_\ell = \Sigma_0 \setminus \bigcup_{i=1}^{\ell} P_i = \left(\Sigma_0 \setminus \bigcup_{i=1}^k P_i \right) \setminus \bigcup_{i=k+1}^{\ell} P_i = \Sigma_k \setminus \bigcup_{i=k+1}^{\ell} P_i. \quad (8)$$

In the discrete setting, the time spans are replaced by a concrete number of steps: Scale ℓ can be reached in ℓ steps from the initial scale or in $\ell - k$ steps from scale k .

Property D3 (Lyapunov Sequences) The same two generic Lyapunov sequences as in the continuous setting apply. The proof for the *decreasing object diameter* remains unchanged and therefore, we do not repeat it here. For the *decreasing object area*, we have slight adaptations due to deviations in the definitions, but the proof remains very similar. Here, the image area is simply defined by the cardinality $|\cdot|$ of the object set.

Proposition 3 (Decreasing Object Area) *The area $a_\ell := |O_\ell|$ of the object decreases as the scale parameter ℓ increases, i.e. $a_k \leq a_\ell$ for $k > \ell$.*

Proof For the path $P = (P_1, \dots, P_m)$ we have $\Sigma_\ell = \Sigma_{\ell+1} \cup P_{\ell+1}$ and thus

$$\bigcup_{x \in \Sigma_\ell} B_{D(x)}(x) = \bigcup_{x \in \Sigma_{\ell+1}} B_{D(x)}(x) \cup \bigcup_{x \in P_{\ell+1}} B_{D(x)}(x).$$

Therefore, we can derive

$$\mathcal{O}(\Sigma_\ell) = \bigcup_{x \in \Sigma_\ell} B_{D(x)}(x) \supseteq \bigcup_{x \in \Sigma_{\ell+1}} B_{D(x)}(x) = \mathcal{O}(\Sigma_{\ell+1}).$$

Thus, we can conclude $a_\ell = |\mathcal{O}(\Sigma_\ell)| \geq |\mathcal{O}(\Sigma_{\ell+1})| = a_{\ell+1}$. $\square$

In general the object area is not strictly decreasing. Although each sparsification step removes at least one skeleton point ($P_{\ell+1} \neq \emptyset$), the reconstructed object may remain unchanged if all discs associated with the removed points are fully covered by discs of the remaining skeleton.

Property D4 (Skeleton Properties) The property follows from the corresponding properties of the discrete skeletonisation and reconstruction operators.

Even though discrete skeletons do not necessarily fulfil the reconstruction criterion S1, $u_\ell = \mathcal{R}(\Sigma_\ell)$ is fulfilled by Definition 10. Since the object is created from the sparsified skeleton itself, there are no discretisation artefacts beyond the potential mismatch to the original shape as discussed in Property D1.

By definition, the initial skeleton Σ has one pixel thickness. Removing pixels cannot increase this thickness and therefore, the property S2 is preserved for all Σ_ℓ , $\ell \in \{0, \dots, m\}$.

Together, these properties ensure that the skeletonisation and reconstruction operators are consistent with each other in the discrete setting, analogous to the continuous case discussed in Property C4.

As in the continuous setting, removing skeleton points without additional constraints can disconnect the skeleton graph and thus change the topology of the skeleton. The homotopy preservation according to S2 is thus not generally fulfilled.

Property D5 (Equivariance) As in the continuous setting, the equivariance properties directly carry over from the medial axis to the scale-space evolution. However, following our discussion of discrete skeletons in Sect. 2.2, this does not apply to all Euclidean transformations, only to on-grid translations, 90° rotations, and mirroring along image axes.

Property D6 (Empty Image as Steady State) Since $P_1, \dots, P_m$ partition Σ and $\Sigma_0 = \Sigma$ by Definition 9 and 10, we can immediately conclude $\Sigma_m = \Sigma \setminus \bigcup_{\ell=1}^m P_\ell = \Sigma \setminus \Sigma = \emptyset$. Since the final skeleton is empty, the steady state is an empty image.

In summary, the discrete scale-spaces capture all major properties of the continuous setting. Only discretisation artefacts in the reconstruction and Euclidean transformations limit properties D1 and D4 slightly compared to their continuous counterparts. This shows that our discrete implementations capture the main ideas of the proposed scale-space framework. While several properties carry over directly to the discrete setting, others only hold in a weaker form or require additional considerations due to discretisation effects.

3.4 A Note on Semi-Discrete Skeletonisation Scale-Spaces

It is also possible to construct scale-spaces that are time-discrete and space-continuous or vice versa. Since properties and proofs closely resemble a combination of C1-C6 and D1-D6, we do not discuss them here in detail to avoid redundancies.

The time-continuous and space-discrete case inherits the drawbacks of both the fully discrete and the fully continuous scale-space: Equivariances and initial state are weakened by spatial discretisation artefacts. Moreover, the time continuous setting requires an additional criterion like the enveloping shrinking discs to make it consistent with the time-discrete steady state.

As expected, the second semi-discrete scale-space enjoys the benefits of both worlds. Since it is spatially continuous, it inherits full Euclidean equivariance and a perfect reconstruction of the initial shape. Since the scale parameter is discrete, we can define the sparsification path directly as in Definition 9. We only need to replace the discrete image f and the domain Ω_D by their continuous counterparts. In our previous conference publication [27], we have discussed this case jointly with the fully discrete scale-space framework.

3.5 Skeleton Densification Scale-Spaces

All previous scale-space frameworks were motivated by the sparsification paradigm that was also applied by Cárdenas et al. [6]. This is a top-down strategy, which starts with the full known data and gradually reduces it until nothing is left. However, in the literature on data optimisation, where sparsification has its origins, there is also an inverse counterpart: Densification [66–68] starts with an empty image and adds new known data in every step. In data optimisation, both strategies co-exist since they have different strengths and weaknesses (see Sect. 1.2.4). We want our scale-space theory to cover as many existing practical applications as possible. Therefore, we define densification scale-spaces as an alternative to sparsification.

To avoid redundancies to previous sections, we limit ourselves to the fully discrete theory and only discuss noteworthy differences. A fully continuous or semi-discrete theory can be constructed in an analogous way.

3.5.1 Densification Scale-Spaces

In a straight-forward way we can define a densification scale-space by inverting the role of the sparsification path. We require a *densification path* to adhere to the same requirements as in Definition 9: The sets $P_1, \dots, P_m$ are all non-empty and form a partition of the skeleton Σ . However, our initial skeleton $\Sigma_0 = \emptyset$ is now empty and the densification path specifies which points are *added* to the skeleton to transition from scale ℓ to $\ell + 1$. This leads us to the following scale-space definition.

Definition 11 (*Discrete Skeleton Densification Scale-Space*) Consider the discrete binary image $f \in \{0, 1\}^{n_x n_y}$ with domain Ω_D , skeleton $\Sigma = \mathcal{S}(f)$, reconstruction $\mathcal{R}(\Sigma)$ and sparsification path $P = (P_1, \dots, P_m)$ partitioning Σ , $m \in \mathbb{N}^+$. The *discrete skeletonisation scale-space* is the family $(u_\ell, \Sigma_\ell)_{\ell=0}^m$ of images u_ℓ and skeletons Σ_ℓ obeying

1. $\Sigma_0 := \emptyset$,
2. $\Sigma_\ell = \Sigma_0 \cup \bigcup_{i=1}^{\ell} P_{m-i}$ for $\ell \in \{1, \dots, m\}$,
3. $u_\ell := \mathcal{R}(\Sigma_\ell)$ and $O_\ell = \mathcal{O}(\Sigma_\ell)$ for $\ell \in \{0, \dots, m\}$.

The causality property D2, the skeleton properties D4 and equivariances D5 all carry over from sparsification scale-spaces directly with the same arguments. All other properties are only slightly adjusted to the new setting.

Property DD1: (Empty Skeleton and Image as Initial State)

By definition, $\Sigma_0 = \emptyset$ at scale $\ell = 0$ and thus the corresponding reconstructed image $u_0 = \mathcal{R}(\emptyset)$ is empty as well.

Property DD3: (Generalised Lyapunov-type Sequences)

The *object diameter* and *object area* are increasing. While Lyapunov sequences are typically decreasing, the reverse direction of the densification space has increasing sequences as a natural counterpart to the densification setting. The corresponding proofs carry over from D3 with minimal adaptations.

Property DD6: (Original Skeleton as Steady State)

Since $P_1, \dots, P_m$ partition Σ , we obtain $\Sigma_m := \Sigma_0 \cup \bigcup_{i=1}^m P_i = \Sigma$ according to Definition 11. The corresponding image $u_m = \mathcal{R}(\Sigma)$ approximates the original object according to the discrete skeletonisation algorithm used to compute Σ .

3.5.2 Extended Densification

The previous section shows that our theory extends to the inverse counterpart of sparsification. This result alone is not very surprising, but it also means that both directions in the scale-space are equally well-posed. In classical scale-space theory this is typically not the case. Going from fine to coarse scales is usually well-posed, but going from coarse to fine scales is often problematic due to ambiguities: For instance, a skeleton consisting of a single point could be a legitimate subset of an infinite number of larger skeletons containing said point. All of them constitute valid densifications of the initial skeleton. This noteworthy fact also opens up some entirely new perspectives.

The steady-state in sparsification is final by design. All points have been removed and there is no interesting information left in the empty image. In contrast, densification stops when the full skeleton is reached. This is a natural stopping point, but there are still interesting points left to be added. While the skeleton might give a full representation of the original object (in the continuous setting), this is not necessarily the case for discrete skeletons. There are many practical reasons why one might want to add more points. For instance, this could be the densification counter-part to pruning: Instead of removing erroneous branches, one could correct overly strict thinning applied by the skeletonisation algorithm. It might also be interesting to consider the transition to more general descriptors such as symmetry sets, a superset of skeletons. In Sect. 4, we highlight a recursive extension of the skeleton that adds additional points that yield

a shape representation that is overcomplete, but relevant, e.g., for 3-D printing.

Such generalisations of the skeleton scale-spaces can be obtained by simply lifting the requirement that the densification path forms a partition of the original skeleton, extending it to the whole object instead.

Definition 12 (Extended Densification Path) Given an object $O \subset \Omega_D$, an *extended densification path* $P = (P_1, \dots, P_m)$ with $m \in \mathbb{N}^+$ is an ordered collection of non-empty sets that form a partition of O , i.e.:

1. $P_k \neq \emptyset$,
2. $P_k \cap P_\ell = \emptyset, k \neq \ell$,
3. and $\bigcup_{\ell=1}^m P_\ell = O$.

At first glance, this seems like a full departure from skeletons to a pure shape densification. However, the scale-space definition still uses the reconstruction transform from an *overcomplete skeleton* Γ . The definition of such a scale-space is identical to Definition 11 except for the relaxed definition of the densification path. Still, we have $\Sigma_0 := \emptyset$, $\Sigma_\ell := \Sigma_0 \cup \bigcup_{i=1}^\ell P_i$, and $u_\ell := \mathcal{R}(\Sigma_\ell)$ for $\ell \in \{0, \dots, m\}$.

The initial state DD1, causality DD2, Lyapunov Sequences DD3, and equivariances DD5 are all not affected by this change, which follows with the same arguments as before.

Property DD4: (Skeleton Properties)

By definition, $u_\ell = \mathcal{R}(\Sigma_\ell)$. Thus, the reconstruction property still applies. In particular, for all ℓ , we still have $\mathcal{R}(\Sigma_\ell) \subset O$ since $\Gamma \subset O$. All discs around object points with radii given by the distance map are fully contained in the object O .

However, S2 is not fulfilled any more since Γ since the extended representation Γ is not restricted to thin skeleton pixels. During densification additional pixels from the reconstructed object may be inserted, so that Γ can locally become thicker than a discrete thin skeleton. This behaviour is illustrated later in Sect. 4.5, where the densification with extension progressively fills parts of the reconstructed object.

Property DD6: (Full Shape and Original Image as Steady State) Since $P_1, \dots, P_m$ partition O , we obtain $\Gamma_m = O$ and thus $u_m = \mathcal{R}(O) = f$. Both the final overcomplete skeleton Γ and its reconstruction are identical to the original object. Note that for a pixel on the object boundary we consider a reconstruction with discs that contain only the boundary pixel itself.

Surprisingly, this sweeping generalisation of the skeletonisation scale-spaces preserves most of its properties.

4 Task-Specific Skeletonisation Scale-Spaces

So far, our frameworks mostly focus on strong architectural and equivariance properties. In practical applications,

it is also desirable to have task-specific, quantifiable simplification. This guarantees that the scale-space provides a coarse-to-fine representation that is useful for the task to be solved.

In this section we demonstrate with proof-of-concept examples from different fields that our framework is highly adaptable. We present three classical applications: *branch pruning* as a pre-processing step of shape recognition, *shape compression*, and *stiffness enhancement* for 3-D printing. The first two applications rely on sparsification and have been published previously in our conference publication [27]. We have considerably extended the evaluations for these methods. For *compression*, we also add densification experiments as an alternative. With *stiffness enhancement*, we demonstrate how we can push the limits of our extended densification framework to provide a scale-space theory for applications that go beyond the classical definition of the medial axis.

4.1 Evaluation Design

While each of our practical examples has different focus, our experiments share common implementation details, data sets, and evaluation metrics. For the sake of reproducibility, we discuss these technical details first.


4.1.1 Skeletonisation Implementation

We require discrete skeletons as a starting point. As discussed in Sect. 2.2, all discrete skeletons are approximations and thus the exact definition of the discrete skeleton set Σ depends significantly on the method of computation. For all of our experiments, we use the same algorithm, namely maximal disc thinning [59] based on the skeleton membership criteria of Remy and Thiel [61]. This method guarantees that the skeleton has a one pixel width and preserves the homotopy of the original shape. Both for the skeleton computation and the reconstruction we require distance maps, which we compute with the method of Meijster et al. [72]. For the reconstruction itself, we rely on a simple disc drawing algorithm according to Definition 4. A sample implementation is available on Zenodo [73].

4.1.2 Data Sets

As an input to our algorithm, we require binary shapes. Therefore, we consider two test datasets: BSDS500 [74] contains ground truth segmentation data for real-world images. This database acts as an example for a typical pipeline in which our scale-space framework would constitute an intermediate step. Before our algorithm is deployed, a segmentation approach extracts a binary shape from a real-world image (see Fig. 2). Our skeleton would then typically feed into other

Fig. 2 We use test images from two databases. **a** With one example from the BSDS500 [74] data set, we illustrate a use case where skeletonisation is part of a pipeline that extracts binary shapes from real-world images. **b** The 1400 images of CE-Shape-1 [69] are the foundation for our quantitative evaluation. Here, we show four examples to illustrate typical shapes. They vary significantly in motive, skeleton complexity and image resolution (Color figure online)



algorithms, e.g. for shape matching. For our visual comparisons, we have extracted individual objects from the ground truth segmentations for our visual comparison.

For our quantitative experiments, we rely on the MPEG7 core experiment CE-Shape-1 database [69] instead. It contains a variety of different classes of synthetic shape images that can be directly used without the need for further preprocessing.

4.1.3 Quality Criteria

An obvious quality measure for skeletons is the reconstruction error as the sum of false negative and false positives in a discrete reconstruction of the initial object O . However, there are a few more quality criteria that are useful for our task-specific scale-spaces [59].

The *reconstruction error* $\mathcal{E}_\ell$ is the number of missing object points

$$\mathcal{E}_\ell := |O_0| - |O_\ell|. \quad (9)$$

Note that this error only considers false negatives in the object representation. False positives cannot occur by definition since we only use this measure for reconstructions from subsets of the full skeleton.

Skeleton minimality relates skeleton points $|\Sigma_\ell|$ to the number $|O_\ell|$ reconstructed image points.

$$\mathcal{M}_\ell := \frac{|\Sigma_\ell|}{|O_\ell|}. \quad (10)$$

Skeleton complexity is the total number of significant points in the skeleton:

$$\mathcal{C}_\ell := |E(\Sigma_\ell)| + |B(\Sigma_\ell)|. \quad (11)$$

This definition relies on the notions of endpoints E and branching points B (see Definition 5).

4.1.4 Baseline: Random Sparsification

As a baseline for our task-specific scale-spaces we consider a sparsification path that at each scale ℓ selects a point $x \in \Sigma_\ell$ uniformly at random. This sparsification path is completely independent of the structure of the object and does not follow any goals. Following the taxonomy of Cárdenas et al. [6], we call such a scale-space *uncommitted*. All our application-specific scale-spaces take the shape or topology of the original object into account and are therefore *committed*.

Note that even this simple uncommitted sparsification path yields a scale-space that obeys the properties D1-D6 from Sect. 3.3.

4.2 Skeletonisation Scale-Spaces for Branch Pruning

First, we investigate classical skeleton pruning [9–14] as a special case of our sparsification framework. Since this is often a preprocessing step for applications such as shape matching, the homotopy preservation property S2 from Sect. 2.1 is important here. The connectivity of the object should be preserved. This constitutes a constraint for our sparsification path.

In addition, end- and branching points from Definition 5 are semantically important for shape matching. Since endpoints and branching points mark where individual arcs of the skeleton intersect or end, they are more relevant for shape recognition than simple points. Thus, skeleton complexity as defined in Sect. 4.1.3 is a good criterion for simplification.

Pruning also aims to remove so-called spurious branches that are the result of small boundary perturbations. These

Algorithm 1: Branch Pruning Path

```

1  $\ell \leftarrow 0$ ;
2 while  $|\Sigma_\ell| > 0$  do
    /* Consider branches with endpoints or
       full skeleton. */
3    $A_\mathcal{E} \leftarrow \{A \in \mathcal{A}(\Sigma_\ell) \mid E(\Sigma_\ell) \cap A \neq \emptyset\}$ , if  $A_\mathcal{E} = \emptyset$  then
      $A_\mathcal{E} \leftarrow \{\Sigma_\ell\}$ ;
    /* Select branch with smallest
       reconstruction impact. */
4    $A_{\min} \leftarrow \operatorname{argmin}_{A \in A_\mathcal{E}} |I_{\ell,A}|$ ;
    /* Remove whole branch excluding
       branching points. */
5    $P_{\ell+1} \leftarrow \{a_i \in A_{\min} \mid a_i \notin B(\Sigma_\ell)\}$ ,  $\Sigma_{\ell+1} \leftarrow \Sigma_\ell \setminus P_{\ell+1}$ ,
      $\ell \leftarrow \ell + 1$ ;

```

can have a large effect on skeleton minimality since they can increase the size of the skeleton significantly without contributing much to the object reconstruction. We formalise this in terms of the reconstruction impact.

Definition 13 (*Reconstruction Impact*) For a set of skeleton points $S \subseteq \Sigma_\ell$ and $\ell \in \{0, \dots, m\}$ we define the *reconstruction impact* $I_{\ell,S}$ by

$$I_{\ell,S} := O_\ell \setminus \left(\bigcup_{i \in \Sigma_\ell \setminus S} B_{D(i)}(i) \right). \quad (12)$$

Thus, our goals for designing a reconstruction path is removing whole branches of the skeleton without disconnecting it and avoiding increases in skeleton complexity. If we have multiple choices for branches to remove, we want to choose the ones with the least reconstruction impact first, since those are most likely to arise from boundary perturbations. We incorporate these goals into our branch pruning path described by Algorithm 1. In the following, we show that the additional requirements of this pruning path lead to skeleton complexity as a new Lyapunov sequence for our task-specific scale-space.

Proposition 4 (Skeleton Complexity is a Lyapunov Sequence) *Skeleton complexity decreases with increasing scale ℓ , i.e. $\mathcal{C}_\ell \geq \mathcal{C}_{\ell+1}$.*

Proof For a scale $\ell \in \{0, \dots, m-1\}$, there are the following possible cases for the next step $P_{\ell+1}$ in the sparsification path:

Case1 : $P_{\ell+1}$ has exactly two endpoints, which are both removed. The branch is not connected to other skeleton components and thus cannot influence other points. Then $\mathcal{C}_{\ell+1} = \mathcal{C}_\ell - 2 < \mathcal{C}_\ell$.

Case2 : By Algorithm 1, $P_{\ell+1}$ has exactly one endpoint, which is removed. The arc is connected to a branching point. Its role can change to an end or simple point, but

Algorithm 2: Skeleton Compression Path

```

1  $\ell \leftarrow 0$ ,  $r \in \mathbb{N}$  user parameter (points to remove per step) ;
2 while  $|\Sigma_\ell| > 0$  do
3    $k \leftarrow |\Sigma_\ell|$ ,  $s \leftarrow \min(r, |\Sigma_\ell|)$ ;
    /* Order all skeleton points by
       reconstruction influence. */
4   Let  $\{c_1, \dots, c_k\} = \Sigma_\ell$  with  $|I_{\ell,c_i}| \leq |I_{\ell,c_j}|$  for  $i \leq j$ ;
    /* Remove  $s$  points with smallest impact
       on reconstruction. */
5    $P_{\ell+1} \leftarrow \{c_1, \dots, c_s\}$ ,  $\Sigma_{\ell+1} \leftarrow \Sigma_\ell \setminus P_{\ell+1}$ ,  $\ell \leftarrow \ell + 1$ ;

```

this cannot increase complexity. Thus, $\mathcal{C}_{\ell+1} \leq \mathcal{C}_\ell - 1 < \mathcal{C}_\ell$.

Case3 : $P_{\ell+1} = \Sigma_\ell$. The remaining skeleton is removed, thus $\mathcal{C}_{\ell+1} = 0 \leq \mathcal{C}_\ell$. □

4.3 Skeletonisation Scale-Spaces for Compression

For shape compression [43], homotopy preservation is not relevant. We merely want to represent the initial shape as accurately as possible with as few skeleton points as possible. More sophisticated versions of such a task-specific scale-space could also incorporate coding aspects such as the entropy of the remaining skeleton points, but this is beyond the scope of our proof of concept. We refer to Mhlhaus [43] for such considerations. Instead, we solely rely on the reconstruction impact from Definition 13. By removing S from the skeleton Σ_ℓ , we remove $I_{\ell,S}$ from the object and thus obtain in every step

$$O_{\ell+1} = O_\ell \setminus I_{\ell,S}. \quad (13)$$

Since we want to minimise a loss of object points for every skeleton part removed, we have to minimise $|I_{\ell,S}|$ in every step. Consequentially, our Algorithm 2 defines a skeleton compression path that minimises this error. Here, we use the shorthand notation $I_{\ell,i} := I_{\ell,\{i\}}$ for a given pixel index $i \in \Sigma$.

By design, the relative error constitutes a Lyapunov sequence as a corollary to Proposition 3 since $|O_\ell|$ is shrinking and thus the relative error is increasing with the scale. This is to be expected, but not a particularly useful property for a compression algorithm. However, we can also guarantee that the increase in error for every step is indeed minimal.

Proposition 5 (Minimal Relative Error Increase) *Among all possible sparsification paths, for all $\ell \in \{0, \dots, m-1\}$, the increase in relative error $\mathcal{E}_{\ell+1} - \mathcal{E}_\ell$ is minimal for a skeleton compression scale-space.*

According to Eq. (13), we have $|O_{\ell+1}| = |O_\ell| - |I_{\ell,P_{\ell+1}}|$. Moreover, Algorithm 2 selects the next step $P_{\ell+1}$ of the sparsification path such that $|I_{P_{\ell+1}}|$ is minimised over all possible

choices of $P_{\ell+1}$. Therefore, each step yields the minimal possible increase in reconstruction error.

Additionally, we want a small number of skeleton points to represent a large number of object points. This is reflected by the skeleton minimality from Sect. 4.1.3. It also constitutes a Lyapunov sequence for our compression scale-spaces.

Proposition 6 (Skeleton Minimality is a Lyapunov Sequence) *Skeleton minimality decreases with increasing scale ℓ , i.e. $\mathcal{M}_\ell \geq \mathcal{M}_{\ell+1}$.*

Proof First, we decompose $\mathcal{M}_\ell$ according to the sparsification path from Definition 10 and the reconstruction impact from Eq. (13). This yields

$$\mathcal{M}_\ell = \frac{|\Sigma_\ell|}{|O_\ell|} = \frac{|\Sigma_{\ell+1}| + |P_{\ell+1}|}{|O_{\ell+1}| + |I_{\ell, P_{\ell+1}}|}. \quad (14)$$

Furthermore, due to the sorting in Algorithm 2, the average unique area of the removed points in $P_{\ell+1}$ is smaller or equal to the average unique area of the remaining skeleton points in $\Sigma_{\ell+1}$, i.e.

$$\frac{|I_{\ell, P_{\ell+1}}|}{|P_{\ell+1}|} \leq \frac{|I_{\ell, \Sigma_{\ell+1}}|}{|\Sigma_{\ell+1}|} \leq \frac{|O_{\ell+1}|}{|\Sigma_{\ell+1}|}. \quad (15)$$

Combining both Eq. (14) and Eq. (15), we can show our claim by

$$|O_{\ell+1}| \geq \frac{|I_{\ell, P_{\ell+1}}| |\Sigma_{\ell+1}|}{|P_{\ell+1}|} \quad (16)$$

$$\iff |O_{\ell+1}| |P_{\ell+1}| \geq |I_{\ell, P_{\ell+1}}| |\Sigma_{\ell+1}| \quad (17)$$

$$\iff |\Sigma_{\ell+1}| |O_{\ell+1}| + |O_{\ell+1}| |P_{\ell+1}| \geq |I_{\ell, P_{\ell+1}}| |\Sigma_{\ell+1}| + |\Sigma_{\ell+1}| |O_{\ell+1}| \quad (18)$$

$$\iff |O_{\ell+1}| (|\Sigma_{\ell+1}| + |P_{\ell+1}|) \geq |\Sigma_{\ell+1}| (|I_{\ell, P_{\ell+1}}| + |O_{\ell+1}|) \quad (19)$$

$$\iff \frac{|\Sigma_{\ell+1}| + |P_{\ell+1}|}{|O_{\ell+1}| + |I_{\ell, P_{\ell+1}}|} \geq \frac{|\Sigma_{\ell+1}|}{|O_{\ell+1}|} \quad (20)$$

$$\stackrel{(14)}{\iff} \mathcal{M}_\ell \geq \mathcal{M}_{\ell+1}. \quad (21)$$

□

4.3.1 A Note on Densification for Compression

Note that we have also implemented a densification alternative to this compression scale-space. The algorithm starts with an empty skeleton and greedily adds the skeleton point in every step that has the largest unique impact on the reconstructed object. This yields a similar trade-off between skeleton size and reconstruction error as the sparsification. Quality-wise, both approaches are on par.

However, in this case, densification is inherently a worse choice due to its prohibitive runtime. We implement both

sparsification and densification with caching techniques that avoid costly re-computations of the unique influence as much as possible. However, there is an inherent difference between both approaches: In the top-down sparsification approach, we benefit from the fact that in the beginning when the skeleton is highly populated, the local influence is small. As it begins to increase, the number of candidates that need to be updated shrinks.

In densification, this relationship is reversed. It is much more costly to update the influence of many potential candidates when the skeleton is still sparse and the influence area large. We have observed multiple orders of magnitude higher runtimes for densification in our experiments. Therefore, we do not recommend it in practice and exclude it from our more detailed evaluations.

4.4 Comparative Evaluation of Sparsification


In the following, we compare our branch pruning and compression scale-spaces against the random sparsification baseline and investigate if they fulfil the goals we set for each application.

4.4.1 Branch Pruning

Figure 3 shows a visual example of branch pruning contrasted to random and compressive sparsification. As the example confirms, the topology is preserved as intended. The *kangaroo* object has a hole in it and therefore, the skeleton contains a loop. Since we aim to reduce skeleton complexity, branches with endpoints are prioritised in the pruning: If only a single branch connected to a loop is left, removing the branch reduces complexity to zero while removing the loop turns one branching point into an end point and keeps minimality constant at two. Both steps would adhere to skeleton minimality as the Lyapunov sequence, but we try to greedily decrease complexity. Thus, as expected, the loop is the last surviving structure of the skeleton. Other branches are pruned one by one according to their impact on the reconstruction. Even with only two branches left, a large part of the shape is retained.

A clear trade-off can be seen in this example. The preservation of homotopy comes at the price of a higher reconstruction error compared to the random or compression sparsification. On the other hand, skeleton complexity is significantly lower for the branch pruning. For applications like shape matching, these metrics are much more important than an accurate reconstruction of the original shape. Therefore, our scale-space is indeed task-adaptive.

Fig. 3 Comparison of sparsification paths on image 69020 of BSDS500 [74] grouped by branch pruning steps. Endpoints are marked in red, branch points in blue. $\mathcal{E}$ denotes the reconstruction error and $\mathcal{C}$ skeleton complexity. Branch pruning preserves homotopy and reduces complexity at the cost of a higher reconstruction error compared to random or compressive sparsification (Color figure online)



4.4.2 Compression

In Figs. 3 and 4, we see that our compression scale-space outperforms both the sparsification baseline and the branch pruning in terms of reconstruction quality. Here, we are only interested in getting the best trade-off between the number of skeleton points that we need to store and the reconstruction error. Lifting the homotopy requirement and prioritising skeleton minimality over skeleton complexity allows superior performance in this regard. This is also confirmed by the quantitative results in Fig. 5. We can see that compression scale-space lead to an error reduction of up to 80% compared to branch pruning and up to 50% compared to the baseline at the same skeleton size.

4.4.3 Quantitative Evaluation

We evaluate reconstruction error $\mathcal{E}$ and skeleton complexity $\mathcal{C}$ for our three sparsification scale-spaces on the CE-Shape-1 database [69]. Since the objects and their associated skeletons have highly variable sizes, the number of discrete steps in scale-space varies as well. Therefore, we cannot average errors over the discrete scale. Instead, we normalise scale to a 0 – 100 percentage scale.

The evolution of the error in Fig. 5a is consistent with the visual evaluation from the previous sections. Branch pruning trades homotopy and monotonically decreasing skeleton complexity for a relatively high reconstruction error. In contrast, the greedy optimisation of our compression scale-space guarantees the minimum increase in the error for every step in scale-space. This is verified by its quantitative performance. It is particularly effective for very coarse scales, where the choices of the preserved skeleton points matter most.

In the case of pruning, the graph of the skeleton complexity in Fig. 5b verifies the expectations that we have for a Lyapunov sequence. It is a meaningful measure for shape matching and declines monotonically. Both compression and random sparsification feature an initial increase in complexity followed by a decline for coarse scales. Since they are not homotopy-preserving, both of these scale-spaces disconnect the skeleton early on, thus creating new end points. However, random sparsification does this without distinction of skeleton points, thus leading to a slightly distorted bell curve. The increase and decline of complexity for compression however is linear. This arises from the fact that there is a certain correlation between reconstruction impact and skeleton complexity: Compression removes the skeleton points with the least impact first. Neighbouring skeleton points often have similar disc radii as well, thus their reconstruction discs are

Fig. 4 Comparison of sparsification paths on image 69020 of BSDS500 [74] with equidistant scale increase. Compared to random sparsification, the compression scale-space removes only redundant skeleton points first and thus does not increase the reconstruction error $\mathcal{E}$ for the first five scales displayed. In general, it offers a quality improvement of up to 50% (Color figure online)





Fig. 5 Quantitative evaluation on the 1400 images of the CE-Shape-1 database [69]. **a** Branch pruning has the highest error due to its preservation of homotopy and reduction of complexity. Compression minimises the reconstruction error and thus performs best. **b** Branch pruning performs best in terms of complexity since it minimises it greedily. Compression implicitly limits the improve in complexity as

well since it removes the least significant points first which often form connected branches in the skeleton. The coloured shading represents the ± 1 standard deviation band. All methods use 100 bins. Due to removal of whole branches these are unevenly populated for pruning. We reduce the noise with a Gaussian filter ($\sigma = 1.5$) applied to all methods (Color figure online)

Algorithm 3: Stiffness Enhancement

```

/* Initialisation */
1  $\ell \leftarrow 0$ ;
2  $A_0 \leftarrow O$ ;
3  $\Gamma_0 \leftarrow \Sigma$ ;
4 while  $|\Gamma_\ell| < |O|$  do
    /* Strengthen current skeleton by
       dilation. */
    /*  $\mathcal{N}(\mathbf{x})$ : 8-neighbourhood of  $\mathbf{x}$  */
5    $\hat{\Gamma}_\ell \leftarrow \Gamma_\ell \cup \{\mathcal{N}(\mathbf{x}) \mid \mathbf{x} \in \Gamma_\ell\}$ ;
    /* Update auxiliary shape by subtracting
       dilated skeleton. */
6    $A_{\ell+1} \leftarrow A_\ell \setminus \hat{\Gamma}_\ell$ ;
    /* Add skeleton of current auxiliary
       shape  $A_\ell$  to evolving generalised
       skeleton. */
7    $\Gamma_{\ell+1} \leftarrow \hat{\Gamma}_\ell \cup SA_\ell$ ;
8    $\ell \leftarrow \ell + 1$ ;

```

heavily overlapping. Therefore, if a skeleton point is insignificant, there is a high chance of its neighbours also being insignificant. Thus, it is likely that whole connected branches are removed first.

Overall, the quantitative evaluation supports our visual observations.

4.5 Stiffness Enhancement with Extended Densification Scale-Spaces

In the following, we demonstrate that our extended densification from Sect. 3.5.2 allows us to perform stiffening [28]. Given the outline of an object to be printed with a 3-D printer, an iterative skeletonisation can improve the stability of the printed object. They have validated that such an approach actually improves the elastic stress properties.

Our highly generalised Definition 12 is flexible enough to cover applications that go far beyond a standard skeletonisation. In Algorithm 3, our first iteration simply adds the original skeleton of the shape to our extended evolving skeleton set Γ_ℓ . This forms the basis for the printed structures. Then, in every step, we first strengthen the pre-existing structure by a one pixel dilation: We thicken the skeleton by adding all direct non-skeleton neighbours. This is a step that highlights the generalisation capabilities of our framework: We

can intentionally sacrifice the thin skeleton property in order to adopt the stiffness enhancement idea that relies on thickening main support structures. Then, we subtract this skeleton from the object, creating an evolving auxiliary object A_ℓ . To this auxiliary object, we apply the skeletonisation algorithm again. By repeating these steps, in the transition from every coarse to the next finer scale, we strengthen the existing support structures and add new thin support beams given by the newly computed skeleton.


Figure 6 shows two examples of stiffness enhancement. As intended, we populate the extended skeleton with new support branches and iteratively strengthen already existing branches.

5 Conclusions and Outlook

We leverage a combination of data optimisation strategies [6, 62–68] and the medial axis transform [8] to establish a new class of scale-space frameworks for shape analysis. With only a few requirements, they offer a broad range of architectural and simplification properties. Compared to many other scale-spaces, they have no ill-posed direction: sparsification and densification scale-spaces are equally well-posed counterparts that allow two approaches for transitioning between coarse and fine shape representations in opposite directions. Together with the sparsification and densification paths as flexible design tools, our framework allows a wide range of practical applications. With proof-of-concept examples from the areas of skeleton post-processing, shape compression, and 3-D printing, we have demonstrated how additional requirements for the scale-space paths yield task-specific theoretical guarantees.

In the future, we plan to explore further practical applications. In particular, we want to investigate scale-spaces for compression that also take coding cost into account, similar to quantisation scale-spaces [65].

Fig. 6 Stiffness enhancement on two images of the CE-Shape-1 database [69]. Each finer scale adds more branches to the extended skeleton and thickens the previous skeleton



Appendix A Maximum Disc Thinning

For the discrete experiments in this paper we employ the maximal disc thinning (MDT) to compute the initial skeletons. As the concepts of the object O , the distance transform D , and homotopic thinning have already been introduced in Sect. 2, we only summarise the algorithmic procedure here. A detailed description and evaluation can be found in [59].

The method extends upon the original approach by Remy and Thiel [61] which proposes to consider the centres of inscribed discs with maximal radius as candidates for the skeleton. However, MDT also incorporates the original grass-fire idea of Blum [8] in a level set paradigm: From the boundary inwards, it considers sets of points that have equal values in the distance transform D (i.e. D is increasing with each iteration).

All object points that are not skeleton candidates are removed in each iteration, unless their removal destroys the homotopy of the remaining shape. To avoid spurious branches, isolated skeleton candidates are pruned if they have no neighbouring candidates in a 5×5 window.

Algorithm 4 contains the pseudocode of the MDT algorithm adapted from [59]. For the definition of simple and end points, please refer to Definition 5. The helper function *isSimple* and *isEndpoint* check the criteria of this definition. Figure 7 illustrates MDT thinning on a toy example.

Algorithm 4: Maximal Disc Thinning

```

/* Inputs (precomputed): order (cached
   ordering values according to  $D$ ), ep
   (endpoints identified with the RT
   algorithm) */
/* hp: min-heap ordered by order[],  $N_8, N_{24}$ :
   8- and  $5 \times 5$ -neighbourhoods */
/* Prune spurious endpoint candidates. */
1 foreach  $p \in O$  do
2   if  $|\{x \in N_{24}(p) \mid ep[x] = \text{TRUE}\}| < 2$  then
3      $ep[p] \leftarrow \text{FALSE}$ ;

/* Initialise heap with boundary points
   marked for deletion. */
4 foreach  $p \in \partial O$  do
5   if isSimple( $p$ ) then
6     insert  $p$  into  $hp$  with key  $order[p]$ ;

/* Homotopy preserving thinning with
   maximal-disk endpoint protection */
7 while  $|hp| > 0$  do
8    $p \leftarrow \text{popmin}(hp)$ ;
9   if isSimple( $p$ ) and ( $\neg \text{isEndpoint}(p)$  or  $ep[p] = \text{FALSE}$ )
   then
10    remove  $p$  from  $O$ ;
11    foreach  $q \in N_8(p)$  do
12      if isSimple( $q$ ) then
13        insert  $q$  into  $hp$  with key  $order[q]$ ;

```




Fig. 7 Toy illustration of thinning on a rectangle. The full rectangle of step 1 is reduced to its thin discrete skeleton in two steps. In each step, we remove one level set of non-skeleton points as in the MDT algorithm

Author Contributions Both authors have contributed equally to both the theoretical content and practical implementations.

Funding Open Access funding enabled and organized by Projekt DEAL.

Data Availability The data sets generated during and/or analysed during the current study are available from the corresponding author on reasonable request.

Declarations

Competing Interests The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Alvarez, L., Guichard, F., Lions, P.-L., Morel, J.-M.: Axioms and fundamental equations in image processing. *Arch. Ration. Mech. Anal.* **123**, 199–257 (1993)
- Iijima, T.: Basic theory on normalization of pattern (in case of typical one-dimensional pattern). *Bull. Electrotech. Lab* **26**, 368–388 (1962). (In Japanese)
- Lindeberg, T.: Generalized Gaussian scale-space axiomatics comprising linear scale-space, affine scale-space and spatio-temporal scale-space. *J. Math. Imaging Vis.* **40**, 36–81 (2011)
- Scherzer, O., Weickert, J.: Relations between regularization and diffusion filtering. *J. Math. Imaging Vis.* **12**(1), 43–63 (2000)
- Weickert, J.: *Anisotropic Diffusion in Image Processing*. Teubner, Stuttgart (1998)
- Cárdenas, M., Peter, P., Weickert, J.: Sparsification scale-spaces. In: Lellmann, J., Burger, M., Modersitzki, J. (eds.) *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science*, vol. 11603, pp. 303–314. Springer, Cham (2019)
- Siddiqi, K., Pizer, S.: Medial Representations: Mathematics, Algorithms and Applications. *Computational Imaging and Vision*, vol. 37, Springer, Dordrecht (2008)
- Blum, H.: A transformation for extracting new descriptors of shape. In: Wathen-Dunn, W. (ed.) *Models for the Perception of Speech and Visual Form*, pp. 362–380. MIT press, Cambridge, MA (1967)
- Bai, X., Latecki, L.J., Liu, W.-Y.: Skeleton pruning by contour partitioning with discrete curve evolution. *IEEE Trans. Pattern Anal. Mach. Intell.* **29**(3), 449–462 (2007)
- Montero, A.S., Lang, J.: Skeleton pruning by contour approximation and the integer medial axis transform. *Comput. Graph.* **36**(5), 477–487 (2012)
- Ogniewicz, R.L.: Automatic medial axis pruning based on characteristics of the skeleton-space. In: *Shape. Structure and Pattern Recognition - Proceedings of 1994 IAPR International Workshop*, pp. 23–32. Nahariya, Israel (1994)
- Shaked, D., Bruckstein, A.: Pruning medial axes. *Comput. Vis. Image Underst.* **69**(2), 156–169 (1998)
- Serino, L., Sanniti di Baja, G.: A new strategy for skeleton pruning. *Pattern Recognit. Lett.* **76**, 41–48 (2016)
- Tam, R., Heidrich, W.: Feature-preserving medial axis noise removal. In: *Proc. 7th European Conference on Computer Vision (ECCV)*, Part II, pp. 672–686., Copenhagen, Denmark (2002)
- Brockett, R.W., Maragos, P.: Evolution equations for continuous-scale morphology. In: *Proceedings of 1992 IEEE International Conference on Acoustics, Speech and Signal Processing*, vol. 3, pp. 125–128., San Francisco, CA (1992)
- Caselles, V., Sbert, C.: What is the best causal scale space for three-dimensional images? *SIAM J. Appl. Math.* **56**(4), 1199–1246 (1996)
- Kimia, B.B., Siddiqi, K.: Geometric heat equation and non-linear diffusion of shapes and images. *Comput. Vis. Image Underst.* **64**, 305–322 (1996)
- Sapiro, G., Tannenbaum, A.: Affine invariant scale-space. *Int. J. Comput. Vis.* **11**, 25–44 (1993)
- van den Boomgaard, R., Smeulders, A.: The morphological structure of images: the differential equations of morphological scale-space. *IEEE Trans. Pattern Anal. Mach. Intell.* **16**, 1101–1113 (1994)
- Soille, P.: *Morphological Image Analysis*, 2nd edn. Springer, Berlin (2003)
- Pizer, S.M., Siddiqi, K., Székely, G., Damon, J.N., Zucker, S.W.: Multiscale medial loci and their properties. *Int. J. Comput. Vis.* **55**, 155–179 (2003)
- Saha, P.K., Borgefors, G., Sanniti di Baja, G.: A survey on skeletonization algorithms and their applications. *Pattern Recognit. Lett.* **76**(C), 3–12 (2016)
- Saha, P.K., Borgefors, G., Sanniti di Baja, G.: Skeletonization and its applications—a review. In: Saha, P.K., Borgefors, G., Sanniti di Baja, G. (eds.) *Skeletonization*, pp. 3–42. Academic Press, London, UK (2017)
- Tsogkas, S., Kokkinos, I.: Learning-based symmetry detection in natural images. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Saito, Y., Schmid, C. (eds.) *Computer Vision - ECCV 2012. Lecture Notes in Computer Science*, vol. 7574, pp. 41–54. Springer, Berlin (2012)
- Pizer, S.M., Burbeck, C.A., Coggins, J.M., Fritsch, D.S., Morse, B.S.: Object shape before boundary shape: scale-space medial axes. *J. Math. Imaging Vis.* **4**(3), 303–313 (1994)

26. Pizer, S.M., Eberly, D., Fritsch, D.S., Morse, B.S.: Zoom-invariant vision of figural shape: the mathematics of cores. *Comput. Vis. Image Underst.* **69**(1), 55–71 (1998)
27. Gierke, J., Peter, P.: Skeletonisation scale-spaces. In: Bubba, T., Gaburro, R., Gazzola, S., Papafitsoros, K., Pereyra, M., Sch nlieb, C. (eds.) *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science*, vol. 15668, pp. 166–178. Springer, Cham (2025)
28. Breuß, M., Buhl, J., Yarahmadi, A.M., Bambach, M., Peter, P.: A simple approach to stiffness enhancement of a printable shape by Hamilton-Jacobi skeletonization. *Procedia Manuf.* **47**, 1190–1196 (2020)
29. Hajłasz, P.: On an old theorem of Erdős about ambiguous locus. *Colloq. Math.* **168**(1), (2022)
30. Erdős, P.: On the Hausdorff dimension of some sets in Euclidean space. *Bull. Am. Math. Soc.* **52**(2), 107–109 (1946)
31. Giblin, P., Kimia, B.: On the intrinsic reconstruction of shape from its symmetries. *IEEE Trans. Pattern Anal. Mach. Intell.* **25**(7), 895–911 (2003)
32. Sethian, J.A.: *Level Set Methods*. Cambridge University Press, Cambridge, UK (1996)
33. Rouy, E., Tourin, A.: A viscosity solutions approach to shape-from-shading. *SIAM J. Numer. Anal.* **29**, 867–884 (1992)
34. Siddiqi, K., Bouix, S., Tannenbaum, A., Zucker, S.W.: Hamilton-Jacobi skeletons. *Int. J. Comput. Vis.* **48**(3), 215–231 (2002)
35. Kuijper, A., Olsen, O.F.: Essential loops and their relevance for skeletons and symmetry sets. In: Olsen, O.F., Florack, A.K.L. (eds.) *Deep Structure, Singularities, and Computer Vision. Lecture Notes in Computer Science*, vol. 3753, pp. 24–35. Springer, Berlin (2005)
36. Damon, J.: Determining the geometry of boundaries of objects from medial data. *Int. J. Comput. Vis.* **63**(1), 45–64 (2005)
37. Bruce, J., Giblin, P., Gibson, C.G.: Symmetry sets. *Proc. R. Soc. Lond. Ser. A* **101**(1–2), 163–186 (1985)
38. Tagliasacchi, A., Delame, T., Spagnuolo, M., Amenta, N., Telea, A.: 3D skeletons: a state-of-the-art report. *Comput. Graph. Forum* **35**(2), 573–597 (2016)
39. Giblin, P.J., Kimia, B.B.: On the local form and transitions of symmetry sets, medial axes, and shocks. *Int. J. Comput. Vis.* **54**, 143–157 (2003)
40. Kimia, B., Tannenbaum, A., Zucker, S.: Shapes, shocks, and deformations I: the components of two-dimensional shape and the reaction-diffusion space. *Int. J. Comput. Vis.* **15**(3), 189–224 (1995)
41. Opiyo, S., Okinda, C., Zhou, J., Mwangi, E., Makange, N.: Medial axis-based machine-vision system for orchard robot navigation. *Comput. Electron. Agric.* **185**, 106153 (2021)
42. Li, X., Song, S., Yao, J., Zhang, H., Zhou, R., Hong, Q.: Efficient collision detection using hybrid medial axis transform and BVH for rigid body simulation. *Graph. Models* **128**, 101180 (2023)
43. Mhlhaus, M.: *Compressing binary images with the medial axis transform*. Bachelor's thesis, Faculty of Mathematics and Computer Science, Saarland University (2020)
44. Mayer, J., Wartzack, S.: Computational geometry reconstruction from 3d topology optimization results: a new parametric approach by the medial axis. *Computer-Aided Des. Appl.* **20**, 960–975 (2023)
45. Witkin, A.P.: Scale-space filtering. In: *Proceedings of Eighth International Joint Conference on Artificial Intelligence*, vol. 2, pp. 945–951., Karlsruhe, Germany (1983)
46. Lindeberg, T.: *Scale-Space Theory in Computer Vision*. Kluwer, Boston (1994)
47. Sporring, J., Nielsen, M., Florack, L.: *Gaussian Scale-Space Theory. Computational Imaging and Vision, Presented at the* (1997)
48. Florack, L.: *Image Structure. Computational Imaging and Vision*, vol. 10. Springer, Dordrecht (2013)
49. Duits, R., Florack, L., de Graaf, J., ter Haar Romeny, B.: On the axioms of scale space theory. *J. Math. Imaging Vis.* **20**, 267–298 (2004)
50. Schmidt, M., Weickert, J.: Morphological counterparts of linear shift-invariant scale-spaces. *J. Math. Imaging Vis.* **56**(2), 352–366 (2016)
51. Felsberg, M., Sommer, G.: Scale-adaptive filtering derived from the Laplace equation. In: Radig, B., Florczyk, S. (eds.) *Pattern Recognition. Lecture Notes in Computer Science*, vol. 2032, pp. 95–106. Springer, Berlin (2001)
52. Burgeth, B., Didas, S., Weickert, J.: Relativistic scale-spaces. In: Kimmel, R., Sochen, N., Weickert, J. (eds.) *Scale Space and PDE Methods in Computer Vision. Lecture Notes in Computer Science*, vol. 3459, pp. 1–12. Springer, Berlin (2005)
53. Andreu, F., Ballester, C., Caselles, V., Mazón, J.M.: Minimizing total variation flow. *Differ. Integral Equ.* **14**(3), 321–360 (2001)
54. Bellettini, G., Caselles, V., Novaga, M.: The total variation flow in R^N . *J. Differ. Equ.* **184**(2), 475–525 (2002)
55. Wirth, B.K.J.: *Variational methods in shape space*, Rheinische Friedrich-Wilhelms-Universität Bonn (2010). (PhD thesis)
56. Zhu, Y., Seneviratne, L.D., Earles, S.W.E.: A fast boundary based thinning algorithm. In: *Proceedings of IAPR Workshop on Machine Vision Applications*, pp. 548–551., Kawasaki, Japan (1994)
57. Palagyi, K., Nemeth, G.: Fully parallel 3D thinning algorithms based on sufficient conditions for topology preservation. In: *Proceedings of 15th IAPR International Conference on Discrete Geometry for Computer Imagery*, pp. 481–492., Montreal, Canada (2009)
58. Pudney, C.: Distance-ordered homotopic thinning: a skeletonization algorithm for 3d digital images. *Comput. Vis. Image Underst.* **72**(3), 404–413 (1998)
59. Peter, P., Breuß, M.: Refined homotopic thinning algorithms and quality measures for skeletonisation methods. In: Breuß, M., Bruckstein, A., Maragos, P. (eds.) *Innovations for Shape Analysis: Models and Algorithms. Mathematics and Visualization*, pp. 77–92. Springer, Berlin (2013)
60. Malandain, G., Fernandez-Vidal, S.: Distance-ordered homotopic thinning: a skeletonization algorithm for 3D digital images. *Image Vis. Comput.* **16**(5), 317–327 (1998)
61. Remy, E., Thiel, E.: Exact medial axis with Euclidean distance. *Image Vis. Comput.* **23**(2), 167–175 (2005)
62. Mainberger, M., Hoffmann, S., Weickert, J., Tang, C.H., Johannsen, D., Neumann, F., Doerr, B.: Optimising spatial and tonal data for homogeneous diffusion inpainting. In: Bruckstein, A.M., Haar Romeny, B., Bronstein, A.M., Bronstein, M.M. (eds.) *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science*, vol. 6667, pp. 26–37. Springer, Berlin (2012)
63. Hoeltgen, L., Mainberger, M., Hoffmann, S., Weickert, J., Tang, C.H., Setzer, S., Johannsen, D., Neumann, F., Doerr, B.: Optimising spatial and tonal data for PDE-based inpainting. In: Bergounioux, M., Peyré, G., Schnörr, C., Caillaud, J.-P., Haberhorn, T. (eds.) *Variational Methods in Imaging and Geometric Control. Radon Series on Computational and Applied Mathematics*, vol. 18, pp. 35–83. De Gruyter, Berlin (2017)
64. Demaret, L., Dyn, N., Iske, A.: Image compression by linear splines over adaptive triangulations. *Signal Process.* **86**(7), 1604–1616 (2006)
65. Peter, P.: Quantisation scale-spaces. In: Elmoataz, A., Fadili, J., Quéau, Y., Rabin, J., Simon, L. (eds.) *Scale Space and Variational Methods in Computer Vision. Lecture Notes in Computer Science*, vol. 12679, pp. 15–26. Springer, Cham (2021)
66. Chizhov, V., Weickert, J.: Efficient data optimisation for harmonic inpainting with finite elements. In: Tsapatoulis, N., Panayides, A., Theo, T., Lanitis, A., Pattichis, C. (eds.) *Computer Analysis of Images and Patterns. Part 2. Lecture Notes in Computer Science*, vol. 13053, pp. 432–441. Springer, Cham (2021)

67. Daropoulos, V., Augustin, M., Weickert, J.: Sparse inpainting with smoothed particle hydrodynamics. *SIAM J. Appl. Math.* **14**(4), 1669–1704 (2021)
68. Karos, L., Bheed, P., Peter, P., Weickert, J.: Optimising data for exemplar-based inpainting. In: Blanc-Talon, J., Helbert, D., Philips, W., Popescu, D., Scheunders, P. (eds.) *Advanced Concepts for Intelligent Vision Systems. Lecture Notes in Computer Science*, vol. 11182, pp. 547–558. Springer, Cham (2018)
69. Latecki, L., Lakämper, R., Eckhardt, U.: Shape descriptors for non-rigid shapes with a single closed contour. In: *Proc. 2000 Conference on Computer Vision and Pattern Recognition*, pp. 1424–1429., Hilton Head, SC (2000)
70. Pfaltz, J.L., Rosenfeld, A.: Computer representation of planar regions by their skeletons. *Commun. ACM* **10**(2), 119–122 (1967)
71. Lieutier, A.: Any open bounded subset of $\mathbb{R}^n$ has the same homotopy type than its medial axis. In: *Proc. 8th ACM Symposium on Solid Modeling and Applications*, Seattle, WA, pp. 65–75 (2003)
72. Meijster, A., Roerdink, J., Hesselink, W.: A general algorithm for computing distance transforms in linear time. In: Goutsias, J., Vincent, L., Bloomberg, D. (eds.) *Proc. 5th International Symposium on Mathematical Morphology and Its Applications to Image and Signal Processing. Computational Imaging and Vision*, vol. 18, pp. 331–340. Springer, Berlin (2000)
73. Gierke, J., Peter, P.: Code for: skeleton sparsification and densification scale-spaces. Zenodo (2026). <https://doi.org/10.5281/zenodo.18982752>
74. Arbelaez, P., Maire, M., Fowlkes, C., Malik, J.: Contour detection and hierarchical image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **33**(5), 898–916 (2011)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.