



HAL
open science

Terrain Amplification using Multi-scale Erosion

Hugo Schott, Eric Galin, Eric Guérin, Adrien Peytavie, Axel Paris

► **To cite this version:**

Hugo Schott, Eric Galin, Eric Guérin, Adrien Peytavie, Axel Paris. Terrain Amplification using Multi-scale Erosion. 2024. hal-04565030

HAL Id: hal-04565030

<https://hal.science/hal-04565030>

Preprint submitted on 1 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Terrain Amplification using Multi-scale Erosion

HUGO SCHOTT, INSA Lyon, CNRS, LIRIS, UMR5205, F-69621 Villeurbanne, France

ERIC GALIN, Université Claude Bernard Lyon 1, CNRS, LIRIS, UMR5205, F-69622 Villeurbanne, France

ERIC GUÉRIN, INSA Lyon, CNRS, LIRIS, UMR5205, F-69621 Villeurbanne, France

ADRIEN PEYTAVIE, Université Claude Bernard Lyon 1, CNRS, LIRIS, UMR5205, F-69622 Villeurbanne, France

AXEL PARIS, Adobe




Fig. 1. Given an input low-resolution terrain, we automatically generate a hydrologically consistent high-resolution model augmented with details by applying multi-scale erosion processes. The landscape designer controls effects by prescribing the maximum amount of erosion, the hardness map to amplify or dampen carving, and may preserve the elevation of specific regions such as peaks or ridge lines by smoothly re-targeting the elevation.

Modeling high-resolution terrains is a perennial challenge in the creation of virtual worlds. In this paper, we focus on the amplification of a low-resolution input terrain into a high-resolution, hydrologically consistent terrain featuring complex patterns by a multi-scale approach. Our framework combines the best of both worlds, relying on physics-inspired erosion models producing consistent erosion landmarks and introducing control at different scales, thus bridging the gap between physics-based erosion simulations and multi-scale procedural modeling. The method uses a fast and accurate approximation of different simulations, including thermal, stream power erosion and deposition performed at different scales to obtain a range of effects. Our approach provides landscape designers with tools for amplifying mountain ranges and valleys with consistent details.

CCS Concepts: • **Computing methodologies** → **Shape modeling**.

Additional Key Words and Phrases: Erosion Simulation, Landscapes

ACM Reference Format:

Hugo Schott, Eric Galin, Eric Guérin, Adrien Peytavie, and Axel Paris. 2024. Terrain Amplification using Multi-scale Erosion. *ACM Trans. Graph.* 1, 1 (May 2024), 12 pages. https://doi.org/0000001.0000001_2

Authors' addresses: Hugo Schott, INSA Lyon, CNRS, LIRIS, UMR5205, F-69621 Villeurbanne, France, hugo.schott@liris.cnrs.fr; Eric Galin, Université Claude Bernard Lyon 1, CNRS, LIRIS, UMR5205, F-69622 Villeurbanne, France, eric.galin@liris.cnrs.fr; Eric Guérin, INSA Lyon, CNRS, LIRIS, UMR5205, F-69621 Villeurbanne, France, eric.guerin@liris.cnrs.fr; Adrien Peytavie, Université Claude Bernard Lyon 1, CNRS, LIRIS, UMR5205, F-69622 Villeurbanne, France, adrien.peytavie@liris.cnrs.fr; Axel Paris, Adobe, aparis@adobe.com.

1 INTRODUCTION

Landscapes are essential in virtual worlds, and high resolution terrain modeling remains a perennial challenge. This challenge stems not only from the diversity of relief, the variety of details at different scales, and the need for geomorphological and hydrological realism but also from controlling the shape and location of landforms to follow the designers' intent.

Landscape designers traditionally start with an initial low-resolution draft model organizing the principal landforms such as mountain ranges, valleys, canyons, and coastlines, by using procedural generation or combining patches extracted from digital elevation models. This initial draft is then processed by surface erosion algorithms as an enhancement post-processing. This step is not only crucial for synthesizing visually appealing and detailed terrain features, but it also guarantees a globally coherent terrain. Those algorithms generate visually convincing effects as long as the initial input terrain features realistic large-scale landmarks.

Our work comes from the observation that existing erosion simulations generate details at the original terrain resolution and often carve drainage patterns at the size of one cell (Figure 2). This limitation is particularly pronounced and noticeable for surface and physically-based hydraulic erosion models that create self-similar fractal or dendritic mountain ranges with river networks. Moreover, existing surface erosion algorithms are difficult to control, require precise and tedious parameter tuning, and generally do not adapt to different scales. Another crucial observation is that designers favor techniques that provide user control, a complex task to achieve using simulations. Even though some methods exist for amplifying terrains with details – in general through procedural and example-based, including deep learning-based multi-scale terrain amplification methods – little attention has been dedicated to amplifying low-resolution terrains using multi-scale erosion




Fig. 2. Traditional erosion simulations either yield exaggerated self-similarities across scales or only produce landforms at one resolution, particularly incisions at the size of one cell (left). In contrast, multi-scale amplification delivers a variety of effects at different scales.

We address these limitations by proposing a novel multi-scale approach for amplifying large-scale input terrains (see Figure 1). Our approach, inspired by previous work in geo-sciences, is in essence a physically-inspired procedural erosion method. We adapt and modify the erosion processes described in geomorphology, primarily focusing on designer control and interactive feedback. Fractal Brownian motion, implemented as sums of scaled-noise functions, has been used intensively in terrain modeling and amplification with the goal of reproducing naturally fractal landform patterns observed in nature. However, it fails to adapt to the underlying hierarchical structure. We fix this by including a multi-scale approach in the erosion process which consists in introducing wavelength and the intensity of erosion at different scales.

We propose a geomorphologically-inspired approach based on the computation of stream power erosion as a scalar field to slightly carve an initial terrain featuring the major landforms and control the amount of removed material. More specifically, our contributions are as follows: 1) We propose an original multi-scale erosion and amplification model that combines different types of simulations; 2) We present a multi-scale breaching algorithm for guaranteeing the final hydrological coherency of the amplified terrain and a diffusion-based re-targeting method for correcting the elevation of the terrain slightly damaged during the amplification; 3) We propose a parallel implementation that allows the user to control the amplification with interactive feedback. Images shown throughout the paper as well as the accompanying video show multi-scale amplification of 256×256 low resolution terrains up to 8192×8192 with a $\times 32$ amplification factor and demonstrate that our method is compatible with other terrain modeling techniques.

2 RELATED WORK

Terrain generation methods are often sorted into one of the following categories: procedural generation, texture synthesis and erosion simulation [Galín et al. 2019]. In this work, we focus on amplification techniques that enhance an input terrain by adding relief details to a previously computed topography. While some procedural and data-based techniques may amplify an initial low-resolution terrain into a high-resolution output, no multi-scale erosion simulation methods exist for augmenting a terrain with hydrologically consistent patterns and landforms.

Procedural terrain amplification often relies on fractal Brownian motion, reproducing the self-similarity across scales, traditionally calculated as a sum of noises at different frequencies and amplitudes [Ebert et al. 2002], to add surface details to an underlying

low-resolution terrain. Directional ridged and erosive noises [de Carpentier and Bidarra 2009] adapt surface details to the direction of the slope, however the employed ridge noise does not introduce convincing patterns. Some complex noise basis functions were developed for reproducing dendritic drainage patterns [Gaillard et al. 2019], however the recursive graph-based distance algorithm remains computationally intensive. A variant of Phasor noise [Grenier et al. 2024] was also proposed to generate high-resolution erosion patterns according to the gradient of the low-resolution terrain. Even though their properties can be modulated spatially to form a multi-fractal signal, most existing noise-based variants fail at synthesizing realistic landforms due to the lack of structure and topography foundation of noise-basis functions.

Example-based amplification methods combine patches extracted from real-world digital elevation models, under the hypothesis that realism can be achieved by combining locally geomorphologically correct real-world patches, which is often not sufficient for generating globally consistent landscapes. Set aside terrain generation techniques inspired by texture-based synthesis methods [Gain et al. 2015; Scott and Dodgson 2021; Tasse et al. 2014; Zhou et al. 2007], sparse representations [Guérin et al. 2016] have been successfully used for terrain amplification. Another recent category of techniques rely on deep learning to increase the resolution of an input terrain and synthesize consistent high resolution landmarks from an otherwise low-resolution model. Conditional Generative Adversarial Networks [Guérin et al. 2017; Perche et al. 2023; Zhang et al. 2022] operate with low-resolution maps, global style information, and local style maps as input, and discriminators capable of classifying different types of terrains. Another approach, specialized in style embedding, was proposed by Zhao *et al.* [2019] to insert the embedded theme into the generation process, the method can amplify a low-resolution input terrain into a high resolution with style variation. Recently, diffusion models were adapted to terrain synthesis [Lochner et al. 2023] with up to $\times 8$ super-resolution amplification in one pass.

Erosion methods can be broadly classified in two sub-categories: surface erosion algorithms and physics-based simulations. The first category typically simulates material detachment, transport and deposition [Musgrave et al. 1989] using grid-based processes to enhance procedurally generated terrain with sedimentary valleys and small-scale features such as gorges and ravines. A geological representation for modeling the strata of the bedrock was introduced by Roudier *et al.* [1993] that considers the characteristics of the different materials during the erosion process. These early approaches were improved in several ways, by computing the acceleration or deceleration of the fluid to erode the bedrock or deposit sediments [Neidhold et al. 2005], combining a shallow water simulation with hydraulic erosion [Benes 2007], or using Smoothed Particle Hydrodynamics [Krištof et al. 2009]. As fluid simulations are computationally intensive, implementations on graphics hardware were proposed to speed-up computations [Mei et al. 2007; Šfava et al. 2008; Vanek et al. 2011]. Surface erosion methods perform the simulations at a limited temporal and spatial scale and then implicitly and often inconsistently (as the equations of the corresponding phenomena are not linear) scale up the results to approximate large scale terrains. Those methods generate small-scale details as long as the




Fig. 3. Given an input low resolution terrain T_0 , we iteratively increase its resolution and apply different erosion models at every resolution to add landscape features. The amplification process is controlled by hardness maps that can be used to prescribe regions which should be protected from erosion, allow, or exaggerate effects. We eventually apply a diffusion-based retargeting to preserve the original elevation of salient features such as peaks, crests lines or saddles followed by a multi-scale breaching step to obtain a hydrologically consistent terrain.

initial input terrain is sufficiently realistic and supplies large-scale landmarks, which are often damaged by the simulation.

Physics-based simulations attempt to reproduce large-scale effects by taking into account the uplift of the bedrock balanced by different types of erosion described in geology, particularly the stream power erosion [Cordonnier et al. 2016] or glacial erosion [Cordonnier et al. 2023]. Contrary to previously described surface methods, they produce large scale mountain ranges with dendritic river networks without the need for an initial elevation. Control is indirectly achieved by prescribing or simulating the folds and faults and the characteristics of the different bedrock strata from the relative movement of the tectonic plates [Cordonnier et al. 2018] or by interactively authoring the uplift [Schott et al. 2023]. Controlling the generation of particular landforms remains difficult and the generated terrains exhibit realistic yet uniform fractal patterns, with exaggerated self similarities across scales.

In contrast, our method bridges the gap between the well-established hydrologically consistent stream power erosion and the need for multi-scale authoring of landform details. We borrow the stream power formalism from physics and propose a multi-scale framework for generating erosion landmarks at different scales. It combines the best of both worlds, allowing the user to control the size of details by prescribing the scale and intensity of effects, their location through user-defined bedrock-hardness maps, and avoids the damaging of features such as peaks or ridges using a retargeting process.

3 OVERVIEW

Without any loss of generality, we consider terrains T over a square domain Ω , with a physical size of $s \times s$ meters, and a resolution denoted as n with a corresponding cell size $c = s/n$. Our objective is to offer an erosion-based amplification process for terrains that preserves their size and increases resolution to introduce realistic erosion features at different scales.

Figure 3 presents an overview of the multi-scale amplification workflow. Given an input low-resolution terrain, we iteratively increase its resolution and apply different erosion models at every resolution to add landscape features. The amplification process is controlled by hardness maps that can be used to prescribe regions that should be protected from erosion, authorize, or embellish effects. Landforms can be composed through the initial terrain, the multi-scale augmentation, and the hardness field. As the terrain is carved and damaged by the erosion simulation, we propose a

retargeting post-processing step based on diffusion to allow the landscape designer to preserve the original relief of feature regions, such as ridges or peaks. Eventually, we correct the remaining hydrological inconsistencies by using a multi-scale breaching.

Inspired from the fractal sum of noises [Ebert et al. 2002] that iteratively amplify the terrain by introducing smaller details at decreasing wavelength, the terrain goes through a series of amplification operations $\mathcal{A}_k, k \geq 0$, with a decreasing intensity and characteristic size. From the input model T_0 , the multi-scale amplification composes a series of amplifications \mathcal{A}_k and generates a series of models T_k at increasing resolutions $n_m > \dots > n_0$ to a final terrain T_m with resolution n_m :

$$T_m = \mathcal{A}_{m-1} \circ \dots \circ \mathcal{A}_0(T_0)$$

Every amplification step \mathcal{A}_k is a single-scale operator composed of an upsampling step and three erosion-based processes:

$$T_{k+1} = \mathcal{A}_k(T_k) = \mathcal{D}_k \circ \mathcal{T}_k \circ \mathcal{E}_k \circ \mathcal{U}_k(T_k)$$

The bicubic upsampling \mathcal{U}_k produces a terrain $T_{k+1} = \mathcal{U}_k(T_k)$ with a new resolution $n_{k+1} > n_k$ without adding any details or specific features. In contrast, the fluvial erosion \mathcal{E}_k (Section 4.2), thermal stabilization \mathcal{T}_k (Section 4.3) and sediment deposition \mathcal{D}_k (Section 4.4) generate landforms and details at the new resolution n_{k+1} , *i.e.*, at the scale of the cell size c_{k+1} .

Unlike existing surface erosion algorithms, we separate detachment and sediment deposition into distinct processes. This choice comes from the observation that when coupled [Cordonnier et al. 2017], the layer of sediments produced by grid-based methods shields the bedrock, preventing detachment. While geomorphologically motivated, this approach is more complex to control and does not allow the user to observe and control the consequences of each separate type of erosion. Our strategy conforms to the original stream power erosion in geo-sciences [Braun and Willett 2013; Whipple and Tucker 1999] that does not consider sediment deposition. The underlying motivation is to obtain controllable stream power erosion \mathcal{E}_k , thermal stabilization \mathcal{T}_k , and deposition \mathcal{D}_k steps, whose action only impacts the elevation in one way (matter removal, slope thresholding, or deposition). The influence of every step is apparent to the user, who can achieve satisfying results faster, with fewer editing steps and simulation time.

The designer can adapt the parameters of every amplification step \mathcal{A}_k (upsampling factor n_{k+1}/n_k , stream power erosion and deposition intensities), which provides control at different scales. In

particular and unlike existing surface erosion methods, it is possible to define the maximum carving depth at every scale, allowing for user-control similar to scaled-noise synthesis (Section 6.2).

Even though the user may limit the amount of erosion, multi-scale amplification still damages the input terrain and may alter the elevation of key landform features such as peaks, saddles, or crest lines which play an important part in the terrain characteristics [Argudo et al. 2019]. Therefore, we introduce a diffusion-based retargeting step (Section 5.1) that corrects the height of peaks, ridges, or even regions of interest, leading to a retargeted terrain $\mathcal{R}(T_m)$. This step can be performed automatically by detecting the ridge positions and diffusing the elevation error around the points of interest.

Finally, although the stream power erosion steps \mathcal{E}_k improve the hydrological characteristics of the terrain at every iteration, the output terrain T_m and its corrected elevation $\mathcal{R}(T_m)$ may still exhibit hydrological inconsistencies. Inspired by the breaching algorithm introduced by Barnes *et al.* [2014], we propose a multi-scale algorithm (Section 5.2) to generate a final hydrologically correct terrain $\tilde{\mathcal{B}} \circ \mathcal{R}(T_m)$ with a consistent drainage. The originality of this method is to avoid carving unrealistic narrow canyons at the size c_m of a cell. Instead, our method controls the extent of incisions while still guaranteeing complete drainage consistency.

4 EROSION SIMULATIONS

Classical surface erosion techniques typically simulate bedrock detachment and sediment deposition processes in a unified framework, which makes control notoriously challenging for the user. In contrast, the stream power model in geomorphology [Braun and Willett 2013; Whipple and Tucker 1999] only simulates detachment and considers sediments transported outside the domain, therefore omitting the deposition process. While some recent advances combine detachment and deposition [Yuan et al. 2019], we take inspiration from the standard stream power erosion and choose to separate the different geological processes in three distinct steps, namely fluvial erosion, thermal stabilization, and sedimentation, performed separately and operating at different scales (see Figure 4). Instead of focusing on mass conservation, this separation allows us to control the impact of every process and provides landscape designers additional control to shape the terrain.

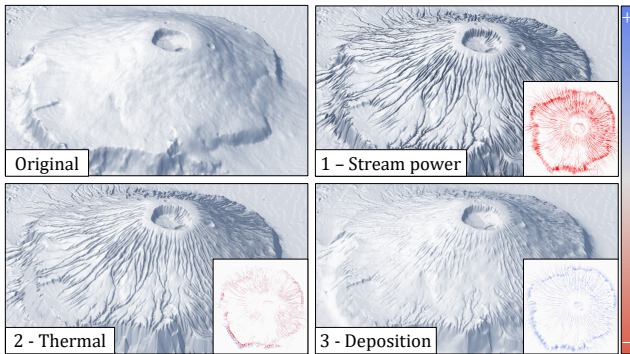


Fig. 4. Action of the three consecutive erosion algorithms, with the height difference at each step depicted in the inset.

Fluvial erosion is implemented as a modified version of the stream power erosion model [Cordonnier et al. 2016], and only removes material from the bedrock (Section 4.2). Thermal erosion, also referred to as debris slope erosion, stabilizes material based on its repose angle, moving matter in the slope direction (Section 4.3). Finally, sediment deposition approximates the flow and deposition of sediments over the terrain (Section 4.4). Separating the different geological processes and applying them one at a time allows for better user control and produces distinct landforms and details. The ordering sequence is also notable: thermal stabilization comes after the carving of fluvial erosion to smooth the resulting steep slopes. After obtaining satisfying drainage and river channels generated by fluvial and thermal erosion, the user may trigger the deposition step to accumulate sediment in the channels and form valleys.

At the core of these processes is the computation of the flow, based on the drainage area, detailed hereafter.

4.1 Flow routing

Computing the water flow (and suspended sediments) over the terrain surface is critical in every erosion algorithm. We rely on the drainage area a that represents the amount of water flowing through each point on the terrain.

We calculate an iterative drainage approximation a_i using a parallel implementation [Schott et al. 2023], updating a_i using a flow routing algorithm parametrized by an exponent p that interpolates flow between diffusive and steepest slope behavior. Let $h_i(\mathbf{p})$ denote the elevation of the terrain at iteration i at a point \mathbf{p} . Let $V(\mathbf{p})$ denote the set of points in the neighborhood of \mathbf{p} (see Figure 5), *i.e.* the eight direct neighbors on a regular grid, $V_i^-(\mathbf{p})$ the subset of lower neighbors and its complementary subset $V_i^+(\mathbf{p})$:

$$V_i^-(\mathbf{p}) = \{\mathbf{q} \in V(\mathbf{p}) \mid h_i(\mathbf{q}) < h_i(\mathbf{p})\}$$

Flow routing determines how a cell distributes its water to its downhill neighbors by computing the coefficients $w_i(\mathbf{p}, \mathbf{q})$ between the center of the cell at point \mathbf{p} and the center of the neighboring cells \mathbf{q} . Let $s_i(\mathbf{p}, \mathbf{q}) = (h_i(\mathbf{p}) - h_i(\mathbf{q})) / \|\mathbf{p} - \mathbf{q}\|$ denote the slope between two cells, we define:

$$w_i(\mathbf{p}, \mathbf{q}) = \frac{s_i^p(\mathbf{p}, \mathbf{q})}{\sum_{\mathbf{p}' \in V_i^-(\mathbf{p})} s_i^p(\mathbf{p}, \mathbf{p}')} \quad (1)$$

The flow routing exponent p , ranging in $[1, \infty[$, defines the flow behavior (see Figure 6): the flow spreads proportionally to the slope if $p = 1$ and down to the steepest cell if $p \rightarrow \infty$. Holmgren [1994] discusses the influence of the parameter p , in our experiments we used $p = 1.3$ to avoid sharp fluvial incision produced by high exponents.

At every iteration, we update the drainage area as the sum of the precipitations (simplified to 1 here, but proportional to the cell area in practice) and the contribution of every uphill neighbor, which




Fig. 5. Flow routing and neighborhoods.

leads to the following equation:

$$a_{i+1}(\mathbf{p}) = 1 + \sum_{\mathbf{q} \in V_i^+(\mathbf{p})} w_i(\mathbf{q}, \mathbf{p}) a_i(\mathbf{q}) \quad (2)$$

This iterative approach lends itself to parallel implementation, allowing us to interleave bedrock detachment or sediment transport processes within every iteration. Since the modifications to the terrain are small between two erosion iterations, the previous drainage approximation a_i remains accurate enough to be reused in the next iteration $i + 1$. Therefore, only one iteration of flow routing is required to maintain satisfying drainage values throughout the erosion algorithms.




Fig. 6. Influence of the flow routing exponent over the drainage area.

4.2 Fluvial erosion

The stream power equation [Whipple and Tucker 1999], which models the interaction between fluvial erosion and tectonic uplift, states that the rate of change of surface elevation results from the balance between the surface uplift $u(\mathbf{p})$ and the fluvial erosion defined by the stream power, denoted as $e(\mathbf{p}) = s^n(\mathbf{p}) a^m(\mathbf{p})$, a function of the slope $s(\mathbf{p})$ and the drainage area $a(\mathbf{p})$:

$$\frac{\partial h}{\partial t}(\mathbf{p}) = u(\mathbf{p}) - k s^n(\mathbf{p}) a^m(\mathbf{p}) \quad (3)$$

Iteratively solving this equation until convergence leads to the formation of large mountain ranges, with only indirect control over the final height. The constants n and m typically fulfill the condition $m/n = 0.5$. Table 4 in Appendix summarizes the constant terms and the coefficients used in our experiments.

In a terrain amplification context, we aim at augmenting an existing low-resolution terrain with fluvial erosion features and patterns while preserving its large-scale structure. To achieve this objective, we ignore the uplift component responsible for the emergence of new mountain ranges. Subsequently, through a few iterations involving only the stream power term, we carve dendritic fluvial paths while maintaining the overall shape of the terrain.




Fig. 7. Visualization of the slope s , the drainage area a and the stream power $e = s^n a^m$ on an eroded terrain.

Our method derives from the parallel implementation proposed in the work of Schott *et al.* [2023]. Let $h_i(\mathbf{p})$ denote the elevation

at point \mathbf{p} and iteration i , k the erosion coefficient, $s_i(\mathbf{p})$ the local slope and $a_i(\mathbf{p})$ the drainage area (see Section 4.1 and Figure 7 for a visual inspection of those quantities). We propose the following iterative erosion algorithm:

$$h_{i+1}(\mathbf{p}) = h_i(\mathbf{p}) - k e_i(\mathbf{p}) \quad e_i(\mathbf{p}) = s_i^n(\mathbf{p}) a_i^m(\mathbf{p}) \quad (4)$$

A commonly observed feature of stream power erosion is the formation of sharp incisions in the terrain, particularly over the steep slopes in mountain ranges. To solve this issue, we introduce a modified stream power definition \tilde{e} with slope and drainage area bounds s_{\max} and a_{\max} clamping e to a controlled maximum value:

$$\begin{aligned} h_{i+1}(\mathbf{p}) &= h_i(\mathbf{p}) - k \tilde{e}_i(\mathbf{p}) \\ \tilde{e}_i(\mathbf{p}) &= \min(s_i^n(\mathbf{p}), s_{\max}^n) \min(a_i^m(\mathbf{p}), a_{\max}^m) \end{aligned} \quad (5)$$




Fig. 8. Comparison between standard (left) and modified stream power (right): bounding produces more realistic and uniform incisions.

Bounds guarantee that N iterations of the modified stream power remove less bedrock than $N k s_{\max}^n a_{\max}^m$, which leads to more uniformly distributed erosion features. Figure 8 illustrates the impact of the bounding terms over the stream power: without clamping, plunging erosion features appear on steep slopes, whereas the peak regions lack erosion landmarks. In contrast, bounding yields a more uniform distribution of components with softer incisions.




Fig. 9. Influence of the hardness function ρ : a constant coefficient k produces self-similar dendritic patterns (left), whereas a high frequency noise with a user-prescribed hard region (right) generates varied landforms.

Another frequently observed limitation of standard stream power erosion is the exaggerated self-similarity of patterns across scales. Therefore, we introduce spatial variations through a varying coefficient $k(\mathbf{p})$ (instead of a constant). We define $k(\mathbf{p})$ by using a user-controlled hardness function, denoted as $\rho : \Omega \rightarrow [0, 1]$, with $\rho(\mathbf{p}) = 0$ for a high erodibility, and $\rho(\mathbf{p}) = 1$ for resistant bedrock:

$$k(\mathbf{p}) = k(1 - \rho(\mathbf{p}))$$

High values of ρ lead to more deeply eroded areas and can be used to create valleys, whereas low values prevent erosion. Discontinuities in the function ρ tend to generate cliffs. A practical method to introduce variety and reduce self-similar erosion landmarks across

scales is to define ρ as a fractal noise (Figure 9). Introducing randomness in the hardness function also reduces the axis-aligned artifacts produced by the regular grid discretization [Schott et al. 2023].

4.3 Thermal stabilization

Stream power erosion usually yields steep incisions in the relief, which can be observed even with our modified version. Thermal erosion introduced by Musgrave *et al.* [1989] and later reused in the context of layered terrain simulation [Cordonnier et al. 2017] typically balances this phenomenon: parts of the terrain with granular material layers at steep slopes slide until reaching a stable talus angle. Formally, let γ_0 denote the talus angle, also referred to as the critical slope $s_0 = \tan(\gamma_0)$; let k_γ the maximum amount of matter moved by the erosion; the evolution of the surface h of the terrain follows the equation:

$$\frac{\partial h}{\partial t}(\mathbf{p}) = -k_\gamma \max(s(\mathbf{p}) - s_0, 0)$$

We propose an iterative algorithm variant that can be implemented in parallel on graphics hardware. Every cell of the terrain changes its elevation depending on how many neighbors are outside the prescribed slope threshold. Let $\alpha_i(\mathbf{p}) = \#\{\mathbf{q} \in V(\mathbf{p}) \mid s_i(\mathbf{p}, \mathbf{q}) < -s_0\}$ (conversely $\beta_i(\mathbf{p})$ denote the number of neighboring cells below \mathbf{p} (above) and with a steeper slope than the critical slope s_0 , we have:

$$h_{i+1}(\mathbf{p}) = h_i(\mathbf{p}) + k_\gamma(\alpha_i(\mathbf{p}) - \beta_i(\mathbf{p}))$$




Fig. 10. Using a space varying critical slope threshold $s_0(\mathbf{p})$ produces irregular slopes with a less artificial aspect.

The balance between the number of neighbors above and below the slope threshold α and β determines if the elevation of a cell should be lowered or raised. Using a constant critical slope s_0 yields uniform slopes over the entire terrain (Figure 10). To overcome this limitation and similarly to the stream power erosion, we define the critical slope as a spatially-varying function $s_0 : \Omega \rightarrow \mathbb{R}_+$ as suggested by Cordonnier *et al.* [2016]. Perturbing the critical slope with high-frequency fractal noise effectively reduces uniform slope artifacts. Note that using a spatially varying critical slope does not guarantee that the prescribed value is respected even if a convergence is observed. From a user perspective, the variations of angles are more valuable than a strict angle guarantee.

4.4 Sediment deposition

The sediment deposition aims at distributing the sediments produced by fluvial erosion and thermal stabilization. Our algorithm takes its inspiration from the capacity-limited models used in computer graphics [Musgrave et al. 1989] and geomorphology [Braun and Sambridge 1997] to transport sediments.

Those models consistently handle bedrock detachment, sediment transport, and deposition in a coherent framework. Instead, we dismiss bedrock detachment (which removes some amount of bedrock transformed into sediments) already performed by fluvial erosion (Section 4.2) and propose sediment creation, transport, and deposition steps. We iteratively perform sediment deposition as follows. Let g_i denote the thickness of the suspended sediment at step i . We first initialize $g_0 = 0$. We then iteratively perform sediment simulation by computing the amount of created sediments c_i , the transport of sediment t_i and the deposition of sediments d_i . The amount of suspended sediment results from sediment creation, transport and deposition:

$$g_{i+1}(\mathbf{p}) = c_i(\mathbf{p}) + t_i(\mathbf{p}) - d_i(\mathbf{p})$$

Simultaneously, the terrain height is updated as follows:

$$h_{i+1}(\mathbf{p}) = h_i(\mathbf{p}) + d_i(\mathbf{p})$$

Sediment creation is proportional to the fluvial erosion intensity: we add a quantity proportional to the stream power (see Table 4 for deposition constants and values):

$$c_i(\mathbf{p}) = k_c e_i(\mathbf{p})$$

We transport sediment by consistently using the same flow routing coefficients $w_i(\mathbf{p}, \mathbf{q})$ (see Equation 1) as the weights used for computing the drainage area a (Section 4.1):

$$t_i(\mathbf{p}) = \sum_{\mathbf{q} \in \mathcal{V}_i^+(\mathbf{p})} w_i(\mathbf{p}, \mathbf{q}) g_i(\mathbf{q})$$

Deposition results from the competition between two quantities: the amount of available suspended sediment t_i and the intensity of stream power e_i . Therefore, we compute a deposition index $\phi_i(\mathbf{p}) = t_i(\mathbf{p}) - e_i(\mathbf{p})$ from which we derive the amount of sediment d_i :

$$d_i(\mathbf{p}) = \min(t_i(\mathbf{p}), k_d \phi_i(\mathbf{p})) \text{ if } \phi_i(\mathbf{p}) > 0, \text{ and } 0 \text{ otherwise}$$

The parameters k_c and k_d denote the creation and deposition coef-




Fig. 11. Impact of different deposition iteration counts: sediments are predominantly deposited in valleys, preserving the critical angle of thermal stabilization in steep and elevated terrain areas.

ficients respectively (Table 4). Figure 11 shows the same terrain progressively covered with an increasing amount of sediments, which is directly controlled by the number of iterations of the algorithm.

5 CONTROL

Our multi-scale erosion method produces a high-resolution amplified that may still exhibit hydrological inconsistencies. Moreover, erosion often inevitably decreases the elevation of feature regions such as ridges or peaks that the user might want to preserve. Therefore, we present two algorithms for correcting the elevation of the

amplified terrain. Height retargeting enforces specific elevation constraints through a diffusion-based retargeting step (Section 5.1). Multi-scale breaching guarantees the hydrological exactness of the final terrain (Section 5.2) with a process that eliminates depressions and pits without introducing artifacts such as cell-size canyons.

5.1 Height retargeting

We propose a method for restoring the elevation of points of interest $C(T_0)$ in the amplified terrain $T_m = \mathcal{A}(T_0)$. Without loss of generality, we describe the case where we automatically detect ridge lines and peaks, but our technique also generalizes to other components such as valleys or cliff faces.




Fig. 12. Thresholding operations on the drainage area a of the terrain, allowing the detection of ridge lines with a varying level-of-detail.

Detecting peaks and ridges on a terrain is a typical difficulty in geomorphology and terrain modeling. Our approach builds on the computation of the drainage area a and exploits that points with low values correspond to crests and peaks. Thus, we define the set of constrained points $C = \{c_k\}$ that have a drainage value lower than a user-defined drainage threshold $a(c_k) < a_0$. This operation comes at no extra cost as a was computed during the amplification steps (see Section 4.2) and allows the detection of the ridge network with a varying level of detail according to the chosen threshold a_0 (see Figure 12).

Let h_0 denote the elevation of the input terrain T_0 , and $h_{\mathcal{A}}$ the elevation of the multi-scale amplified terrain $\mathcal{A}(T_0)$. For every control point $c_k \in C$, we define the elevation error:

$$\delta(c_k) = h_0(c_k) - h_{\mathcal{A}}(c_k)$$

Restoring the elevation of every point in C by adding the error $\delta(c_k)$ yields discontinuities. Argudo *et al.* [2019] partially addressed the case of distant peak and saddle control points by elevating the heightfield in a compactly defined region of influence around control points according to a decreasing falloff function of the distance.

One crucial limitation of this method is that it cannot handle crest lines or close control points with overlapping regions of influence or different elevation errors $\delta(c_k)$. Here we consider the general problem of interpolating irregular data: we aim at interpolating the error everywhere between all points $\delta(c_k)$ to build a continuous error map $E : \Omega^2 \rightarrow \mathbb{R}$ that we can add to $h_{\mathcal{A}}$ (see Figure 13). Although this is




Fig. 13. Retargeting process.

a problem with a wide range of solutions, we found that a simple diffusion process delivered satisfying results. Let E_0 the initialization of the error map:

$$\forall c_k \in C, E_0(c_k) = \delta(c_k) \quad \forall p \notin C, E_0(p) = 0$$

We build E using a diffusion-based error propagation approach while preserving the error at the constraint points. Let Δ denote the Laplacian operator, handling the diffusion, and $m : \Omega^2 \rightarrow \{0, 1\}$ the indicator function of C , *i.e.*, $m(p) = 1$ only for constrained points. At every diffusion step, the map E_{i+1} is iteratively defined as:

$$E_{i+1}(p) = m(p) E_0(p) + (1 - m(p)) (E_i - \Delta E_i)(p)$$




Fig. 14. Use of the retargeting algorithm as post-processing on an eroded terrain. The peaks and ridges elevations are restored, without damaging the features of the terrain.

In our experiments, we used $n = 500$ iterations without letting the diffusion converge onto the full domain Ω . The diffusion can also be replaced by any diffusive operation (gaussian blur, mean filtering, ...). The smooth diffused error is finally applied to obtain the final retargeted elevation of $\mathcal{R} \circ \mathcal{A}(T_0)$:

$$h_{\mathcal{R}}(p) = h_{\mathcal{A}}(p) + E(p)$$

Using a constrained diffusion process allows us to specify elevations for peaks, saddle points, crest lines, or even entire regions, automatically detected or prescribed by the user (see Figure 14), while maintaining coherent values outside those points. The computed error map is smooth enough to allow the retargeting of the constrained points without damaging erosion features. Moreover, each erosion process only affects the elevation of peaks and crests by a small amount. Therefore, applying one retargeting step at the end of the amplification is sufficient to fulfill every constraint.

5.2 Hydrological coherency

Hydrological coherency is a fundamental aspect of geological realism. Except for the rare presence of endorheic lakes [Cordonnier *et al.* 2016], water should flow towards the boundaries of the terrain and reach the seashores.

A straightforward strategy to obtaining consistent drainage is to apply a breaching algorithm [Barnes *et al.* 2014; Garbrecht and Martz 1997]. Unfortunately, the one-cell-wide carved channels produce unrealistic sharp ravines or canyons crossing the previously generated landforms (see Figure 16, left image). Although multi-scale erosion amplification improves the drainage area consistency (see Figure 17, center image), the global drainage network still contains local depressions and pits.

We modify the original breaching algorithm to avoid one-cell-wide channels with a multi-scale approach. Let \mathcal{B} denote the breaching operator from Barnes *et al.* [2014]. We define the multi-scale




Fig. 15. Results of the amplification process, applied on a real low resolution terrain (original resolution 128×128).




Fig. 16. Comparison between the original breaching algorithm (with one-cell-wide channel highlighted in red) and multi-scale breaching on a synthetic fractal terrain (with carved regions outlined in green).




Fig. 17. Influence of our erosion method and multi-scale breaching on the water flow; regions with inconsistent drainage are highlighted in red.

breaching $\tilde{\mathcal{B}}$ by combining a series of *partial* breaching operations $\mathcal{P}_k, k \in [0, n - 1]$ at different scales with decreasing radius of influence r_k . In our implementation, we use the geometric law $r_{k+1} = 0.5 r_k$.

$$\tilde{\mathcal{B}}(T) = \mathcal{P}_{n-1} \circ \dots \circ \mathcal{P}_0(T)$$

At every scale k , the *partial* breaching operator \mathcal{P}_k modifies the elevation by computing the amount of bedrock that should be removed $T_k - \mathcal{B}(T_k)$ and smoothing it using a compactly supported kernel g_k or radius r_k :

$$T_{k+1} = \mathcal{P}_k(T_k) = T_k - (T_k - \mathcal{B}(T_k)) * g_k$$

We call \mathcal{P}_k a *partial* breaching operator because the convolution reduces the impact of the effective full breach operator \mathcal{B} and diffuses the error within the range r_k . This process produces a succession of carvings with decreasing radii. The final *partial* breaching step \mathcal{P}_n is performed with $r_n = 1$, *i.e.*, $\mathcal{P}_n = \mathcal{B}$, which yields a final hydrologically consistent terrain. This multi-scale breaching produces carvings without artifacts (see Figure 16, right image).

6 RESULTS

We implemented our method in C++ and the presented algorithms were coded as GLSL compute-shaders. Stream power erosion \mathcal{E} ,

thermal erosion \mathcal{T} and sediment deposition process \mathcal{D} lend themselves for parallel implementation on graphics hardware, which in turn allows for integration into a real-time authoring framework.

Experiments were performed on a desktop computer equipped with Intel® Core i7, clocked at 4 GHz with 16 GB of RAM, and an NVidia GTX 3080 graphics card. The terrains were rendered using Mitsuba and streamed into Eon-Software Vue to produce photo-realistic landscapes (Figure 25). The multi-scale amplification code is available at <https://github.com/H-Schott/MultiScaleErosion>.

6.1 Amplification

Landscape designers traditionally draft an initial low-resolution model by organizing a variety of landforms, such as principal mountain ranges, valleys, canyons, and coastlines. They combine procedural primitives or patches extracted from digital elevation models and then apply erosion algorithms to add realistic and detailed terrain features. Figure 18 shows that our method perfectly adapts to this workflow to enhance low-resolution terrains produced with a variety of techniques, using procedural noise [Ebert et al. 2002], orometry [Argudo et al. 2019], example-based models [Zhou et al. 2007], and large-scale tectonic-simulation [Schott et al. 2023].

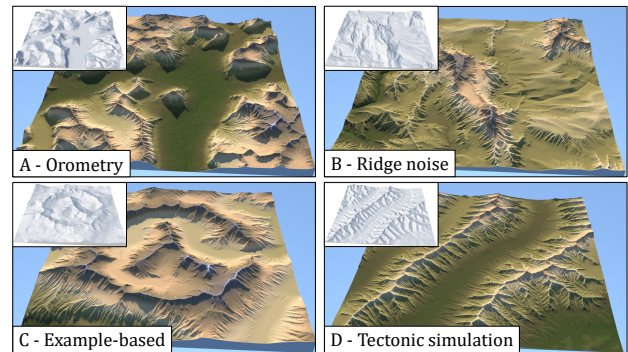


Fig. 18. Results of our amplification on various types of low resolution terrains: procedural, example-based, orometry and tectonic-simulation-based.

Readily available low-resolution digital elevation maps often serve as a basis for authoring more complex terrains but lack details and need to be amplified. Our multi-scale approach can also be used in this context for generating hydrologically consistent details while preserving the input terrain structure. Figure 19 shows the example of a simple low-resolution fault progressively amplified into a realistic canyon. Figure 15 shows a typical amplification process where




Fig. 19. Amplifying a procedurally generated incision into an eroded canyon with a $\times 16$ amplification factor; without relief re-targeting or breaching.

the user progressively upsampled and eroded the terrain at different scales, leading to a total amplification factor of 32 (going from $r_0 = 128$ to $r_n = 4096$). Figure 23 illustrates the difference between our multi-scale amplification and a standard erosion simulation producing one-cell-wide drainage patterns.

6.2 Control

Our multi-scale amplification comes with control parameters for authoring details. The intensity of fluvial erosion can be controlled through the use of the hardness map ρ , which not only removes regular grid discretization artifacts (see Figure 9) but also allows the user to increase or reduce the erosion intensity in specific regions. Figure 20 illustrates this case.




Fig. 20. Impact of the hardness control map ρ on a $\times 16$ amplification.

Separating the different erosion processes (fluvial and thermal erosion, and sediment deposition) also provides additional control over the different amplification steps. Our method allows the user to incrementally amplify the terrain by checking the output of every step and adjusting parameters, such as the amount of sediments and the intensity of fluvial erosion, without changing the previously created details.

Working with a multi-scale process also provides better control over the erosion landmarks scale. In practice, the user can specify the amplification factor of every amplification step \mathcal{A}_k . For instance, terrain B in Figure 21 was amplified using two amplification steps with factors $\times 4$ (see Table 1), whereas other terrains (A, C and D) used four $\times 2$ amplification steps. The user can also set the erosion and deposition intensities.

Finally, the re-targeting post-process offers a way to precisely set a chosen elevation on any selected area without damaging the erosion features resulting from the amplification.

6.3 Comparison with other techniques

Several methods have been proposed for amplifying terrains with details, a process sometimes referred to as terrain augmentation

Terrains		Erosion iterations					
		128	256	512	1024	2048	4096
Figure 15		2600	1800	4200	1600	700	400
Figure 18	A	None	6000	2200	1500	700	400
	B		4200	2800	1500	600	300
	C		3500	2200	1500	600	400
	D		4200	2800	1500	600	300
Figure 21	A	None	5300	1900	1400	1700	800
	B		0	-	4800	-	600
	C		1200	1200	700	500	400
	D		7400	2200	2700	2000	750

Table 1. Number of erosion iterations along with the resolution at which the amplification was performed.

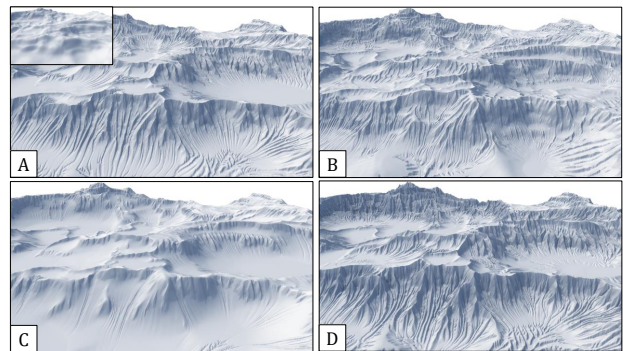


Fig. 21. Comparison of $\times 16$ amplifications with different parameter sets, applied on the same low-resolution terrain.

in the literature. Here, we compare our approach to some of the most recent methods: learning-based [Perche et al. 2023], gradient-aligned procedural noise [Grenier et al. 2024], and example-based sparse modeling [Guérin et al. 2016]. Figure 22 shows three different terrains with various landforms and topography.

Sparse amplification relies on a dictionary of terrain patches constructed from digital elevation models and produces realistic local features. However, contrary to our method, the output may not satisfy hydrological constraints (see Section 5.2), and assembled patches do not produce coherent erosion patterns. The rationale for this is that the patch selection and blending process does not guarantee spatial coherency. Augmenting the patches with orientation and drainage information improves the matching process and yields a more coherent terrain [Argudo et al. 2017], at the cost of a larger patch database. Terrains A and C show ridge lines damaged by the sparse-amplification, in particular some sharp ridge lines were replaced by gentle-slope patches in the sparse dictionary.

The procedural technique proposed in the work of Grenier et al. [2024] generates ravines and elevation details that resemble erosion patterns through carefully designed Phasor noise kernels. Compared to standard isotropic noise-based amplification approaches, this technique synthesizes details aligned with the slope of the terrain. Still, the resulting features are not as realistic as our multi-scale

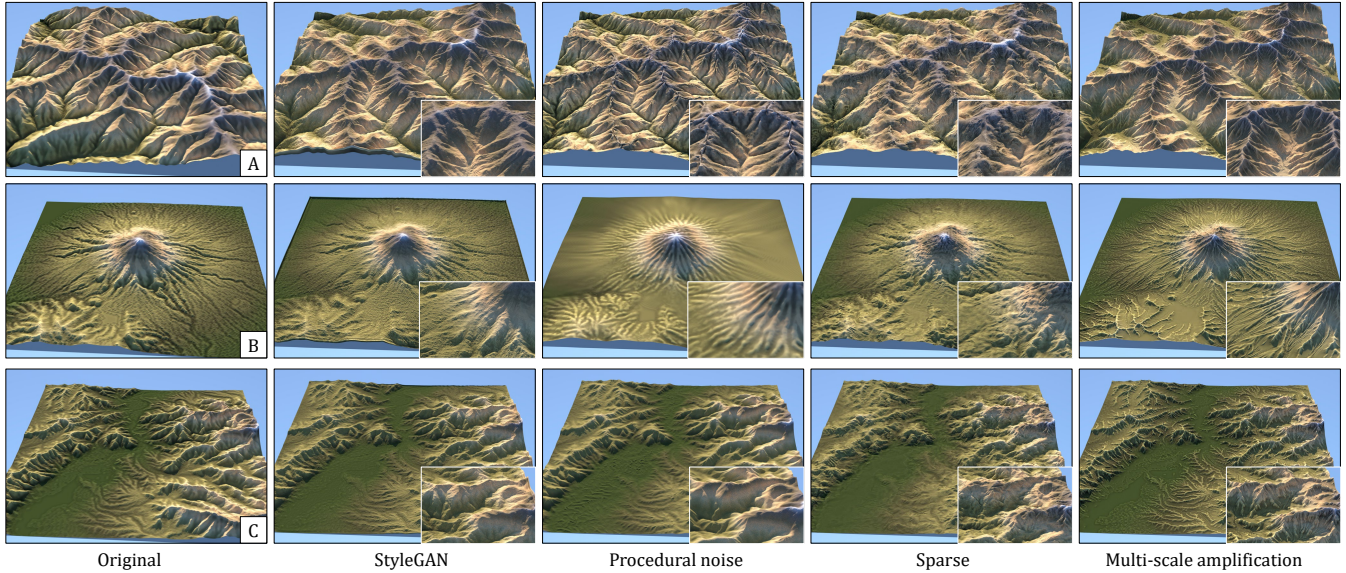


Fig. 22. Comparison of our multi-scale amplification with recent amplification methods: learning-based, procedural, and sparse modeling. The amplified terrains have 2048×2048 resolution with a $\times 8$ amplification factor, obtained from 256×256 input terrains.

erosion-based method: in particular, the lack of a sediment deposition process leads to the absence of valleys with gentle slopes. Terrain A shows noticeable artifacts over ridge lines and characteristic noise patterns located over the smooth slopes of terrain B.




Fig. 23. Comparison between stream power erosion performed at 4096×4096 resolution, and multi-scale amplification with a $\times 16$ factor, from 256×256 to 4096×4096 resolution.

One of the best terrain amplifications that adapt to various inputs and coherently increase resolution is the StyleGAN-based deep learning approach proposed by Perche *et al.* [2023]. Our multi-scale amplification also compares favorably to this approach and produces a wider range of landform details. In particular, fluvial erosion creates realistic ravines and the sediment deposition step produces hydrologically coherent reliefs with sedimentary valleys at the bottom of mountain ranges. This type of amplification only enhances existing landforms but does not augment the relief with detailed erosion landmarks.

6.4 Performance

We implemented the erosion operators \mathcal{E} , \mathcal{T} and \mathcal{D} on graphics hardware. Table 2 reports the performance statistics for one iteration at 256 to 4096 resolution. Timings show that amplification

steps become computationally intensive above 2048 resolution. Still, the average time for computing one erosion step remains below 11 milliseconds. Note that the number of erosion iterations also decreases as resolution increases (see Table 1), which reduces the impact of high-resolution amplification steps. With a focus on user control, our method delivers an interactive amplification framework that allows the designer to modify parameters with direct visual feedback (see accompanying video).

Process	128	256	512	1024	2048	4096	8192
Erosion	0.025	0.027	0.093	0.33	1.27	5.07	19.2
Thermal	0.006	0.007	0.018	0.06	0.23	0.90	3.51
Deposition	0.030	0.055	0.188	0.69	2.70	10.7	41.5

Table 2. Timings (in milliseconds) for the fluvial and thermal erosion and sediment deposition processes at 7 different resolutions.

In contrast, the elevation retargeting and multi-scale breaching processes, implemented on the CPU, take a few seconds to complete at 4096×4096 resolution. This delay is not a limitation in practice as those post-processing steps are performed only once at the end of the multi-scale amplification. A complete parallel implementation on the GPU would significantly improve the overall performance and user experience but remains beyond the scope of this paper.

6.5 Validation

Hydrological consistency is a required criterion for determining geological soundness and perceptual realism of a synthetic terrain. One form of evaluation is to compute and visualize the drainage of a terrain.

In Table 3, we choose to provide a numerical comparison of the drainage of our multi-scale erosion process, without multi-scale

Terrain	Learning	Procedural	Sparse	Ours
A	23	41	30	0.05
B	120	1.7	1.0	0.02
C	62	41	79	4.8

Table 3. Average breaching volume \bar{b} ($\times 10^3$ for readability) for the terrains presented in Figure 22.

breaching, against reference amplification alternatives (see Figure 22 for a visual comparison). We compute the average breaching volume \bar{b} for each case, defined as the volume of material removed from h by breaching required to ensure free drainage of the terrain [Barnes et al. 2014] divided by terrain area Ω .

$$\bar{b} = \frac{1}{\Omega} \int_{\Omega} |h(\mathbf{p}) - h_{\mathcal{B}}(\mathbf{p})| d\mathbf{p}$$

Results demonstrate that our method outperforms other techniques, even before multi-scale breaching that ultimately guarantees a consistent hydrology, thus $\bar{b} = 0$.

6.6 Limitations, future work and extensions

Multi-scale amplification, which relies on different categories of erosion, requires the processing of the entire elevation map at every step, notably to evaluate the drainage (Section 4.2) and the diffusion of sediments (Section 4.4). We limited our experiments to a maximum grid size of 8192×8192 , on a single up-to-date GPU. Due to the memory limits of graphic hardware, amplifying terrains beyond this size would require decomposing maps into patches, as advocated by Vanek *et al.* [2011], before assembling them. This decomposition introduces a series of technical challenges, particularly the propagation of drainage across the boundaries of the patches, and their seamless blending after erosion.

The features produced by erosion algorithms and control tools determine their style and range. As a result, the results presented in this work demonstrate typical hydraulic erosion landforms and patterns. Simulating a broader range of geological phenomena would diversify the landmarks but remains an open challenge.




Fig. 24. Synthesized terrain generated from an 128×128 input map featuring blended Gaussian bumps and prescribed peaks elevations; the final resolution is 4096×4096 with a $\times 32$ amplification factor.

The multi-scale amplification framework can be adapted to terrain synthesis by increasing the intensity of the erosion effects applied to the input topography. This extension comes from the stream-power-based erosion process already used in interactive terrain modeling [Schott et al. 2023]. In this context, the retargeting (Section 5.2) and multi-scale breaching (Section 5.1) steps play a crucial part

in following the landscape designer’s intent and guaranteeing the hydrological consistency of the output terrain. Figure 24 shows mountain ranges produced by our approach. Starting from a few noised Gaussian bumps defining the initial elevation, the algorithm automatically carves fluvial valleys, performs downstream sediment deposition in flats, an models ravines and gullies at different scales.

7 CONCLUSION

We introduced a novel terrain amplification method based on a multi-scale erosion framework. Starting from a low-resolution terrain, we iteratively increase its resolution and augment it with landmarks and complex patterns across a range of scales by applying physically-based erosion models. Our method controls the amplitude and resolution of the generated features by separating the different erosion processes and using hardness maps. The diffusion-based retargeting process compensates for bedrock detachment and allows the landscape designer to prescribe the elevation of crucial terrain components such as peaks or saddles. Finally, the multi-scale breaching algorithm guarantees a hydrologically consistent output without one-cell canyon artifacts.

Our method runs at interactive rates and amplifies terrain with factors up to $\times 32$ with a final 8192 resolution and compares favorably to procedural and example-based amplification techniques, including deep learning approaches, concerning the consistency of generated features.

Future research directions include extending our technique to large maps (16k and above). Integrating other erosion phenomena into the multi-scale framework world would be worth investigating, particularly aeolian phenomena generating various landscape features such as dunes, yardangs, and generally ventifacts.

ACKNOWLEDGMENTS

This work was supported by the projects AMPLI ANR-20-CE23-0001 and EOLE ANR-23-CE56-0008 of Agence Nationale de la Recherche, France.

APPENDIX

Here we summarize the simulations parameters with their corresponding notations, nominal values, and interpretation.

Symbol	Value	Name
p	1.3	Flow routing exponent
n	2	Stream Power slope exponent
m	0.8	Stream Power drainage exponent
k	5.10^{-4}	Stream Power erosion coefficient
s_{max}	1	Stream Power maximum slope
a_{max}	250	Stream Power maximum drainage area
k_{γ}	5.10^{-5}	Thermal erosion coefficient
γ	[0.8, 1.4]	Thermal noisy angle of repose
k_d	0.1	Deposition strength
k_c	0.1	Deposition creation coefficient
a_0	2	Ridge detection threshold

Table 4. List of constants and simulation parameters.




Fig. 25. Photo-realistic landscapes obtained by the multi-scale erosion.

REFERENCES

- Oscar Argudo, Carlos Andujar, Antonio Chica, Eric Guérin, Julie Digne, Adrien Peytavie, and Eric Galin. 2017. Coherent multi-layer landscape synthesis. *The Visual Computer* 33, 6 (2017), 1005–1015.
- Oscar Argudo, Eric Galin, Adrien Peytavie, Axel Paris, James Gain, and Eric Guérin. 2019. Orometry-Based Terrain Analysis and Synthesis. *ACM Transactions on Graphics* 38, 6, Article 199 (2019), 12 pages.
- Richard Barnes, Clarence Lehman, and David Mulla. 2014. An efficient assignment of drainage direction over flat surfaces in raster digital elevation models. *Computers and Geosciences* 62 (2014), 128–135.
- Bedrich Benes. 2007. Real-Time Erosion Using Shallow Water Simulation. In *Proceedings of Virtual Reality Interactions and Physical Simulations*. Eurographics Association, Dublin, Ireland, 43–50.
- Jean Braun and Malcolm Sambridge. 1997. Modelling landscape evolution on geological time scales: a new method based on irregular spatial discretization. *Basin Research* 9, 1 (1997), 27–52.
- Jean Braun and Sean D. Willett. 2013. A very efficient $O(n)$, implicit and parallel method to solve the stream power equation governing fluvial incision and landscape evolution. *Geomorphology* 180–181 (2013), 170–179.
- Guillaume Cordonnier, Jean Braun, Marie-Paule Cani, Bedrich Benes, Eric Galin, Adrien Peytavie, and Eric Guérin. 2016. Large Scale Terrain Generation from Tectonic Uplift and Fluvial Erosion. *Computer Graphics Forum* 35, 2 (2016), 165–175.
- Guillaume Cordonnier, Marie-Paule Cani, Bedrich Benes, Jean Braun, and Eric Galin. 2018. Sculpting Mountains: Interactive Terrain Modeling Based on Subsurface Geology. *IEEE Transactions on Visualization and Computer Graphics* 24, 5 (2018), 1756–1769.
- Guillaume Cordonnier, Eric Galin, James Gain, Bedrich Benes, Eric Guérin, Adrien Peytavie, and Marie-Paule Cani. 2017. Authoring Landscapes by Combining Ecosystem and Terrain Erosion Simulation. *ACM Transactions on Graphics* 36, 4 (2017), 134.
- Guillaume Cordonnier, Guillaume Jouviet, Adrien Peytavie, Jean Braun, Marie-Paule Cani, Bedrich Benes, Eric Galin, Eric Guérin, and James Gain. 2023. Forming Terrains by Glacial Erosion. *ACM Transactions on Graphics* 42, 4, Article 61 (2023), 14 pages.
- Giliam J. P. de Carpentier and Rafael Bidarra. 2009. Interactive GPU-based Procedural Heightfield Brushes. In *Proceedings of the International Conference on Foundations of Digital Games*. ACM, Orlando, USA, 55–62.
- David S. Ebert, F. Kenton Musgrave, Darwyn Peachey, Ken Perlin, and Steven Worley. 2002. *Texturing and Modeling: A Procedural Approach* (3rd ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- Mathieu Gaillard, Bedrich Benes, Eric Guérin, Eric Galin, Damien Rohmer, and Marie-Paule Cani. 2019. Dendry: A Procedural Model for Dendritic Patterns. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*. ACM, New York, NY, USA, 1–9.
- James Gain, Bruce Merry, and Patrick Marais. 2015. Parallel, Realistic and Controllable Terrain Synthesis. *Computer Graphics Forum* 34, 2 (2015), 105–116.
- Eric Galin, Eric Guérin, Adrien Peytavie, Guillaume Cordonnier, Marie-Paule Cani, Bedrich Benes, and James Gain. 2019. A Review of Digital Terrain Modeling. *Computer Graphics Forum (Proceedings of Eurographics 2019)* 38, 2 (2019), 553–577.
- Jürgen Garbrecht and Lawrence W. Martz. 1997. The assignment of drainage direction over flat surfaces in raster digital elevation models. *Journal of Hydrology* 193, 1 (1997), 204–213.
- Charline Grenier, Eric Guérin, Eric Galin, and Basile Sauvage. 2024. Real-time terrain enhancement with controlled procedural patterns. *Computer Graphics Forum* 43, number (2024).
- Eric Guérin, Julie Digne, Eric Galin, and Adrien Peytavie. 2016. Sparse representation of terrains for procedural modeling. *Computer Graphics Forum (Proceedings of Eurographics 2016)* 35, 2 (2016), 177–187.
- Eric Guérin, Julie Digne, Eric Galin, Adrien Peytavie, Christian Wolf, Bedrich Benes, and Benoit Martinez. 2017. Interactive Example-Based Terrain Authoring with Conditional Generative Adversarial Networks. *ACM Transactions on Graphics (Proceedings of Siggraph Asia 2017)* 36, 6, Article 228 (2017), 13 pages.
- Peter Holmgren. 1994. Multiple flow direction algorithms for runoff modelling in grid based elevation models: An empirical evaluation. *Hydrological Processes* 8 (1994), 327–334.
- Peter Krištof, Bedrich Benes, Jaroslav Krivánek, and Ondřej Šřava. 2009. Hydraulic Erosion Using Smoothed Particle Hydrodynamics. *Computer Graphics Forum* 28, 2 (2009), 219–228.
- Joshua Lochner, James Gain, Simon Perche, Adrien Peytavie, Eric Galin, and Eric Guérin. 2023. Interactive Authoring of Terrain using Diffusion Models. *Computer Graphics Forum (Proceedings of Pacific Graphics)* 42, 7 (2023), 13.
- Xing Mei, Philippe Decaudin, and Baogang Hu. 2007. Fast Hydraulic Erosion Simulation and Visualization on GPU. In *Pacific Graphics*. IEEE, 47–56.
- Forest Kenton Musgrave, Craig E. Kolb, and Robert S. Mace. 1989. The synthesis and rendering of eroded fractal terrains. *Computer Graphics* 23, 3 (1989), 41–50.
- Benjamin Neidhold, Markus Wacker, and Olivier Deussen. 2005. Interactive physically based Fluid and Erosion Simulation. In *Proceedings of the Eurographics Workshop on Natural Phenomena*. Eurographics Association, Dublin, Ireland, 25–32.
- Simon Perche, Adrien Peytavie, Bedrich Benes, Eric Galin, and Eric Guérin. 2023. Authoring Terrains with Spatialised Style. *Computer Graphics Forum (Proceedings of Pacific Graphics)* 42, 7 (2023).
- Pascale Roudier, Bernard Peroche, and Michel Perrin. 1993. Landscapes synthesis achieved through erosion and deposition process simulation. *Computer Graphics Forum* 12, 3 (1993), 375–383.
- Hugo Schott, Axel Paris, Lucie Fournier, Eric Guérin, and Eric Galin. 2023. Large-Scale Terrain Authoring through Interactive Erosion Simulation. *ACM Transactions on Graphics* 42, 5, Article 162 (2023), 15 pages.
- Joshua J. Scott and Neil A. Dodgson. 2021. Example-based terrain synthesis with pit removal. *Computers and Graphics* 99, C (2021), 43–53.
- Ondřej Šřava, Bedrich Benes, Matthew Brisbin, and Jaroslav Krivánek. 2008. Interactive Terrain Modeling Using Hydraulic Erosion. In *Proceedings of the Symposium on Computer Animation*. 201–210.
- Flora Ponjou Tasse, Arnaud Emilien, Marie-Paule Cani, Stefanie Hahmann, and Adrien Bernhardt. 2014. First Person Sketch-based Terrain Editing. In *Proceedings of Graphics Interface*. Canadian Information Processing Society, Montreal, Canada, 217–224.
- Juraj Vanek, Bedrich Benes, Adam Herout, and Ondřej Šřava. 2011. Large-Scale Physics-Based Terrain Editing Using Adaptive Tiles on the GPU. *IEEE Computer Graphics and Applications* 31, 6 (2011), 35–44.
- Kelin X. Whipple and Gregory E. Tucker. 1999. Dynamics of the stream-power river incision model: Implications for height limits of mountain ranges, landscape response timescales, and research needs. *Journal of Geophysical Research: Solid Earth* 104, B8 (1999), 17661–17674.
- Xiaoping P. Yuan, Jean Braun, Laure Guerit, Delphine Roubey, and Guillaume Cordonnier. 2019. A new efficient method to solve the stream power law model taking into account sediment deposition. *Journal of Geophysical Research: Earth Surface* 124, 6 (2019), 1346–1365.
- Jian Zhang, Chen Li, Peichi Zhou, Changbo Wang, Gaoqi He, and Hong Qin. 2022. Authoring Multi-style Terrain with Global-to-Local Control. *Graphical Models* 119 (2022), 101–122.
- Yiwei Zhao, Han Liu, Igor Borovikov, Ahmad Beirami, Maziar Sanjabi, and Kazi Zaman. 2019. Multi-Theme Generative Adversarial Terrain Amplification. *ACM Transactions on Graphics* 38, 6, Article 200 (2019), 14 pages.
- Howard Zhou, Jie Sun, Greg Turk, and James M. Rehg. 2007. Terrain Synthesis from Digital Elevation Models. *IEEE Transactions on Visualization and Computer Graphics* 13, 4 (2007), 834–848.