

Volcanic Skies: coupling explosive eruptions with atmospheric simulation to create consistent skiescapes

P. C. Pretorius¹, J. Gain^{1,2}, M. Lastic², G. Cordonnier³, J. Chen², D. Rohmer², and M.-P. Cani²,

¹ University of Cape Town, South Africa

² LIX, École Polytechnique and CNRS, IP Paris, France

³ Inria, Université Côte d'Azur, France



Figure 1: Impressive phenomena arise from the interaction between the plume of explosive volcanoes and the atmosphere. Our framework is able to capture (from left to right): embedded skirt clouds, cap clouds, ash rain, and volcanic lightning.

Abstract

Explosive volcanic eruptions rank among the most terrifying natural phenomena, and are thus frequently depicted in films, games, and other media, usually with a bespoke once-off solution. In this paper, we introduce the first general-purpose model for bi-directional interaction between the atmosphere and a volcano plume. In line with recent interactive volcano models, we approximate the plume dynamics with Lagrangian disks and spheres and the atmosphere with sparse layers of 2D Eulerian grids, enabling us to focus on the transfer of physical quantities such as temperature, ash, moisture, and wind velocity between these sub-models. We subsequently generate volumetric animations by noise-based procedural upsampling keyed to aspects of advection, convection, moisture, and ash content to generate a fully-realized volcanic skyscape. Our model captures most of the visually salient features emerging from volcano-sky interaction, such as windswept plumes, enmeshed cap, bell and skirt clouds, shockwave effects, ash rain, and sheathes of lightning visible in the dark.

CCS Concepts

• **Computer Graphics** → Physical simulation; Procedural animation;

1. Introduction

The portrayal of volcanic eruptions using computer graphics has taken centre stage in a range of digital media, including films and games. From a storytelling perspective, the catastrophic power of volcanoes has been embedded in collective myth from the time of the eruption of Vesuvius in 79 AD and before. There is also a place for animating and rendering volcanic eruptions in both public education and risk forecasting. The challenge in depicting this phenomenon lies in the multifaceted interaction between the eruption

plume and the atmosphere on the one hand, and between lava flows and the terrain on the other. For instance, the force of the explosion alone can deform the landscape (as happened with the caldera breach of Mt St. Helens in 1980) and the flow of lava and lahar (a slurry of water, ice, mud, ash, and rock fragments) can reshape terrain through erosion and accretion.

In this paper, we set aside effusive eruptions to focus on explosive (so called Plinian) eruptions dominated by the ejection of rock fragments, gas, and ash. Here the short to medium term bi-

directional interaction between the eruption plume and the surrounding atmosphere is key and gives rise to effects such as entrainment of the surrounding air and moisture that combines with convective uplift to form transitory clouds enmeshed with the plume (see Figure 2 and 3), lightning caused by the convective friction of ash particulates (see Figure 4), and the transport of fine ash around which moisture accumulates leading to ash rain downwind (see Figure 5).



Figure 2: Convective uplift can create embedded lenticular cap clouds. Sarychev Peak eruption, Russia, in 2009. ©NASA.




Figure 3: Skirt clouds are formed by the uplift and condensation of air layers ahead of and around the plume. Eruption of Hunga Tonga–Hunga Ha’apai, Tonga, in 2022, courtesy of the Tonga Geological Services.

Existing eruption simulations [MDN02, LRC*22] focus on plume dynamics and treat the atmosphere as a static boundary condition. For example, a constant wind velocity at discrete altitudes is assumed. This neglects bi-directional effects where the plume influences the atmosphere through heat, moisture, and ash transfer



Figure 4: A net of lightning encases the volcanic plume, an effect that is more visible during a night-time eruption. Eyjafjallajökull eruption, Iceland, in 2010. ©Terje Sörgjerd.



Figure 5: Ash is transported from the plume, mixes with water, and falls as ash rain. Eruption of Shinmoedake, Japan, in 2011, courtesy of the U.S. Geological Survey.

and the atmosphere acts on the plume in turn through advective and convective forces.

On the other hand, atmospheric models devoted primarily to cloud simulation [VGL*20, HMP*20] are also insufficient because they are not configured to handle the extremes of heat at the ejection site, do not account for ash content in the atmosphere, and generally update too slowly to capture fine plume detail.

To address the specific challenges of plume-atmosphere interaction, we propose a two-component coupled model, based on a hybridisation of Lagrangian plume dynamics [LRC*22] and atmo-

spheric Eulerian advection and procedural convection [VGL*20]. Specifically, we approximate the plume as an evolving stack of cylinders whose motion is determined primarily by ejection and air entrainment forces and secondarily by temperature, moisture, and wind velocity from the atmosphere. These cylinders spawn Lagrangian spheres which provide both a coarse visual hull for the plume and also transfer heat, moisture, and ash to the atmosphere. For computational efficiency, our atmospheric model is structured as a stack of Eulerian grids containing pertinent atmospheric quantities (temperature, water vapour, condensed moisture, ash, and advective and convective velocity). Horizontal wind is simulated through Eulerian advection, while vertical convection due to heat imbalance and uplift from terrain slope is handled pseudo-physically using parcel transport and pressure adjustment. Heat, moisture and ash from the plume are transferred to this layered model and induce convection, cloud formation, and, ultimately, ash fall. Finally, from this coupled simulation outputs are fed to a procedural upsampling process. This provides a finely-sampled 3D density grid for the atmosphere and a collection of ellipsoids of varying density for the plume, suitable for direct volumetric rendering or export to external applications for use in games, film, or education.

Such an upsampling approach enables relatively coarse-scale simulation with attendant efficiency benefits, while generating a plausible fine-scale final animation. However, it does mean sacrificing some of the complex turbulence detail, which is replaced by noise processes.

In developing our framework we build on existing models in computer graphics for plume dynamics [LRC*22] and skyscape simulation [VGL*20]. In doing so, it is necessary to extend them to handle extreme temperatures, encompass a variety of previously unsupported phenomena, and enable effective coupling between Eulerian and Lagrangian strategies.

In summary, our contribution is an efficient hybrid model adapted to high-temperature physics that guides plume dynamics using advective and convective atmospheric velocity; transfers heat, moisture and ash from plume to atmosphere; and supports procedurally amplified volumetric rendering, thereby enabling for the first time a variety of attendant phenomena, including enmeshed clouds, ash rain, and plume lightning.

2. Related work

While significant effort within computer graphics has been devoted to the simulation of clouds [WCGG18, HMP*20, VGL*20, HHP*21], smoke [FSJ01, WXCT19] and lava flow [SAC*99, NKL*07, ZKL*17], there has to date been relatively little focus on volcanic plumes. While plumes share some characteristics with related phenomena, such as the convective velocity of cumulonimbus clouds, they are unique in regard to the extremes of temperature and vertical scale.

In terms of plume simulation, Mizuno *et al.* [MDN02] employ a fully three-dimensional Eulerian grid combined with semi-lagrangian advection. Despite efforts to accelerate this through a coupled map lattice approach, for memory and computation reasons the resolution is limited to $100 \times 100 \times 100$ cells. To over-

come this, Lastic *et al.* [LRC*22] adopt a coarse-grained strategy using Lagrangian cylinders that extend like a spinal column along the core line of the eruption to model the dynamics. The ascending cylinders are refined into a hierarchy of rotating spheres, and procedural noise-based upsampling is used to introduce rendering detail. Although our model is partly inspired by this proxy-geometry approach, none of these models address the two-way interaction between the plume and the atmosphere.

Another source of inspiration is the recent work in hybrid atmospheric simulation. For example, Fire in Paradise [HBP*21] couples a module-based ecosystem with fluid solvers for the spread of fire, smoke, and clouds in the service of capturing wildfire dynamics. This enables complex feedback effects, such as flammagenitus clouds that arise from smoke and convective uplift and can subsequently lead to condensation and rain capable of extinguishing fires at other sources. EcoClimates [PMG*22] captures the feedback cycles between the atmosphere and biosphere that give rise to microclimates and vegetation patterning. This is achieved by using temperature and moisture maps as an exchange mechanism between soil, vegetation, and atmospheric sub-models. As in our case, these simulations are tailored to the requirements of specific phenomena.

Unsurprisingly, simulation-based studies of volcanic eruption are common in the geosciences. Typically, these model the volcanic column as a multiphase flow and discretize the domain using a regular three-dimensional grid [DR06, NEOdMVC22]. These simulations are capable of capturing a wide-range of effects, including volcanic jets, convective plumes, collapsing columns, and ballistic ejecta. In addition, geophysicists have explored related emergent phenomena in detail. This includes the transport of water vapour from ejection and ambient sources to ultimate condensation and cloud formation at higher altitudes [GBW97]; the entrainment of surrounding air into the column [MTT56, SKI0]; the accretion of ash, ice and other particulates within the plume [GL94], and mechanisms for electrical charging that lead to volcanic lightning [JWL*08, SGVE*21].

Unfortunately, transplanting these techniques directly to Computer Graphics is not viable. First, the CFD models involved are accurate but heavyweight, and often require days to execute. Second, due to the size of the domain (around 60×60 km horizontally and 20 km vertically) and simulated timespan (in the range of several hours), a coarse resolution of around $100 \times 100 \times 100$ m per grid cell is often employed. This suffices for meso-scale studies of thermodynamics and chemical composition [TGL*05, SCC*16] but cannot convey the fine-scale turbulence necessary for visual applications. Nevertheless, these studies do provide a rich source for validation of emergent properties, such as the temporal evolution of plume shape, vertical velocity, and ash deposition.

3. Overview

3.1. A primer on explosive volcanic eruption

It is worth briefly explaining the physics of explosive volcanic eruptions, as illustrated in Figure 6. Here, we concentrate on the evolution of the plume and de-emphasise surface and subsurface effects, such as the formation of magma chambers and calderas, and the flow of pyroclastic clouds, lava, and lahar mudslides. For a more

complete treatment of volcanology, the reader is referred to works by Parfitt and Wilson [PW09] and Cas *et al.* [CGW18].




Figure 6: The progression of a Plinian eruption through the jet, convective, and umbrella phases, including the formation of skirt clouds and ashfall from the downwind plume.

Volcanic eruptions that follow the so-called Plinian or ultra-Plinian regime, as contrasted with Peléan, Hawaiian, and other forms, are caused by pressure from sources such as the separation out of solubles or the melting and expansion of rock during the formation of magma. This builds behind a plug sealing the volcanic vent. The eventual eruption is an expulsion of a mixture of gas, lava, and tephra at high velocities (50-200 m/s) and high densities (several hundred kg/m^3).

The column emitted from the vent progresses through three stages. In the initial low-altitude jet phase it is lifted by the primary momentum of the explosion. Next, in the convective phase ash sediments out and the surrounding air is entrained and heated, thereby making the plume more buoyant than the surrounding atmosphere. If the ejection velocity is insufficient to lift the plume to a level where entrainment can occur then a Peléan regime applies and the plume will collapse and spread across the terrain as a pyroclastic flow. Otherwise, the plume reaches the umbrella phase where buoyancy is neutral with respect to the atmosphere, and spreads out in the direction of the prevailing wind.

Depending on the force of the eruption this may be at an altitude as high as 45km, making it one of the few natural phenomena capable of significantly breaching the tropopause - i.e., the atmospheric boundary at about 12-17km high that demarcates the troposphere below, where temperature decreases with altitude, from the stratosphere, where it increases again. In fact, it is this breaching and the injection of sulphur and other particulates into the stratosphere that can lead to long-term consequences, such as a volcanic winter.

The most striking visual consequences of the explosion are:

- **Embedded clouds** (see Figures 2 and 3): Apart from air entrained from the sides, air layers can also be driven upwards ahead of the advancing plume, until eventually the uplifted layers and plume intermingle. Given sufficient lift the water vapour in these layers condenses to form clouds. Depending on the structure of the at-

mosphere and plume these can present as layered "skirt clouds" that overhang the plume [Bar82] or saucer-like cap clouds.

- **Transient condensation clouds:** In common with nuclear explosions, the rarefaction phase of a volcanic shockwave causes a drop in pressure and temperature that can lead to a transient and spherical imprint of condensation [YI07].
- **Ashfall** (see Figure 5): One effect of an eruption is to transport pyroclasts consisting of ash, cinders, and blocks from the vent into the atmosphere. The heavier fragments tend to fall close to the vent, contributing to formation of a volcano cone, but finer ash can be carried far and wide, to later rain down as ashfall.
- **Lightning** (see Figure 4): The circulation within a volcanic plume leads to friction of ash, ice, and tephra particles and often leads to discharges within and around the plume in the form of volcanic lightning. Although fleeting, this phenomenon can be spatially extensive, sometimes taking the appearance of a web or sheath that encloses the plume.

3.2. Framework

Our strategy, summarised in Figure 7, is to combine three sub-models to achieve a balance between computational efficiency and plausible fine-scale detail. The volcanic plume is captured through cylinder-and-sphere proxy geometry driven by Lagrangian dynamics, the atmosphere is represented by discrete layers subject to within-layer semi-Lagrangian advection and between layer procedural transfer, and lastly a noise-based amplification process provides detail prior to volumetric rendering.

Despite their structural differences, these sub-models share the same physical quantities (see Table 1), which makes their combination possible through transportation and state transformation during the simulation. These physical quantities include: temperature T , which peaks at the volcano vent and is spread through the agency of the plume to the atmosphere; water vapour W present in both the atmosphere and erupting column, which with cooling due to vertical transport forms condensed water C visible as clouds; ash content A contained in the plume parcels, transferred to the atmosphere, and precipitating as ashfall, and 3D velocity (or wind), which is separated in the atmospheric model into horizontal advective H and vertical convective U components and transports the other quantities.

These quantities induce bi-directional feedback in the sub-models, with the exception of ash, which is uni-directionally transferred from plume to atmosphere. For example, atmospheric temperature, water vapour, moisture and wind directly affect the evolving shape of the plume, and at the same time the plume in its jet phase acts to block the wind, and throughout heat, water vapour, moisture, and ash transfer induce convection, cloud formation, and ashfall in the atmosphere.

To summarise, our framework builds on the work of Vimont *et al.* [VGL*20] and Lastic *et al.* [LRC*22] to combine a layered atmospheric model reworked to handle high-temperature physics and to include ashfall and transient condensation clouds (Section 4) with a plume dynamics model that, for the first time, represents embedded clouds and lightning (Section 5). The physical interaction between these sub-models is managed using a transfer mechanism

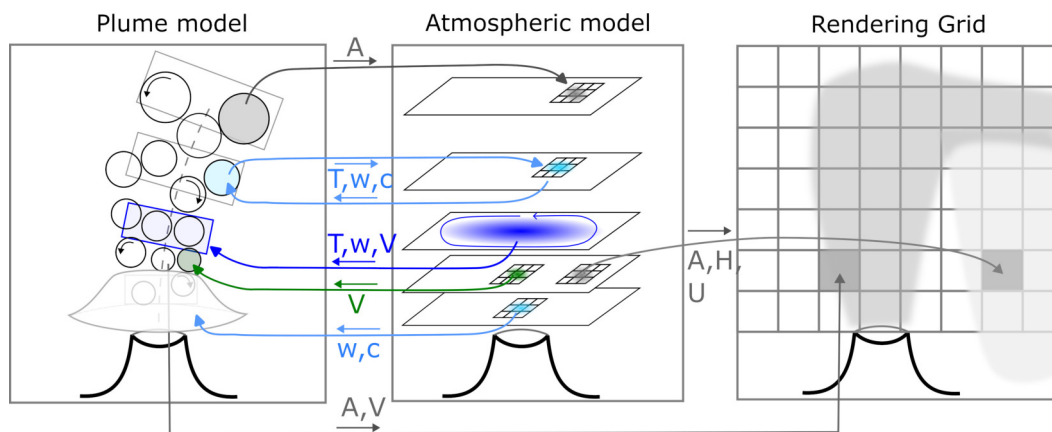


Figure 7: Overview: a layered atmospheric model is coupled to a proxy-geometry plume model with two-way flow of temperature, moisture, ash, and velocity components. Both coupled models feed data to a procedural upsampling process for final detailed off-line animation.

Notation	Quantity	Units
Θ	altitude	m
T	temperature	K
A	ash	kg/m^3
W	water vapour	kg/m^3
C	condensed water	kg/m^3
L	static charge	unit-less
\mathbf{V}	velocity	m/s (3D)
H	horizontal velocity	m/s (2D)
U	vertical velocity	m/s (1D)

Table 1: A summary of quantities, notation and units used in this paper. Subscripts are used to indicate attribution to atmospheric layers (ℓ), grid cells (i), plume spheres (p), or vapour sheet samples (s).

(Section 6) that accounts for air entrainment and ash transfer at inherently different update rates. Lastly, a new procedural amplification mechanism is introduced to provide sub-cell detail suited to volumetric rendering of the plume, clouds and ash-rain, and procedural wire structures for lightning (Section 7).

4. Atmospheric model

We utilise an improved version of the hybrid modelling system created by Vimont *et al.* [VGL*20] to model the weather and generate a consistent skyscape. The atmosphere is divided into several horizontal layered grids at a mesoscale (roughly 60 metres by 60 metres horizontal resolution per 2D grid cell). Each layer is simulated separately in terms of horizontal transport, phase transitions, and radiative temperature changes, followed by a more procedural coupling between layers to model convection and ash and rain precipitation. Using layers limits the memory requirements of the simulation. It also improves efficiency by allowing 2D rather than 3D solves of the Poisson equation. The number of layers and their altitudes are specified to conform to particular meteorological features (see Table 2 for the layer structure used in our scenarios).

Each atmospheric layer consists of a grid of cells i , containing the following parameters (see Table 1): temperature T_i in Kelvin (K) reflecting the mean temperature in the cell; absolute humidity W_i in $kg \cdot m^{-3}$ as the quantity of water vapour present in the cell; water content C_i indicating the quantity of condensed (liquid) water in $kg \cdot m^{-3}$; the particulate content A_i in $kg \cdot m^{-3}$ as the quantity of solid matter, such as volcanic ash in the cell; and the 2D velocity H_i , which is the local planar wind driving horizontal transport (advection).

We alter the phase change system of Vimont *et al.* [VGL*20] to one more suited to the environment in and around a volcanic plume. In addition to the standard evaporation-condensation cycle, we allow for evaporation by boiling of condensed water if the saturation vapour pressure exceeds atmospheric pressure. To support this, we change the saturation vapour pressure equation from the Tetens to Arden Buck formulation [Buc96] because the latter is implemented as curves fitted to empirical data over a wider range of temperatures [RLB19].

We also upgrade the treatment of fluid dynamics inside the layers. We retain a combination of Semi-Lagrangian advection and pressure projection [Sta99], but incorporate a staggered grid, where the velocities are separated, offset, and stored on the edges of the cells, while scalar values (such as temperature, pressure, and moisture) remain within the cell. For example, the x -velocity component is stored on the left edge of cell i , and the y -component on the top edge. The split is required for accurate solves of the pressure update equation, for which we use a Jacobi-preconditioned Conjugate Gradient solver. This solution is very accurate, especially with regard to obstacles such as terrain – a notable shortcoming of the previous approach [VGL*20]. An additional benefit is that it can be used for grids of arbitrary dimension, in contrast to Vimont *et al.*'s method [VGL*20] that requires a square grid with sides of a power of 2.

As in Vimont *et al.* [VGL*20], inter-layer exchange is modelled as a procedural, vertical velocity. This modifies the 2D velocity fields of the horizontal layers to account for convection and they are, therefore, no longer divergence free. The updated advection

scheme is more sensitive to instabilities, which can be overcome by running vertical transport at a lower frequency than horizontal advection. In practice, we found that a ratio of 1 vertical coupling step per 3 advection steps is sufficient to sustain stability indefinitely.

4.1. Precipitation

We add precipitation to the inter-layer exchange mechanism, suitable for both rain and ash particles. The rain is modelled with a physical approximation of coalescence, rather than the probabilistic models common in meteorology [RLW*21, SEH*20]. This is needed to accurately transfer physical quantities between the discrete layers using our procedural approach to inter-layer exchange.

First, potential precipitation ρ_{pot} is calculated using the current density ρ_i of precipitable matter in the cell scaled by a coalescence factor C_{coal} , according to:

$$\rho_{pot} = \rho_i \times C_{coal}, \quad (1)$$

$$C_{coal} = \min(|\vec{H}_i| \times F_{hor} + |U_i| \times F_{ver}, 1.0). \quad (2)$$

The coalescence depends on the magnitude of velocity in the cell (see Table 1 for relevant notation). We separate horizontal F_{hor} and vertical F_{ver} scaling factors, since the vertical velocity is a procedural rather than physical value.

Second, if this potential precipitation exceeds a calculated resistance Ω_{fall} , a portion of the precipitation $\Delta\rho$ falls to the layer below, with scaling factors from the timestep (Δt) and the height between the layers (h) to prevent instantaneous transport:

$$\Delta\rho = \begin{cases} \rho_{pot} \cdot F_{base} \cdot \Delta t \times \frac{1}{h} & \text{if } \rho_{pot} > \Omega_{fall} \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

$$\Omega_{fall} = \Omega_{base} + U_i \times C_L \quad (4)$$

If the precipitation reaches the ground, it is removed from the atmospheric system and added to the terrain.

4.2. Transient condensation clouds

In explosions, such as bomb blasts and volcanic eruptions, the initial shock wave can be powerful enough to cause transient condensation clouds to form in humid conditions during pressure wave propagation. The condensation cloud is formed by an adiabatic drop in air temperature as the pressure wave passes. This causes the moisture in the humid air to condense out and form a cloud of droplets. After a few seconds the air temperature returns to normal and the droplets evaporate again.

One way to account for this would be to directly decrease and then restore pressure in the atmospheric system for the affected cells. However, as was the case for Vimont *et al.* [VGL*20], there is no direct conversion from dimensionless pressure values to kilo-Pascal (kPa) units, which renders this approach unsuited to the phase transition equations. Instead, we apply a temporary downward adjustment in temperature used during phase transition (with pressure values sourced from standard barometric equations [BC09]). This is realized by simulating a pressure wave centred on the volcano vent Q with initial non-dimensional strength N .

Once initiated, the pressure wave propagates outward at the speed of sound. For each cell C_i that is within the spherical domain defined by the pressure wave, we calculate the distance D_i between C_i and Q . The strength decays as the square of the radius, and the influence on a cell's temperature calculated for use in the phase transition equations is N/D_i^2 . Some seconds after the wave first passes cell i , its effect on temperature is removed and adiabatically condensed water in cell i is allowed to evaporate.

5. Plume model

To represent the evolving volcanic plume we adapt the proxy-geometry approach of Lastic *et al.* [LRC*22], in which the core dynamics of the eruption are approximated using cylindrical slices, with temperature, ash density and moisture attributes, simulated as Lagrangian particles subject to gravity, buoyancy and global wind forces. The acceleration of each slice is computed by applying all contributing forces, and then integrated to update velocity and position.

We also simulate air entrainment, which is vital to the sustained rise of the plume in a Plinian regime. Due to the turbulence of the column, the surrounding air is incorporated, rapidly decreasing the plumes density and increasing its buoyancy. We consider the entrainment velocity U_e , which is the velocity of the air gathered inwards and upwards as a function of prevailing horizontal wind velocity, vertical plume velocity and plume orientation, and use this to compute the volume of entrained air. This leads to a revision of the temperature, volume, and density of slices.

Next, a two-layers hierarchy of primary and secondary spheres is attached to the slices, both to delineate the plume's shape for rendering purposes and to model local convection. The latter harnesses the angular velocity of primary spheres Ω_p to rotate the attached secondary spheres in an orbital fashion, with $\Omega_p = (U_p/r_p)\mathbf{T}_p$, where U_p is the vertical velocity of the sphere, r_p its radius, and \mathbf{T}_p is the unit vector tangent to the sphere. Primary spheres are originally emitted by the slices, and then move freely, using the velocity of the closest slice.

We make two significant alterations to this plume model, detailed next: First, we enhance the proxy spheres so that they both accumulate electrical charge and transport and deposit ash, elements necessary for lightning and ashfall effects, respectively. Second, we add deformable water vapour sheets, initially anchored to their source atmospheric layers, but which are subsequently uplifted by the plume to form skirt, bell and cap clouds.

5.1. Sphere enhancement

We promote the primary and secondary spheres to a more central rôle by having them track the local quantities necessary to render lightning and ash-fall, as well as for more accurate atmospheric transfer during simulation.

To be precise, in addition to the slice temperature present in the prior model, we now track ash A , water vapour W and electrical charge L for both cylindrical slices and spheres. Simulation updates for temperature, ash and water vapour are vested in the slices, copied to the spheres and then transferred to the atmosphere

(Section 6). Primary spheres handle temperature and water vapour, where precision is not needed, while secondary spheres are used for ash loss, which requires a precise spatial location. Electrical charge is a special case and is simulated entirely within the spheres with no transfers to or from slices or the atmosphere (Section 5.2).

In addition, during the umbrella phase, once the plume reaches neutral buoyancy, spheres gradually shrink vertically as a function of time and wind speed to replicate the vertical thinning and elongation of the plume as it moves downwind away from the vent. This effect is known as a gravity current intrusion [WK94] and is caused by movement along the plane of neutral buoyancy in a stratified environment.

This enhanced sphere model has two benefits: first, it allows capture, transport and dispersion of atmospheric quantities at a finer resolution than possible with slices, thereby improving simulation accuracy and, second, it enables a better visual depiction of the variation between individual spheres during rendering.

5.2. Lightning

The occurrence and causes of volcanic lightning have been studied both through radio-frequency measurement of live volcanoes and shock-tube and current-impulse laboratory experiments [CG22]. This has led to a number of working hypotheses as to the dominant electrification mechanisms in volcanic plumes, including:

- **Fracto-electrification:** This charge accumulates during tephra formation as hardened magma shatters in the conduit and on exit from the vent. We approximate this by providing slices with an initial constant charge at vent altitude.
- **Tribo-electrification:** In the jet and convective phases, friction between heterogeneous materials with differing chemical composition (glass, minerals, and rock fragments) and granularity (size, density and shape), leads to a build-up of charge. Interestingly, laboratory experiments [SCG*19, MHCMD20] have shown that in the absence of ice, water vapour has a significant dampening effect on this charge. We factor this in by adding an exponential decay based on water vapour content for charging below the -20C isotherm.
- **Water and hydro-meteors:** Friction between supercooled water droplets and ice flakes occurs at higher altitudes typical of the umbrella phase in a similar fashion to electrification in thunderstorms, except that the effect is magnified by the accretion of ice around ash particles. Evidence points to this charging effects being quite a bit stronger than tribo-electrification [CG22].

To account for these sources of friction, we represent electrical charge using an accumulated dimensionless charge value tracked for all sphere centres p and updated by δL_p on each plume timestep, as follows:

$$\delta L_i = \begin{cases} k_1 \cdot \bar{A}_p(k_{hA}) \cdot \|\mathbf{V}_p\|(k_{hV}) \cdot e^{-10\bar{W}_p(k_M)} & \text{if } T > 253.15\text{K} \\ k_2 \cdot \bar{C}_p(k_C) \cdot \bar{H}_p + \bar{U}_p(k_{cV}) & \text{otherwise} \end{cases} \quad (5)$$

This distinguishes between tribo-electrification below and hydro-meteor electrification above the -20C (253.15K) isotherm. The former includes terms for ash density, sphere velocity, and

an exponential decrease with water vapour, while the latter accounts for atmospheric velocity, ice and ash (with notation defined in Table 1). The bar function denotes a capped normalization $\bar{f}(k) = \min(f, k)/k$. We use the following coefficients in practice: $k_1 = 0.1$, $k_2 = 0.5$, $k_{hA} = 100$, $k_{cA} = 5$, $k_{hV} = 400$, $k_{cV} = 20$, $k_W = 50$, $k_C = 50$.

The actual discharge in the form of lightning is managed through procedural amplification (see Section 7).

5.3. Skirt, bell, and cap clouds

While the horizontal layers used in our atmospheric model (Section 4), positioned at appropriate altitudes, are set to capture the stratification of typical weather conditions, they are poorly suited for modeling air parcels in the immediate vicinity of the plume, which can be subject to significant uplift and distortion.

Instead, we represent these structurally as deformable sheet surfaces sampled vertex-wise at a much finer resolution than the atmospheric grid. Sheets are instantiated as horizontal discs centered on the vent and spaced to represent naturally occurring alternating layers of moist and dry air prior to eruption. This represents finer detail than is present in the atmosphere stack, with up to 3 or 4 sheets per atmosphere layer.

Each sheet vertex s is initialised with water vapour W_s and condensed water C_s taken from the closest atmospheric grid cell. During simulation sheet vertices are subject to vertical uplift forces from the plume spheres. On each plume step, a sphere p has its vertical velocity V_{pz} transferred to a sheet vertex s , weighted according to the distance between them $\phi(s-p)$. This weighting function is set to unity within the inner radius of the sphere r_s and tails off differently within an outer horizontal r_h and vertical r_v offset region, according to:

$$\phi(v) = \begin{cases} 1 & \text{if } \|v\| \leq r_s \\ S_1(1-d) & \text{if } \|v\| > r_s \text{ \& } d = \|(\bar{v}_x/r_h, \bar{v}_y/r_h, \bar{v}_z/r_v)\| < 1 \\ 0 & \text{otherwise,} \end{cases} \quad (6)$$

where $S_1(t) = 3t^2 - 2t^3$ is the smoothstep function and $\bar{v} = v - v/\|v\| \cdot r_s$ is the vector from the sphere surface to the sheet sample. The ellipsoidal outer weighting region (r_h, r_v) allows us to procedurally model the difference between air parcels pushed ahead of the plume and those dragged alongside it, which do not rise as fast. This process continues until the sheet reaches a preset neutral buoyancy altitude.

During uplift the sheet vertices are subject to adiabatic cooling and potential condensation of their water vapour content. Their deformation then leads to the formation of characteristic skirt, bell and cap clouds.

6. Plume-atmosphere transfers

Effectively coupling the atmosphere and plume sub-models requires a transfer of physical quantities (temperature, moisture, wind, and ash) to replicate the processes that take place in nature.

Our first challenge is handling differences in simulation fre-

quency. The volcanic plume involves higher velocities that necessitate smaller timesteps. In our framework, we employ a ratio of 50 to 150 plume timesteps ($dt = 0.02s$) for every atmospheric timestep ($dt = 1$ to $3s$). We synchronise based on the lower frequency, transferring data from plume to atmosphere immediately before each atmosphere update step and in reverse from atmosphere to plume immediately thereafter.

The second challenge lies in the differences in spatial structure: a collection of cylindrical slices and spheres versus stacked grids.

To transfer data from plume to atmosphere, for each air layer ℓ , we consider the plume slices that intersect the layer's altitude range and select the one that lies closest to the layer's central altitude Θ_ℓ . Then grid cells i intersecting the plume slice s obtain values (for T, C, W) from s . The adjustments made to a given cell i 's values are a linear function of the distance from the slice centre to the cell centre, with the adjustment equal to the relevant value of s if the distance is zero and no adjustment if the distance is greater than the radius of s . Note that this transfer relies only on plume slices and not plume spheres.

Ash is a special case, in that there needs to be a close visual correspondence between plume spheres and the source of the ashfall in the atmosphere. Between atmosphere updates, each plume sphere p keeps an accumulated count of ash loss A_p , which is transferred to the closest grid cell i during the plume to atmosphere transfer stage. This ash is then subject to precipitation in a similar fashion to rain once sufficient coalescence has occurred (see Section 4.1).

To transfer data from atmosphere to plume, we consider an annular region in the air layers, corresponding roughly to the region where air is entrained by the plume. This region has an inner radius equal to the extent of the plume r_{plume} and an external radius of $3r_{plume}$. For each air layer, we compute a mean value (for H, T, C, W) over the cells in this annulus and use linear interpolation to deduce values for altitudes that lie inbetween the layer centers Θ_ℓ , before assigning the averaged and interpolated values to a slice.

In summary the following quantities are transferred:

- Temperature (in both directions) to recreate the mixing of plume and atmosphere that gives rise to plume buoyancy and localised atmospheric convection.
- Water vapour and condensed water (in both directions) so that the entrainment and release of moist air can be modelled.
- Horizontal wind velocity (from atmosphere to plume) so that the plume is subject to the appropriate downwind push.
- Ash (from plume to atmosphere) so that ash can mix with water vapour in the atmosphere and fall subject to wind advection.

7. Procedural amplification

The simulation models are run at a level of detail sufficient to capture the key dynamics of plumes, embedded clouds and atmospheric layers, while ensuring manageable execution times. It is, however, necessary to incorporate finer-scale detail for rendering purposes, which we do by procedurally upsampling the simulation data. Fortunately, beyond ash and cloud density, this process can be supported by a rich set of ancillary data (altitude, charge, water

vapour, advective and convective velocities). Moreover, temporal consistency can be ensured by backtracking the velocity field.

Plume: Solid spheres do not combine well with volumetric rendering, since a naïve rasterization would require an impractically dense voxel sampling. Instead, we combine a distance field, flow noise [PN01], and vector backtracking to achieve plausible and consistent animation at reasonable voxel resolutions (100m on a voxel side). This requires storing, per voxel: (1) the distance d_p to the closest plume sphere p ; (2) sphere density A_p , and (3) a vector displacement from the current to initial sphere position \mathbf{B}_p .

During rendering, we use ray-tracing on the distance values in the voxel grid to estimate the location of hard transitions between the plume and surrounding air. This opacity is then modulated according to sphere density. To introduce high frequency detail we subtract flow noise with a period of 50/4 simulation steps from the distance value. This choice of noise function provides the necessary undulating billowy character. For temporal consistency, the noise is seeded based on the original position of the sphere, obtained using an interpolation of \mathbf{B}_p .

Clouds: We create a full volumetric cloud density field by linearly interpolating condensed water between atmospheric layers. Specific cloud types (cirrus, stratus, and cumulus) are then assigned for values above a visibility threshold on the basis of layer altitude and per cell convection [VGL*20]. One issue with computing the rendered atmospheric volume using such linear upsampling is that the vertical density gradients are much weaker than the lateral ones. To compensate, we laterally blur the layers prior to upsampling.

Detail is introduced by warping density with flow noise, based on a combination of density values and cloud type: high frequencies and turbulence for cumulus, lower turbulence for stratus, and lower frequencies for stratus clouds. Density is scaled up for cumulus and down for stratus, to sharpen and smooth them, respectively.

Finally, We use wind velocity $\mathbf{V}(t)$ to recursively backtrack a displacement $displ(t)[x]$ at voxel position x to the cloud source, using: $displ(t+dt)[x] = displ(t)[x - dt\mathbf{V}(t)] - dt\mathbf{V}(t)$.

Over time, this backtracking can lead to discontinuities in the displacement and consequent rendering artefacts. This is overcome by relaxing the displacement – we linearly interpolate between $displ(t+dt)[x]$ and the identity. This interpolation depends on the cloud type: cumulus and stratus clouds are biased towards the identity because their billowy appearance depends primarily on local turbulent phenomena, while cirrus clouds are wind swept, with a low relaxation.

Ash rain: We reproduce the characteristic striated appearance of ash fall by introducing noise with a high horizontal (1 – 3m period) and low vertical frequency (1 – 3km period). To ensure temporal consistency we advect the noise as before, except that the velocity is now predominantly vertical, with an horizontal component matched to the prevailing wind.

Lightning: Drawing inspiration from Reed *et al.*'s work [RW94], we model lightning as a branching wire structure comprising a primary central branch with multiple radiating secondary branches. These secondary branches exhibit random angles (below 45 degrees) with respect to the main branch. Originating within the

plume, each lightning strike randomly connects two highly charged particles. The presence of such a connection is probabilistic, occurring with probability of 0.1 at each time step. To enhance the realism of our visual simulation, we introduce fractional Brownian motion noise [MVN68] as a function of spatial position and time to perturb the line segments composing the wire structure.

8. Results

The simulation sub-models are coded in C++ 20. OpenGL is used for pre-visualization. The procedural upsampling and volumetric rendering are implemented in Houdini 19.5. To aid replication, our code and test scenarios are available at: <https://github.com/jgain/VolcanicSkies.git>.

Results show that our model is able to capture the most salient visual effects induced by the interaction of a volcanic eruption with the ambient weather, including: volcanic lightning (Figure 12), skirt clouds (Figure 13), a cap cloud (Figure 14), transient condensation clouds (Figure 15), and, finally, an eruption whose duration extends to the umbrella phase leading to ash rain (Figure 16). Note that in the earlier phases ash fall is usually hidden within the plume.




Figure 8: Ash deposition on the terrain 20 minutes after eruption. [right] Ash accumulation and colour-coded transects overlaid on the terrain. [left] Graphs of ash deposition along the transects as a function of distance from the volcano vent.

8.1. Validation

In terms of validation we consider: ash transmission through the atmosphere and deposition on the landscape, the incremental build up of electrical charge in the plume, and an ablation of the effect of the advected windfield as compared to a uniform wind vector.

Ashfall: We performed qualitative validation of ash precipitation by investigating the behaviour of ash subsequent to atmospheric transfer both within atmospheric layers and after deposition onto the terrain (see Figure 8 and Figure 9). By recording the accumulation of ash on the terrain, we see that ash fall takes place surrounding the vent from early in the ash rain scenario (Figure 9 a), $t = 204s$). As is to be expected, the intense uplift directly above the vent prevents ash from accumulating there. In its umbrella phase, the plume evolves more gradually and ash saturates the atmosphere to the extent that it becomes visible in the air. On settling out of the plume, ash movement is dictated by atmospheric convection and advection. This is illustrated by Figure 9 - right, which show ash density at the same timestep $t = 477s$, for two layers in elevation




Figure 9: Ash fall density in the atmosphere under a wind from left to right, with darker red indicating greater density. a) Ash accumulated on the terrain before the formation of the umbrella; b) Ash in the top most layer below the umbrella, where most of the ash transfer takes place; c) Ash in the layer immediately below b).

sequence, b) at 17.5km and c) below it at 16.2km. Swept by the wind the ash in c) is less dense and further away from the plume.

In terms of ash deposition over the landscape, Figure 8 shows a map and transects for the depth of ash in centimetres accumulated during the first 20 minutes of a volcanic eruption. In the longer term, ash depth can peak at a few metres for major eruptions, but this involves volcanic activity that can continue for days or even weeks. Unsurprisingly, the deposition profile depends on wind direction. In our case, the peak is some distance from the vent along the wind axis because we only model fine-shard ash as opposed to coarser grains, much of which is born aloft by the convection of the plume. However, the overall deposition does follow an exponential thinning pattern that accords with field studies [Py189].

Lightning: Figure 10 depicts both the incremental and accumulated charge in a medium stage plume. Two aspects are particularly worth noting. First, the charge increment per timestep (Figure 10[left]) increases by about an order of magnitude above the $-20C$ isotherm, indicated by the black rectangle frame in the figure. This is because of the significant increase in charge due to supercooling at this altitude and above (see Equation 5). Second, there are discontinuous patches within the accumulated charge. This is due to discharge as a result of lightning strikes, which we model by zeroing the charge in a spherical region around the anchor points, before resuming charge accumulation.

Windfield ablation: The atmospheric sub-model introduces considerable implementation complexity and so it is worth assessing whether it adds commensurate value to the rendered results. We in-

Scenario	Scene Dimensions	Grid size	Layers	Timestep computation (in s)			
				Plume	Skirt	Atmosphere	Total
Skirt	60 × 60 × 20km	512 × 512	20 { 0.4, 1.2, 2, 2.7, 3, 3.5,	0.44s × 50	1.53s × 50	8.7s	92.2s
Cap and condensation	60 × 60 × 20km	512 × 512	4.1, 4.7, 5.5, 6.5 km, then	0.11s × 50	0.7s × 50	19.1s	59.6s
Ash rain	60 × 60 × 20km	1024 × 1024	evenly spaced up to 20km }	0.53s × 50	No skirts	52.9s	72.9s

Table 2: Parameters for 3 scenarios. The average performance per timestep ($dt = 1s$) is reported, with a breakdown by sub-model.




Figure 10: A visualization of charge accumulation in plume particles, using a colour scale from yellow (low charge) to red (high charge). The outlined frame positioned at 4100m roughly indicates the $-20C$ isotherm. [left] The change in charge over a single timestep, which is approximately an order of magnitude greater above the isotherm. [right] Total accumulated charge. Interspersed lighter regions are a result of discharge due to lightning.

		Resolution		
		256 ²	512 ²	1024 ²
Layers	10	0.86s	3.65s	21.64s
	20	1.66s	8.51s	51.83s
	30	3.12s	15.23s	98.5s

Table 3: Total performance cost per timestep ($dt = 1s$) as a function of the number of atmospheric layers and layer resolution for the ash rain scenario.

investigate this through an ablation study comparing the impact of our advected windfield against a single uniform wind direction per layer. Figure 11 shows the effect this choice has on the ash distribution, with complex eddies and currents evident in the case of advection that are simply absent in the case of uniform wind.

8.2. Performance

Timing tests were conducted on an Intel[®] Core™ i9-9900K CPU @ 5GHz and an NVidia[®] RTX 2080M (8GB). Table 2 shows the computational performance of the sub-models in our simulation. We found that for a single step, the atmosphere model takes much longer than the plume or skirt models, but is run at a much lower frequency, which partially mitigates the issue.

Component-wise, computation of the atmospheric model is dominated (70%) by the pressure solve. Adding layers linearly increases computation time, as does increasing grid resolution (see Table 3). From a quality perspective we have found that around 20 layers and a cell size of 60 – 100m per side is sufficient to capture relevant atmospheric behaviour.

The atmosphere model is more expensive in the cap and condensation scenario than the skirt scenario despite the grid size remaining constant. This is because the phase transition (see Section 4.2) and cloud classification subsystems must be run with a smaller

timestep ($dt = 0.1s$) to adequately capture the transient nature of condensation clouds. In comparison, the rest of the atmospheric model runs at a timestep of $dt = 2s$, and in the other scenarios, all parts of the atmosphere model run at the same timestep ($dt = 2s$ for the skirts scenario and $dt = 3s$ for the ash rain scenario).

The computation cost for the plume increases linearly with the number of slices and spheres added as part of plume expansion. The cost of this sub-model, negligible to begin with, can require as much as a second of computation per timestep $\Delta t = 0.02s$ when the simulation reaches $t = 600s$. The total number of timesteps in a scenario in Table 2 thus accounts for differences in this average computational time per timestep. Deforming skirt surfaces is particularly expensive because it requires subsampling the plume trajectory. A more sophisticated scheme for calculating per timestep intersection of skirt and plume would compensate for this. However, once the skirt reaches neutral buoyancy this component of the computation falls away.

8.3. Limitations

While our solution for coupling explosive eruptions with the atmosphere represents a reasonable compromise between accuracy and computation cost, it remains far slower than real time. The resolution of the atmospheric model is also coarse, with cell sizes of 60 or 100m on a side, which can cause some inaccuracies around the plume. Finer resolutions are possible, but come at a significant computation cost. This is infeasible for simulation times of $t = 60s$ or more. Both a GPU implementation and multiresolution atmospheric grids are possible given the nature of the algorithms, and would speed up computation significantly.

In our plume model, we do not differentiate between positive and negative charge and simply connect highly charged regions. While visually acceptable, this is not physically accurate. In the atmospheric model, the inter-layer exchange is procedural and could be




Figure 11: An ablation study contrasting the pattern of ash in the atmosphere under constant wind [left] and using an advection field [right] over a time span of 180s. Advection introduces downwind turbulence and variation in ash concentration.

improved by using a more accurate, 3D model of the atmosphere, again at the cost of more computation time.

Lastly, as with many simulation methods involving fluid dynamics, there is an overhead associated with setting up and tuning the initial conditions and parameters. This process may require some experimentation and adaptation. In particular, the air layers need to be situated at altitudes that best match atmospheric features, such as inversion layers.

9. Conclusion

In this paper, we have developed an integrated framework for the animation of explosive volcanic eruptions that combines efficient, interacting models for plume and atmosphere simulation with procedural amplification. For the first time, this enables the graphical



Figure 12: Volcanic lightning.

portrayal of salient secondary effects, including ashfall, enmeshed clouds, and lightning. Furthermore, the inclusion of a windfield allows the plume to more realistically track crosswinds and eddies. Performance-wise, the atmospheric simulation is stable, despite the high energy dynamics. Computation times depend on the number of layers and their resolution, but are generally under 2 minutes per second of simulation. Since execution is dominated by advection this would be improved by a GPU implementation [CLT07].

An exciting future extension would be to integrate sub-models for lava flow and topographic deformation to our system, in order to create a unifying framework that incorporates all the visual aspects of volcanic eruption, across the range of eruption types. In particular, while lava flows have been previously modeled in computer graphics [SAC*99, ZKL*17], this is not the case for lahar flows, which are challenging because they combine pyroclasts, water, and debris, with significant variations in temperature and viscosity.

Acknowledgements

This project was funded by the Agence Nationale de la Recherche (Invterra ANR-22-CE33-0012-01), M.-P. Cani's Hi!Paris fellowship, and the National Research Foundation of South Africa (Grant Number: 129257). J. Chen acknowledges support from the Uber chair at LIX.

References

- [Bar82] BARR S.: Skirt clouds associated with the soufriere eruption of 17 april 1979. *Science, New Series* 26, 4550 (1982), 1111–1112. 4
- [BC09] BARRY R., CHORLEY R.: *Atmosphere, Weather and Climate*, 9 ed. Routledge, 2009. 6
- [Buc96] BUCK A.: Buck research cr-1a user's manual. *Buck Research Instruments: Boulder, CO, USA* (1996). 5
- [CG22] CIMARELLI C., GENAREAU K.: A review of volcanic electrification of the atmosphere and volcanic lightning. *Journal of Volcanology and Geothermal Research* 422 (2022), 107449. 7
- [CGW18] CAS R. A., GIORDANO G., WRIGHT J. V.: *Volcanology: Processes, Deposits, Geology and Resources*. Springer, 2018. 4
- [CLT07] CRANE K., LLAMAS I., TARIQ S.: Real-time simulation and rendering of 3d fluids. In *GPU Gems 3*. 2007, ch. 30, pp. 633–675. 11




Figure 13: A triple layered skirt cloud pushed ahead and then merging with the plume, over a span of $t = 60s$.



Figure 14: Cap cloud formation and absorption over a span of $t = 30s$.




Figure 15: Transient condensation cloud formation over a span of $t = 6.5s$ resulting from an explosion at the plume base. [Far right] A closeup image around the crater (indicated by the white triangle) of the moisture layer at 3000m showing that the transient clouds mirror atmospheric structure. Lighter blue indicates higher moisture content, particularly towards the left. The blue line in the left most image shows the moisture layer location relative to the plume.

- [DR06] DOBRAN F., RAMOS J.: Global volcanic simulation: Physical modeling, numerics, and computer implementation. *Developments in Volcanology* 8 (2006), 311–372. 3
- [FSJ01] FEDKIW R., STAM J., JENSEN H. W.: Visual simulation of smoke. *SIGGRAPH* (2001), 15–22. 3
- [GBW97] GLAZE L. S., BALOGA S. M., WILSON L.: Transport of atmospheric water vapor by volcanic eruption columns. *Journal of Geophysical Research: Atmospheres* 102, D5 (1997), 6099–6108. 3
- [GL94] GILBERT J., LANE S.: The origin of accretionary lapilli. *Bulletin*

of Volcanology 56 (11 1994), 398–411. 3

- [HBP*21] HÄDRICH T., BANUTI D. T., PALUBICKI W., PIRK S., MICHELS D. L.: Fire in paradise: Mesoscale simulation of wildfires. *ACM Trans. Graph.* 40, 4 (jul 2021). 3
- [HHP*21] HERRERA J. A. A., HÄDRICH T., PALUBICKI W., BANUTI D. T., PIRK S., MICHELS D. L.: Weatherscapes: Nowcasting heat transfer and water continuity. *ACM Trans. Graph.* 40, 6 (dec 2021). 3
- [HMP*20] HÄDRICH T., MAKOWSKI M., PALUBICKI W., BANUTI




Figure 16: Ash fall in the umbrella phase of an eruption over a span of $t = 200s$.

- D. T., PIRK S., MICHELS D. L.: Stormscapes: Simulating Cloud Dynamics in the Now. *ACM TOG, Proc. SIGGRAPH* (2020). 2, 3
- [JWL*08] JAMES M., WILSON L., LANE S., GILBERT J., MATHER T., HARRISON R. G., MARTIN R.: Electrical charging of volcanic plumes. *Space Science Reviews* 137 (2008), 399–418. 3
- [LRC*22] LASTIC M., ROHMER D., CORDONNIER G., JAUPART C., NEYRET F., CANI M.-P.: Interactive simulation of plume and pyroclastic volcanic ejections. *Proc. ACM Comput. Graph. Interact. Tech.* 5, 1 (may 2022). 2, 3, 4, 6
- [MDN02] MIZUNO R., DOBASHI Y., NISHITA T.: Volcanic smoke animation using CML. *Proc. Int. Computer Symposium* (2002). 2, 3
- [MHCDM20] MÉNDEZ-HARPER J., COURTLAND L., DUFEK J., MCADAMS J.: Microphysical effects of water content and temperature on the triboelectrification of volcanic ash on long time scales. *Journal of Geophysical Research: Atmospheres* 125, 14 (2020), e2019JD031498. 7
- [MTT56] MORTON B. R., TAYLOR G., TURNER J.: Turbulent gravitational convection from maintained and instantaneous sources. *Proc. R. Soc. London* 234 (1956), 1—23. 3
- [MVN68] MANDELBROT B. B., VAN NESS J. W.: Fractional brownian motions, fractional noises and applications. *SIAM review* 10, 4 (1968), 422–437. 9
- [NEOdMVC22] NERI A., ESPOSTI ONGARO T., DE' MICHELI VITURI M., CERMINARA M.: *Multiphase Flow Modeling of Explosive Volcanic Eruptions*. Springer International Publishing, Cham, 2022, pp. 243–281. 3
- [NKL*07] NARAIN R., KWATRA V., LEE H.-P., KIM T., CARLSON M., LIN M. C.: Feature-guided dynamic texture synthesis on continuous flows. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques* (2007), EGSR'07, Eurographics Association, p. 361–370. 3
- [PMG*22] PAŁUBICKI W., MAKOWSKI M., GAJDA W., HÄDRICH T., MICHELS D. L., PIRK S.: Ecoclimates: Climate-response modeling of vegetation. *ACM Trans. Graph.* 41, 4 (jul 2022). 3
- [PN01] PERLIN K., NEYRET F.: Flow noise. In *Siggraph Technical Sketches and Applications* (Aug 2001), p. 187. 8
- [PW09] PARFITT L., WILSON L.: *Fundamentals of physical volcanology*. John Wiley & Sons, 2009. 4
- [Pyl89] PYLE D.: The thickness, volume and grain size of tephra fall deposits. *Bulletin of Volcanology* 51 (jan 1989), 1–15. 9
- [RLB19] RUMBLE J. R., LIDE D. R., BRUNO T. J.: *CRC handbook of chemistry and physics: a ready-reference book of chemical and physical data, 2019–2020, 100th edition*. ed. CRC Press, Boca Raton, FL, 2019. 5
- [RLW*21] RAVURI S., LENC K., WILLSON M., KANGIN D., LAM R., MIROWSKI P., FITZSIMONS M., ATHANASSIADOU M., KASHEM S., MADGE S., ET AL.: Skillful precipitation nowcasting using deep generative models of radar. *Nature* 597, 7878 (2021), 672–677. 6
- [RW94] REED T., WYVILL B.: Visual simulation of lightning. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques* (1994), pp. 359–364. 8
- [SAC*99] STORA D., AGLIATI P.-O., CANI M.-P., NEYRET F., GASCUÉL J.-D.: Animating Lava Flows. In *Graphics Interface (GI'99) Proceedings* (Kingston, Ontario, Canada, 1999), pp. 203–210. 3, 11
- [SCC*16] SUZUKI Y., COSTA A., CERMINARA M., ONGARO T. E., HERZOG M., EATON A. V., DENBY L.: Inter-comparison of three-dimensional models of volcanic plumes. *Journal of Volcanology and Geothermal Research* 326 (2016), 26–42. 3
- [SCG*19] STERN S., CIMARELLI C., GAUDIN D., SCHEU B., DINGWELL D.: Electrification of experimental volcanic jets with varying water content and temperature. *Geophysical Research Letters* 46, 20 (2019), 11136–11145. 7
- [SEH*20] SØNDERBY C. K., ESPEHOLT L., HEEK J., DEGHANI M., OLIVER A., SALIMANS T., AGRAWAL S., HICKEY J., KALCHBRENNER N.: Metnet: A neural weather model for precipitation forecasting. *arXiv preprint arXiv:2003.12140* (2020). 6
- [SGVE*21] SMITH C. M., GAUDIN D., VAN EATON A. R., BEHNKE S. A., READER S., THOMAS R. J., EDENS H., MCNUTT S. R., CIMARELLI C.: Impulsive volcanic plumes generate volcanic lightning and vent discharges: A statistical analysis of sakurajima volcano in 2015. *Geophysical Research Letters* 48, 11 (2021), e2020GL092323. 3
- [SK10] SUZUKI Y., KOYAGUCHI T.: Numerical determination of the efficiency of entrainment in volcanic eruption columns. *Geophysical Research Letters* 37, 5 (2010), L05302–. 3
- [Sta99] STAM J.: Stable fluids. *SIGGRAPH* (1999), 121–128. 5
- [TGL*05] TEXTOR C., GRAF H.-F., LONGO A., NERI A., ONGARO T. E., PAPAŁE P., TIMMRECK C., ERNST G.: Numerical simulation of explosive volcanic eruptions from the conduit flow to global atmospheric scales. *Annals of Geophysics* 48 (2005). 3
- [VGL*20] VIMONT U., GAIN J., LASTIC M., CORDONNIER G., ABIODUN B., CANI M.-P.: Interactive meso-scale simulation of skyscapes. *Computer Graphics Forum* 39, 2 (2020), 585–596. 2, 3, 4, 5, 6, 8
- [WCGG18] WEBANCK A., CORTIAL Y., GUÉRIN E., GALIN E.: Procedural Cloudscapes. *Computer Graphics Forum* 37, 2 (2018). 3
- [WK94] WOODS A., KIENLE J.: The dynamics and thermodynamics of volcanic clouds: Theory and observations from the April 15 and April 21, 1990 eruptions of Redoubt Volcano, Alaska. *Journal of Volcanology and Geothermal Research* 62 (1994), 273–299. 7
- [WXCT19] WERHAHN M., XIE Y., CHU M., THUERER N.: A Multi-Pass GAN for Fluid Flow Super-Resolution. *Symp. on Computer Animation* (2019). 3
- [YI07] YOKOO A., ISHIHARA K.: Analysis of pressure waves observed in Sakurajima eruption movies. *Earth Planets Space* 59 (2007), 177–181. 4
- [ZKL*17] ZHANG S., KONG F., LI C., WANG C., QIN H.: Hybrid modeling of multiphysical processes for particle-based volcano animation. *Computer Animation & Virtual Worlds* 28 (2017). 3, 11