




An empirically derived adjustable model for particle size distributions in advection fog

M. Kolářová¹ , L. Lachiver¹  and A. Wilkie¹ 

¹ Faculty of Mathematics and Physics, Charles University, Prague, Czech Republic



Figure 1: A foggy urban scene rendered with our proposed model (left), and a thin (middle) and dense (right) homogeneous fog layer for comparison. Our model features a realistic vertical density profile: note that visibility at ground level is significantly better than e.g. for the top of the high rise buildings. Like in many real fog situations, the tops of high buildings disappear in the fog, while horizontal visibility along the street is still reasonable. Of course it is also possible to model a vertical fog density gradient in an ad hoc fashion: but our model is based on real fog particle distribution measurements, and can therefore be used for predictive and reference rendering purposes.

Abstract

Realistically modelled atmospheric phenomena are a long-standing research topic in rendering. While significant progress has been made in modelling clear skies and clouds, fog has often been simplified as a medium that is homogeneous throughout, or as a simple density gradient. However, these approximations neglect the characteristic variations real advection fog shows throughout its vertical span, and do not provide the particle distribution data needed for accurate rendering. Based on data from meteorological literature, we developed an analytical model that yields the distribution of particle size as a function of altitude within an advection fog layer. The thickness of the fog layer is an additional input parameter, so that fog layers of varying thickness can be realistically represented. We also demonstrate that based on Mie scattering, one can easily integrate this model into a Monte Carlo renderer. Our model is the first ever non-trivial volumetric model for advection fog that is based on real measurement data, and that contains all the components needed for inclusion in a modern renderer. The model is provided as open source component, and can serve as reference for rendering problems that involve fog layers.

CCS Concepts

• *Computing methodologies* → *Computer graphics*;

1. Introduction

Realistic computer graphics has progressed to the point where most synthetic scenes that 3D artists can think of can be rendered to look perfectly convincing to a human observer. Challenges arise, though, when the resulting images additionally need to be reliable predictions of scene appearance. Such accuracy is not frequently needed for mainstream CG work, so most current rendering software is not designed with this in mind.

However, predictive rendering can be essential for a variety of

specialised tasks, such as virtual training of autonomous vehicles and aircraft. There, systematic exploration of the parameter space of outdoor scenery is important: this includes weather and lighting conditions. The use of synthetic imagery for the training of such systems allows exact control over the training process, and gives confidence in the reliability of the trained system. But this only holds if one is able to generate the training imagery in a truly predictive fashion in the first place: and reliably accurate models of all weather phenomena that can occur in real environments are one of the needed components in such a system.

In this paper, we propose a reference model for advection fog that is based on real world measurements, and that captures the characteristic profiles of particle distribution and density that are typical for advection fog (in particular, that fog density increases with altitude, before falling off again at the top of the fog layer). Our model is formulated in a way that allows inclusion in modern path tracers, and will be provided as an Open Source component.

In renderings that use the new model, the visual appearance of the resulting fog is not markedly different from simpler approximations: but as Figure 1 shows, some differences are still apparent. Prior to us deriving this model, it was an open question how large the difference between a truly realistic fog model that is based on actual measurements, and the currently commonly used simple approximations would be. Our work closes this knowledge gap, and allows those who are working in the domain of predictive atmospheric rendering to switch to a reliably realistic fog model in those cases where the additional accuracy is needed.

2. Optical properties of atmospheric phenomena

Here we briefly provide background on the sort of interactions light can have with atmospheric components, how one can model the distribution of such constituents of a clear atmosphere, and what the properties of fog layers are.

2.1. Atmospheric scattering

From the viewpoint of light transport, there are three main mechanisms of light interacting with matter in the atmosphere: scattering on larger particles (Mie scattering), scattering by molecules (Rayleigh scattering), and plain absorption [EBSK*16]. As each constituent element of the atmosphere gives rise to one main interaction mode, atmospheric data models such as OPAC [HKS98] sort atmospheric constituents according to the main interaction they cause. Fog, like clouds in general, usually consists of water droplets (rarely also ice crystals), for which the interaction of light can be described by Mie scattering.

2.2. Clear atmospheres

As discussed in great detail in [EBSK*16], and summarised from a graphics viewpoint in [WVBR*21], even the composition of clear atmospheres is far from trivial. The thing to note with regard to our goal (fog rendering) is that as shown in Figures 6 and 13 of [WVBR*21], even realistic vertical profiles of Mie scatterer density in seemingly simple clear atmospheres exhibit a non-exponential profile that makes analytical integration of atmospheric light transmission impossible, and that requires a tabulation approach to handle correctly. So we can also expect the structure of realistic fog layers to potentially be more complex than one would naively expect.

2.3. Fog formation

Fog is formed by a phenomenon called heterogeneous nucleation [Pru10]. With the decrease in air temperature higher relative humidity is achieved and water vapour starts to condensate around

condensation nuclei. These are mainly aerosol particles naturally present in clear atmospheres that are capable of initiating drop formation even at low supersaturation (relative humidity 101%). Eventually, if supersaturation is high enough, all aerosol particles are able to initiate drop formation. The concentration of aerosol particles in a particular air mass can be seen as an indicator of the number of drops that will be formed. Concentration of fog drops can be expressed in terms of number density which defines the number of drops per chosen unit volume, e.g. $[\text{m}^{-3}]$.

2.4. Fog structure

Fog is described as a cloud of small water droplets that is near ground level and sufficiently dense to reduce visibility to < 1000 m.

While there are scenarios where the differences between fog and clouds can become blurry (e.g. in the case of rising fog banks), the main differences to clouds are that fog is an aerosol which usually is

1. horizontally spread out more or less evenly over a larger area (unlike individual clouds, which are more localised), and
2. in most cases is reasonably close to the ground.

There are several fog types and they are usually defined by mechanism that led to their formation, though several mechanisms combined can lead to fog formation. Most common ones over land and most stable fog types are *radiation fog* and *advection fog*.

Radiation fog is formed by emissive cooling of the ground after sunset, usually on clear nights. It is often shallow, with a thickness of a few metres, to maybe tens of metres. It is often called "ground fog".

By contrast, advection fog layers frequently reach several hundred metres of thickness: these are formed when a large moving mass of moist air is cooled down by external factors. The typical scenario is moist air being pushed over cooler water or land by the wind. For instance, along the coast of California (see Figure 2), ocean currents bring cold water to the surface next to the shore, so advection fog based on moist sea air being pushed inland over this cool water by the prevailing westerly winds is a regular occurrence there.

2.5. Fog microstructure

The crucial characteristic of fog from which the optical properties can be modelled is the particle size distribution. Particles in fog are represented by water droplets with the diameters reaching up to $50 \mu\text{m}$.

One of the most common realistic particle size distribution models for homogeneous fog has been proposed by Deirmendjian [Dei69]. The distribution curve is characterised by a modified Γ -distribution. The same distribution type is used by Shettle and Fenn [SF79]. Moreover, the authors have proposed a four class topology model for two types of most stable fogs, advection and radiation fog. Both of these types have two sets of estimated parameters, each set describing their different densities. The particle size range this model is applicable to is $2\text{--}50 \mu\text{m}$. In this model, it is also assumed that fog density is uniform throughout the entire

volume, and that the distribution curve shape does not vary with altitude. However, the value of the b parameter, which determines the shape of the curve, varies among the same fog type for the different densities. This indicates that the shape of the curve is actually influenced by the fog density.

Beside the modified Γ -distribution, that is mostly used, another common type of distribution function used to describe fog particles is the lognormal distribution [KD14]. In order to fit the observed multimodality, the sum of several lognormal distributions is used.

While these realistic models are based on actual data measurements, they only accurately characterize the low altitudes of fog, as the data is acquired at the ground level. The measurements discussed in section 4.1 show that particle size distributions can vary considerably with altitude: so a more powerful model that also takes this parameter into account is needed.

3. Atmospheric rendering

Here we provide a very brief summary of the state of the art with regard to realistic rendering of atmospheric phenomena in graphics.

3.1. Clear sky models

As rendering of outdoor scenes has been a topic of interest for as long as graphics has been done, there are by now many models of varying accuracy and speed which can be used to describe clear, cloudless skies. An overview of the state of the art, in particular that for high quality offline rendering, can be found in [WVBR*21], where a recent fitted, pre-computed model for the radiance patterns found on prototypical clear days is presented. There is also a large number of more approximative techniques for interactive graphics, with the technique by [Hil20] as a very powerful recent example.

Recent improvements in path guiding [RGH*20, Her22] mean that even for high quality results, brute force computations of realistic atmospheres will likely become practical in the foreseeable future. As a full, unbiased solution that is not based on pre-computations is preferable for predictive rendering purposes anyway, the importance of fitted models like [WVBR*21] will likely decrease in the future. An additional benefit of this will be that no special casing of various atmospheric phenomena is needed anymore: rendering software for outdoor scenes will get simpler once all one has to do is to feed it a model of the atmosphere, including clouds, haze, fog, and whatever phenomenon one wishes to include.

3.2. Fog rendering

So far, computer graphics mostly uses simple phenomenological models for fog: usually, 3D artists tweak the appearance of a colourless scattering medium until the desired appearance has been reached. Most publications deal with rendering fog approximations: initially to get them to render at all [KVH84, PPS97], and then mostly dealing with the issues that arise in real time rendering [LMKA01, BM02, DPLK04, SRNN05]. Other works that involve fog are usually for specific purposes [DPLK05], and those that explicitly deal with fog modelling do so in an ad hoc fashion that is mostly appearance driven [ZHG*07, GB10]. Overall, it is

currently still rare to use physically plausible droplet distributions to define fog: partly because no such model exists that goes beyond providing a single particle size distribution for all altitudes and locations within a fog layer (i.e. homogeneous fog).

4. Our advection fog model

4.1. Measurement data of real fog layers

We use data from a measurement campaign conducted for NASA in 1994 [Zak94] for our work: goal of these measurements was to obtain a better understanding of fog layer composition for aviation safety purposes.



Figure 2: Aerial photograph of advection fog in the California coastal region, similar to the fog that was measured in Arcata. A slightly turbulent top layer of the fog is evident in this image: the fluctuations in the particle size distributions measured during different measurement approaches by [Zak94] are likely mostly caused by similar phenomena. Our fitted model does not cover the turbulence aspect of advection fog: but adding such perturbations to a fog layer is comparatively easy, if and when this is needed. Photograph by Brian Uretsky, used by permission.

For these measurements, a business jet aircraft equipped with external sensors that were able to record particle size distribution data at 10 m altitude increments flew instrument approaches to several airports across the United States. This was done when suitably interesting fog conditions were in place: and at each location, several such approaches were flown in succession. Not all datasets gathered during this campaign are for advection fog: but for those which are, the datasets nicely show the temporal development of a rising advection fog layer. For our work, we concentrate on the dataset gathered at Arcata, CA: this is an example of classic advection fog that can serve as an exemplar of the phenomenon. As discussed in section 4, the measured data was somewhat noisy, due to advection fog layers sometimes being slightly billowy, like in the photograph in Figure 2. However, this does not affect our ability to extract useful particle size distribution trends from the data.

We chose the Arcata fog dataset because it exhibits consistent

behaviour throughout its growth from 100 m to 260 m. Multiple approaches were made during the fog's development phase, and we categorised the dataset into three subsets, each corresponding to a specific overall height: 100 m, 180 m and 260 m. In contrast, the Vandenberg fog measurements were conducted at a constant fog height, resulting in a single dataset corresponding to 180 m. We observed similarities in its properties with those of the Arcata fog dataset at 180 m, and as Figure 5 shows, the fit obtained from the Arcata data actually aligns pretty well with the single dataset from Vandenberg. The Santa Maria advection fog exhibited comparable properties to that of Vandenberg and Arcata fogs, as reported in [Zak94]. The other types of fog in the paper had either not enough data to create a general model, or there were significant inconsistencies among measurement approaches.

4.2. Data Fitting

As mentioned in previous section, three datasets corresponding to a specific overall height were chosen to be analyzed: 100 m, 180 m and 260 m.

The function used for the fitting of Arcata data collected in [Zak94] is a linear combination of lognormal and Gaussian function and the formula is as follows

$$f_{LG}(d) = L \cdot f_{logn}(d) + G \cdot f_{Gauss}(d) \quad (1)$$

$$f_{logn}(d) = \frac{1}{\sqrt{2\pi} \cdot \sigma_L \cdot d} \cdot \exp\left(-\frac{(\ln(d) - \mu_L)^2}{2\sigma_L^2}\right) \quad (2)$$

$$f_{Gauss}(d) = \frac{1}{\sqrt{2\pi} \cdot \sigma_G} \cdot \exp\left(-\frac{(d - \mu_G)^2}{2\sigma_G^2}\right) \quad (3)$$

where coefficients L and G are the scaling parameters of particular curves, providing the areas under lognormal and Gaussian peak, σ_L and μ_L are parameters indirectly characterising the shape of lognormal curve, σ_G and μ_G are parameters that give the shape of Gaussian curve directly, with σ_G being the width of the Gaussian peak and μ_G being the position of Gaussian peak maximum.

More details on the fitting procedure can be found in sections 1.1 and 1.2 of the supplement.

4.3. Developing the general model of advection fog

Following the fitting procedure, the obtained parameters were analyzed and compared among the three datasets, i.e. three development stages of Arcata advection fog. Similar shape development of the particle size distribution curves over the span of its total height has been observed. This led to developing a model, that can generalize this behavior for any given advection fog height.

The model development can be summarized as follows:

- For each of the parameters obtained from the data fitting the dependencies on altitude were normalized
- Normalized data from all 3 datasets was merged for each parameter to create one general dataset

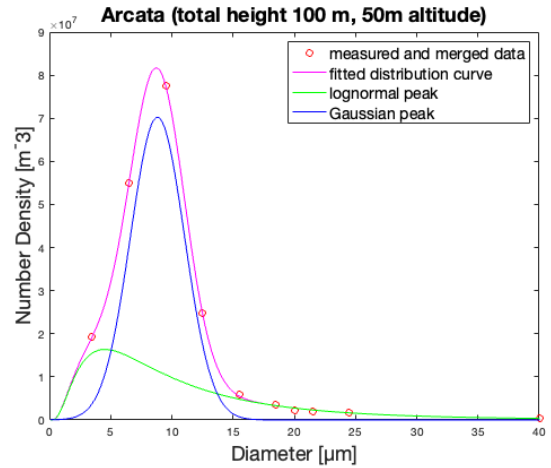


Figure 3: Example of data fit using the sum of lognormal curve (green) and Gaussian curve (blue) in order to obtain the best fit of the data (pink). The actual measured data obtained by averaging measurement approach 1 and 2 are depicted as red circles. This dataset represents Arcata particle distribution at 50 m altitude for the case where fog height reached the total height of 100 m. Estimated parameters for L, μ_L and σ_L were $1.99 \cdot 10^8$, 2.13 and 0.79. Parameter values for G, μ_G and σ_G were $3.89 \cdot 10^8$, 8.82 and 2.21. RRSE of the regression was 2.3%.

- Parameters corresponding to a curve shape (μ_L , σ_L , μ_G and σ_G and G:L area ratio) were fitted in order to get one general dependence on altitude
- The dependencies of parameter corresponding to number density were fitted separately because they are depending on total height of the fog
- Scaling parameter matrix was calculated for the number density dependence, allowing to return a particular number density curve for given total height of the fog

Following all the necessary parameters being fitted, the only input for our model is the total fog height that can be chosen by the user. Model then generates particle size distribution curve for any requested altitudes within its total height.

Detailed description of the model development can be found in the supplement.

4.4. Verification of our results

4.4.1. Reproducibility of the original data

We first generated 3 distribution sets of 3 different fog of heights that exactly correspond to the heights of those 3 input datasets, i.e. we checked if the fitted model (which can be evaluated for arbitrary fog thicknesses) actually returns the original particle distributions when queried for the originally measured thicknesses. The example of resulting 3D plots from our general model in comparison to the actual datasets is shown in Figure 4 for fog height of 180 m. Plots corresponding to the two remaining datasets (100 m and 260 m) can be found in the supplemental.




Figure 4: An example of generated curves for the fog height of 180 m (right column) in comparison to the original data fit of 180 m dataset (left column). An important point to note is that in the bottom view, our fitted model looks like it increases the particle counts beyond what is found in the original data: the "ridgeline" of the plot is significantly higher. But the fit is reasonable insofar as it preserves the area under each measured distribution curve: the shape of the distribution (and in particular, its peak) is also affected by noise, which is smoothed out by the fit. So the overall transmission characteristics of the fog are still properly captured by the fit.

The generated distributions are smoother due to the parameter fitting which eliminated the outliers, but the general trend in shape has been successfully reproduced in all 3 fog datasets. The original fits have been rescaled according to the correct values of particle densities so they can be compared also in terms of their magnitude. This can be seen in the side view of the graphs. However, one must take into consideration that the peak heights might not be exactly the same due to slight changes in the width of the peaks. Generated curves slightly differ in shape parameters due to the fitting procedure. For example, higher values of σ_G in the generated set cause a decrease in the peak height since it enlarges the width of the peak while keeping the same magnitude. Therefore the seemingly lower height of the generated peaks is not an indication of the bad fit. In developing the model, we generalised not only the shape but also the magnitude of the curves, thus suppressing the variation of the measurement data.

4.4.2. Comparison to Vandenberg advection fog data

To claim generality, we analyzed the data from Vandenberg fog and compared the results to our generated curves. The number density profile of Vandenberg data from [Zak94] show that the fog reached 180 m total height and then approach 7 was performed when the fog height was 240 m. The average number density of approaches 3 to 6 was taken to obtain the profile, with the elimination of extreme outlier values. We performed the same data fit as with the Arcata data, and the resulting curve was compared to our modelled curve for 180 m fog. The results shown in Figure 5 indicate that

even at different locations, advection fog type keeps similar number density profile.

In terms of the actual size distribution shapes, the comparison of our modelled curve with the Vandenberg data fit can be seen in Figure 6. While our model very nicely describes the curve shape of the actual Arcata data (left column), in the case of Vandenberg advection fog (right column), we can notice that the distributions are slightly wider. However, the areas of the fit and the modelled curve are comparable which can be expected based on similarity in number densities as shown in Figure 5. This shows that even though advection fog comes from different locations and its shape of distribution might slightly vary, its overall properties are similar and our model can be used for their particle distributions generation.




Figure 5: Comparison of advection fog data from Vandenberg location to our modelled number density profile. Averaged data for Vandenberg fog that reached 180 m are depicted as blue points. Blue line represents the actual fit of these data. Pink line is our modelled density profile for general 180 m advection fog.

5. Using the model in a renderer

In this section, we discuss how our model can be easily implemented into an existing rendering engine.

5.1. Describing a participating medium

In most renderers, participating media are described using three parameters: attenuation coefficient (σ_t), a phase function and an albedo. The albedo of fog is commonly defined as 1, given that it is the albedo of water. Fog is a participating medium whose absorption is negligible, meaning that only scattering coefficient (σ_s) contributes to attenuation and therefore $\sigma_t = \sigma_s$.

5.2. Ray and medium interactions

In volumetric path tracing, the interaction between the ray and the participating medium is usually handled this way [NGH*18]:

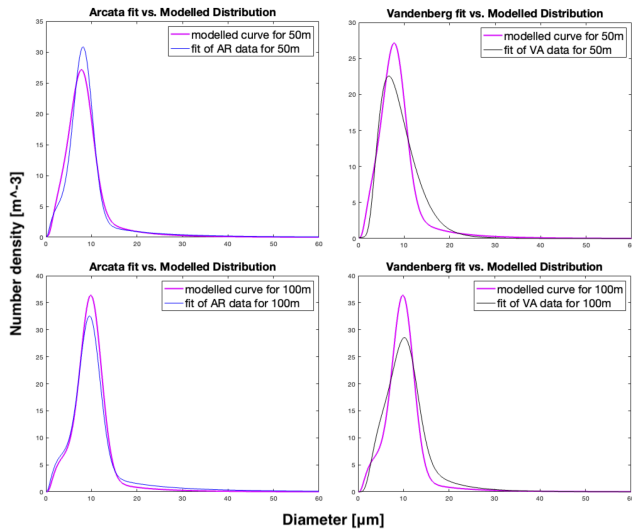


Figure 6: Curves generated from our model (pink) for fog of 180 m height in comparison to actual data fits from Arcata (blue) and Vandenberg (black) for the altitude of 50 m (first row) and 100 m (second row).

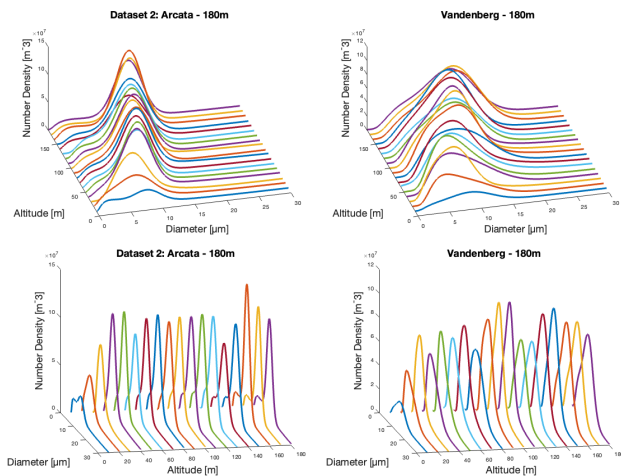


Figure 7: Comparison of 3D plots for advection fogs from different locations that both reached 180 m height. Plots represent fits of actual measurement data. A similar altitude behaviour can be observed, even though the Vandenberg distributions have a more prominent tail, and a slightly wider distribution curve. However, the overall number density per each layer (and by extension, visibility with in the fog) is comparable, as shown in Figure 5.

- Check the intersection between the ray and the medium bounding box
- If there is, sample a distance t along the ray
- Retrieve the medium scattering coefficient and the phase function value for this position and adjust the ray weight and direction accordingly

From the position of the intersection, we calculate the altitude

which is used to get the correct distribution parameters from our polynomials. We can then sample a particle diameter and compute its associated density using the PDF. In our method, the sampling is performed using normal and lognormal distributions from `std::random` library in C++. The algorithm for our particle diameter sampling method can be found in section 3 of the supplement.

Once the diameter and density are sampled, we can calculate the scattering coefficient.

5.3. Calculating the scattering coefficient

The scattering coefficient σ_s can be calculated from a certain particles density N with the following equation:

$$\sigma_s = N * \sigma(d) \quad (4)$$

where N is the number density for a particle of diameter d and its value is obtained from our modelled distributions. σ is the scattering cross section and its value depends on particle's diameter d . However, the value of σ , is trickier to calculate and relies on Mie scattering. We thus pre-calculated its values using `miepython` module [mie23] for a range of diameters between 0.5 to 60 μm and use these tabulated values for equation 4 to calculate the scattering coefficient on the fly. The scattering cross section is normally wavelength-dependent, but we averaged its values since our model is not (yet) intended for spectral rendering.

5.4. Phase function

As stated in subsection 2.1, the interaction between light and fog particles is modelled by Mie scattering. Jendersie et al. [Jd23] propose an analytical model for evaluating the Mie phase function for a given particle diameter. In the supplemental, we provide our C++ implementation of this. For the purpose of this paper we opted for a tabulated phase function. Mitsuba 3 ships a `tabphase` plugin taking phase function values as an array. With the help of `miepython` [mie23] module, we calculated the Mie phase function values for the mean diameter of our distributions and with an angular resolution of 1° .

Also, we used `miepython` to calculate the average asymmetry factor g over the altitude, which gives **0.83**. We can then give this to the frequently used Henyey-Greenstein phase function as an approximation of Mie scattering. We compare renders with these two phase functions in the results section.

5.5. Defining our fog model in a 3D scene description

Once the fog medium class has been added to the renderer, the user can add the fog model in the scene description. To remain physically correct, the fog bounding box must be scaled accordingly to the desired fog height. This fog height is the only parameter needed to describe our fog model, since the scattering coefficients will be calculated on the fly.

5.6. Layered fog model

Instead of implementing our diameter sampling algorithm, one can also extract the average scattering coefficients for a set of altitudes and store them in a volumetric grid. In Mitsuba 3 [JSR*22], heterogeneous media are described by a volumetric grid storing their parameters (scattering coefficient, phase function and albedo). In the supplemental archive, we provide the code for generating this volumetric grid out of our average scattering coefficients, which can then be loaded in a scene. This method produces way faster renderings and can be implemented in a shader for real-time purposes. However, in the results section, we demonstrate that it does not yield the same results as our more physically-based diameter sampling method. We call this model "layered fog model" since it has one scattering coefficient per layer.

5.7. Implementation

Our Mitsuba 3 [JSR*22] plugin implementing our fog model is provided in the supplemental archive.

6. Results

6.1. Visual differences between models and evaluations

For comparisons, we defined a simple Mitsuba 3 [JSR*22] scene for "fog debugging": it shows a stylised street-like setting, with some buildings on the left and right and a tall pole at the end of the street. All these scene elements roughly have the dimensions of real buildings, but are textured by coloured checkerboard patterns of different bright colors, to better show the different effects the various fog evaluation techniques have. There are two versions of this scene, one for a day setting, and a night version. The day version is lit by an environment map with a constant radiance of 1.0 (i.e. diffuse lighting outside the fog layer), while the night version has no atmospheric light sources and instead features the checkerboard textures as emitters. At night, the idea behind using the glowing checkerboards is to maximise visibility of the forward scattering behaviour of Mie scattering, compared to Henyey-Greenstein; and having them emit in red and blue is intended to highlight wavelength-dependent differences. Since it is difficult to obtain a ground truth fog image to compare to, we perform qualitative comparisons of our model against others.

In Fig. 12, we present some renderings obtained with different fog models. The homogeneous fog is described with a unique scattering coefficient which is, in this case, the mean of our averaged scattering coefficients as presented in section 5.6. We also show a render obtained with a linear gradient, since it is a straightforward way to mimic a non-trivial fog that is denser at the top of the fog layer. To define this medium, we exported a volumetric grid storing scattering coefficients which range from our minimum to maximum average ones (averaging scattering coefficients is presented in section 5.6).




Figure 8: From top to bottom: rendering of our fog model using Henyey-Greenstein phase function, using Mie scattering, and the ΔE_{2000} between the two, for the day (left) and night scene (right). The value of g , the Henyey-Greenstein asymmetry factor, is set to 0.83 for a reason explained in section 5.4.




Figure 9: Figure illustrating the differences in visibility predictions using Koschmieder's law. The comparison is made between our model and a baseline model that employs a linear gradient of scattering coefficients.

We employed an error metric known as ΔE_{2000} , widely recognized for its capability to quantify perceptual differences between color variations. Utilizing this metric, we compared the rendering obtained with the linear gradient model against our fog model. The linear gradient, while offering a simplified representation, failed to capture the density variations specific to the fog.

We also compared the renderings using Henyey-Greenstein phase function and Mie scattering in Fig. 8. A notable difference, quantified using the DeltaE metric, became apparent. This difference is especially pronounced during nighttime scenarios. One plausible explanation for this divergence lies in the inherent characteristics of Mie scattering which tends to be more forward scattering. This can lead to a more concentrated brightness in the direction of the light source.




Figure 10: From top to bottom: rendering of our fog model using diameter sampling, using our layered fog model, and the ΔE_{2000} between the two methods, for the day (left) and night scene (right) discussed at the beginning of Section 6.1.

6.2. Visibility predictions

Moreover, we present differences in visibility predictions using two models: Koschmieder's [Kos24] and Allard's [All76] laws, two de facto standard measures of visibility [DB77] that play a significant role in calculations used in the aviation industry. While Koschmieder's law (Fig. 9) offers a relationship between meteorological visibility and the atmospheric extinction coefficient, par-

ticularly useful for daytime scenarios, Allard's law (Fig. 6.2) extends this understanding to nighttime conditions, factoring in the influence of varying light intensities. By juxtaposing the outcomes derived from both models, we provide a comprehensive view of visibility estimations for varying altitudes, and under different lighting conditions. Our realistic fog model leads to significantly different visibility predictions using these two standard approaches, which underlines the point that using such a non-trivial fog model is likely justified for aviation applications.

In both cases, we observe on Fig 9 and Fig. 6.2 that the difference in predicted horizontal visibility is roughly 100 meters at 25 meters of observer altitude, which is already a significant difference for some applications. Furthermore, the visibility for our model is slowly increasing towards the maximum altitude, a behaviour not represented by a simple gradient of densities. All the visibility calculations are provided in our supplemental archive.




Figure 11: Figure illustrating the differences in visibility predictions using Allard's law. The comparison is made between our model and a baseline model that employs a linear gradient of scattering coefficients. Allard's law is designed for nighttime scenarios and thus calculates the visibility of a light source for varying intensity. We here calculate the visibilities for light intensities of 100, 1000 and 10000 candelas.

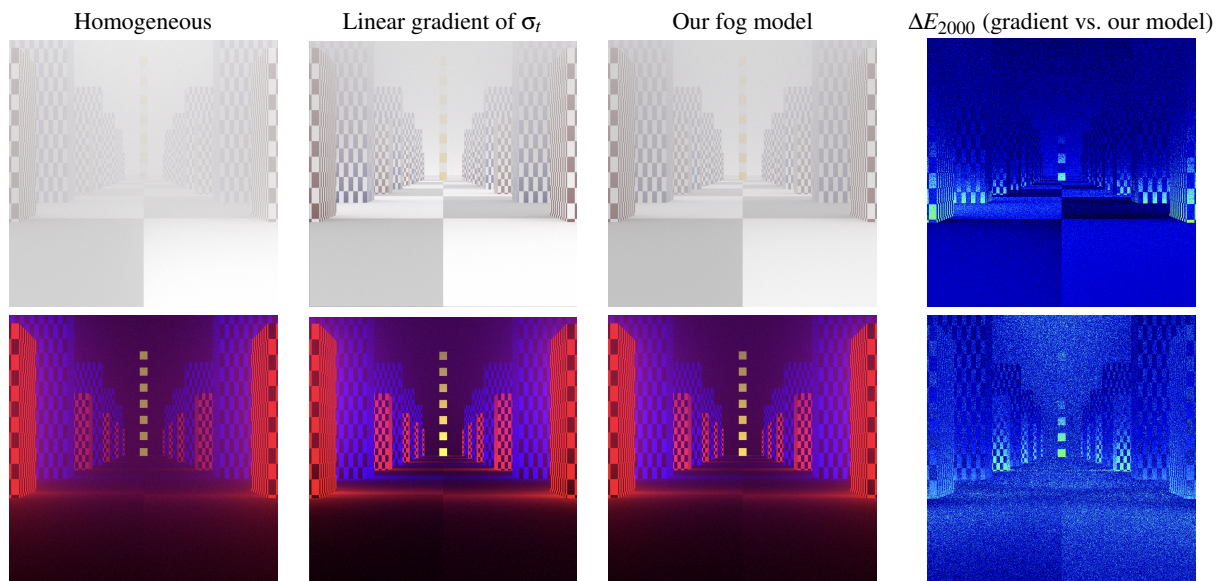


Figure 12: The scene described at the beginning of Section 6.1 (top: day, bottom: night). From left to right: a homogeneous medium, a medium with a linear gradient of scattering coefficients, our 260 m fog model coupled with Mie scattering and the ΔE_{2000} metric between our model VS. linear gradient of scattering coefficients. The scattering coefficient of the homogeneous medium is the mean of our averaged scattering coefficients (see section 5.6). The linear gradient ranges from the minimum to the maximum of our averaged scattering coefficients.

7. Conclusion and outlook

We presented an analytical model of the particle distributions within advection fog, parameterised by the thickness of the fog layer. The model is based on real measurement data, and provides, to our knowledge for the first time within graphics, a possibility to include realistic fog layers in predictive renderings. We also discussed how to use this model in modern Monte Carlo renderers. We provide the full code source of our model compatible with Mitsuba3 [JSR*22] in the supplemental archive.

Future work will be to obtain data for other fog types, and create similarly flexible models for these as well: with radiation fog as the most obvious target for such an effort. Looking forward, we are also keen to take the wavelength-dependent effects of Mie scattering into account.

Acknowledgements

We acknowledge funding by the Grantová Agentura České Republiky with grant number GAČR-22-22875S, the Horizon 2020 Framework Programme of the European Union (grant number 956585, the PRIME ITN), and by Charles University under internal grant SVV-260699.

References

- [All76] ALLARD E.: *Memoire Sur l'Intensite et la Portee des Phares*. Dunod, Paris, France, 1876. 8
- [BM02] BIRI V., MICHELIN S.: Real time animation of realistic fog. In *Poster Session of 13th Eurographic Workshop on Rendering* (2002), pp. 9–16. 3

- [DB77] DOUGLAS A., BOOKER R. L.: *Visual Range: Concepts, Instrumental Determination, and Aviation Applications*. National Bureau of Standards, Washington, USA, 1977. 8
- [Dei69] DEIRMENDJIAN D.: *Electromagnetic Scattering on Spherical Polydispersions*. RAND Corporation, Santa Monica, CA, 1969. 2
- [DPLK04] DUMONT E., PAULMIER G., LECOCQ P., KEMENY A.: Computational and experimental assessment of real-time front-lighting simulation in night-time fog. In *Proceedings of the Europe Driving Simulation Conference* (2004). 3
- [DPLK05] DUMONT E., PAULMIER G., LECOCQ P., KEMENY A.: Assessment of the simulation of low beams back-scattering in fog. In *6th International Symposium on Automotive Lighting, ISAL 2005* (2005). 3
- [EBSK*16] EMDE C., BURAS-SCHNELL R., KYLLING A., MAYER B., GASTEIGER J., HAMANN U., KYLLING J., RICHTER B., PAUSE C., DOWLING T., BUGLIARO L.: The libradtran software package for radiative transfer calculations (version 2.0.1). *Geoscientific Model Development* 9, 5 (2016), 1647–1672. 2
- [GB10] GIROUD A., BIRI V.: Modeling and rendering heterogeneous fog in real-time using b-spline wavelets. 3
- [Her22] HERHOLZ S.: Intel Open Path Guiding Library, 2022. <https://github.com/OpenPathGuidingLibrary/openpgl>. 3
- [Hil20] HILLAIRES S.: A scalable and production ready sky and atmosphere rendering technique. *Comput. Graph. Forum* 39, 4 (2020), 13–22. 3
- [HKS98] HESS M., KOEPKE P., SCHULT I.: Optical properties of aerosols and clouds: The software package opac. *Bulletin of the American Meteorological Society* 79, 5 (1998), 831–844. 2
- [Jd23] JENDERSIE J., D'EON E.: An approximate mie scattering function for fog and cloud rendering. In *ACM SIGGRAPH 2023 Talks* (New York, NY, USA, 2023), SIGGRAPH '23, Association for Computing Machinery. URL: <https://doi.org/10.1145/3587421.3595409>, doi:10.1145/3587421.3595409. 6
- [JSR*22] JAKOB W., SPEIERER S., ROUSSEL N., NIMIER-DAVID M.,

- VICINI D., ZELTNER T., NICOLET B., CRESPO M., LEROY V., ZHANG Z.: Mitsuba 3 renderer, 2022. <https://mitsuba-renderer.org>. 7, 9
- [KD14] KLEIN C., DABAS A.: Relationship between optical extinction and liquid water content in fogs. *Atmospheric Measurement Techniques* (02 2014), 11. 3
- [Kos24] KOSCHMIEDER H.: Theorie der horizontalen sichtweite. *Beiträge zur Physik der freien Atmosphäre* (1924), 33–53. 8
- [KVH84] KAJIYA J. T., VON HERZEN B. P.: Ray tracing volume densities. *ACM SIGGRAPH computer graphics* 18, 3 (1984), 165–174. 3
- [LMKA01] LECOCQ P., MICHELIN S., KEMENY A., ARQUES D.: Lighting simulation with the presence of fog: a real-time rendering solution for driving simulators. In *Proceedings of the Driving Simulation Conference, DSC 2002* (2001). 3
- [mie23] miepython module, 2023. <https://miepython.readthedocs.io/en/latest/>. 6
- [NGH*18] NOVÁK J., GEORGIEV I., HANIKA J., KRIVÁNEK J., JAROSZ W.: Monte Carlo methods for physically based volume rendering. In *ACM SIGGRAPH Courses* (Aug. 2018). doi:10/c5fj. 5
- [PPS97] PÉREZ F., PUEYO X., SILLION F. X.: Global illumination techniques for the simulation of participating media. In *Rendering Techniques '97: Proceedings of the Eurographics Workshop in St. Etienne, France, June 16–18, 1997* 8 (1997), Springer, pp. 309–320. 3
- [Pru10] PRUPPACHER K.: *Microphysics of Clouds and Precipitation*. Springer, Netherlands, 2010. 2
- [RGH*20] RATH A., GRITTMANN P., HERHOLZ S., VÉVODA P., SLUSALLEK P., KRIVÁNEK J.: Variance-aware path guiding. *ACM Trans. Graph.* 39, 4 (aug 2020). 3
- [SF79] SHETTLE E., FENN R.: Models for the aerosols of the lower atmosphere and the effects of humidity variations on their optical properties. *Environ. Res.* (09 1979), 94. 2
- [SRNN05] SUN B., RAMAMOORTHY R., NARASIMHAN S. G., NAYAR S. K.: A practical analytic single scattering model for real time rendering. *ACM Transactions on Graphics (TOG)* 24, 3 (2005), 1040–1049. 3
- [WVBR*21] WILKIE A., VEVODA P., BASHFORD-ROGERS T., HOŠEK L., ISER T., KOLÁŘOVÁ M., RITTIG T., KRIVÁNEK J.: A fitted radiance and attenuation model for realistic atmospheres. *ACM Transactions on Graphics* 40, 4 (Aug. 2021), 135:1–135:14. 2, 3
- [Zak94] ZAK J. A.: Drop size distributions and related properties of fog for five locations measured from aircraft. NASA Contractor Report 4585, DOT/FAA/CT-94/02, April 1994. Prepared for Langley Research Center under Contract NAS1-19341. 3, 4, 5
- [ZH*07] ZHOU K., HOU Q., GONG M., SNYDER J., GUO B., SHUM H.-Y.: Fogshop: Real-time design and rendering of inhomogeneous, single-scattering media. In *15th Pacific Conference on Computer Graphics and Applications (PG'07)* (2007), IEEE, pp. 116–125. 3