How Wildfires Spread and Why: Visual Multi-field Analysis of Vorticity-driven Lateral Spread Ensembles

Gabriel Borrelli* Lars Hagemann Jannik Steinkühler Adrian Derstroff Marina Evers
Karim Huesmann Simon Leistikow Hennes Rave Reyhaneh Sabbagh Gol Lars Linsen*

ABSTRACT

We present an interactive visual analysis tool for the spread of wildfires and what influences their evolution. Multiple spatio-temporal scalar and vector fields are investigated and related to each other to identify causes of atypical fire spread. Our tool allows for the comparative analysis of multiple runs of a simulation ensemble.

1 INTRODUCTION

Understanding the interactions of various physical and chemical phenomena is essential for the study of wildfires. Of those phenomena, *vorticity-driven lateral spread (VLS)* has been singled out as one important factor for causing atypical fire propagation [2]. Gaining insight into the emergence of VLS requires an analysis of potential relationships between the fire's environment, such as vegetation and atmospheric conditions, and the fire's evolution. We propose an interactive workflow for a visual analysis of those features and their relationships. We apply our methods to the IEEE SciVis 2022 Contest data to analyze interactions between the wind and other environmental features, especially in relation to fire spread.

2 METHODS

The given simulation ensemble consists of seven runs, where three can be categorized as *headcurve* (fire starts on the windward side of a mountain ridge), three as *backcurve* (fire starts on the leeward side), and one as a simulation of a *valley*. The headcurve and backcurve runs are differentiated by the curvature of the mountain ridge. Each run contains six scalar fields and one vector field. We present methods to analyze both the fields and their interactions with each other. Our work focuses on the analysis of flow features and especially the vorticity of the wind to understand the influence of the shape of the terrain.

2.1 Fire Spread

The spread of the fire over time can be analyzed and quantified by observing the behavior of the dry-fuel (rhof_1) field. For each run, we detect the time step in which the dry-fuel density inside a voxel starts to diminish and encode it using a sequential transfer function (see Fig. 1). Summing up the quantity by which the value of each voxel diminishes in one time step and multiplying it by the volume of the voxel ($1.8m^3$) allows us to further compare the mass of the burned dry fuel over time for each run (see Fig. 2). To compare the different runs' fire spread over the whole simulation time, we visualize the contours of the regions in which the dry fuel density changed between the last and the first time step (see Fig. 3).

2.2 Global Comparison of Simulation Runs

Similarity plots [1] are used to compare the evolution of all runs over time in all fields. They depict each simulation run by a curve parametrized over time in a low-dimensional embedding, where the closeness of curves in the embedding represents field similarity of the respective simulation runs' time steps. To choose the dimensionality of the embeddings, we plot the principal components (see Fig. 4), which indicates that the first principal component is always dominant despite not capturing the full variance. Still, we decided to use a 1D embedding, which allows us to intuitively plot the evolution over a time axis, see Fig. 5(a), Fig. 6(a), Fig. 7(a) and Fig. 8 for the evolution of the ensemble runs' similarity for different fields.

2.3 Correlations

To visually analyze the correlations and relations between given and derived scalar fields (including vector field magnitude), we use *parallel coordinates*.Interactively changing the order of axes as well as brushing on the axes to highlight selected data items using color allow us to investigate different relations and properties for a selected time step (see Fig. 9, Fig. 10). To observe how correlations change over time, we calculate the *Pearson correlations* among the whole fields at different timesteps (see Fig. 11). To analyze the impact of the vegetation, we create histograms to compare the fuel distribution of voxels that start burning to the overall distribution (see Fig. 12).

2.4 Volume Rendering

To analyze the spatial behavior of scalar fields, we use volume visualizations based on multi-field GPU ray-casting. To define a meaningful color mapping for each scalar field, we use piecewise linear transfer functions [3]. Dry-fuel, temperature, and O2 scalar fields, e.g., allow for a visual comparison of a chosen time step for different runs (see Fig. 13). We also use volume renderings of the fire to provide context in other visualizations (see Fig. 14).

2.5 Flow Integration

Line- and surface-based flow visualizations show the wind's behavior, especially in the vicinity of the fire. Streamlines and *stream surfaces* (see Fig. 14) can be used to analyze the current state of the flow field , while pathlines and *path surfaces* provide additional temporal information that can be color-coded into the surfaces (see Fig. 15). For flow integration, we use a fourth-order Runge-Kutta scheme. The seeding curves can be placed on arbitrary, user-defined positions inside the given volume. Alternatively, a GPU implementation of the Marching Cubes algorithm for isosurface extractions of the rhof_1 field approximates a region of interest near the fire where the seeding curve can be placed automatically based on a few parameters, like the distance from the region of interest (see Fig. 16).

2.6 Vorticity and Divergence

Vorticity is a measure of the curl or rotation of the wind vector field that is of particular interest in the context of vorticity-driven lateral spread (VLS). Vorticity of a flow volume is given by $\omega = \nabla \times V$ where $V = (u, v, w)$ denotes a 3D vector field and ∇ the gradient operator. Due to the fact, that the vorticity does not describe the direct movement of particles anymore, we cannot utilize stream- or pathlines. Instead, we follow the idea of Sharples et al. [2] to differentiate between ambient and pyrogenic vorticity. The ambient vorticity is defined as the v-component of the vorticity ω_v, while the pyrogenic vorticity ω_p is defined as the vector $\omega_p = (\omega_u, 0, \omega_w)$. Sharples et al. describe regions of interest on the right and left flanks

*e-mail: {gabriel.borrelli | linsen}@uni-muenster.de
All authors are affiliated with the University of Münster.

of the fire, where both components of the pyrogenic vorticity are positive or negative, respectively. The interaction between ambient and pyrogenic vorticity drive the occurrence of VLS. We apply a volume renderer to the ambient vorticity ω_v (see Fig. 17). Additionally, we categorize the PV into the sets $\omega_{p<0} = (\omega_u < 0, 0, \omega_w < 0)$ and $\omega_{p>0} = (\omega_u > 0, 0, \omega_w > 0)$ and map the result to -1 if $\omega_p \in \omega_{p<0}$, 1 if $\omega_p \in \omega_{p>0}$, and 0 otherwise. Then, we apply again a volume renderer (see Fig. 18).

Divergence is a measure of a flow to compress or spread out at a certain point. Divergence $\mathrm{div}(V)$ of a vector field V is defined as $\mathrm{div}(V) = \nabla \cdot V$. We visualize it using volume rendering (see Fig. 19).

2.7 Simulation Run Prediction

Since the analysis of the influence of simulation parameters would benefit from more simulation runs, we aimed to extend the simulation ensemble. As a proof of concept, we apply a conditional variational auto-encoder (CVAE) [4] that is trained on the u, v, and w component of the vector field of time steps 6 - 70, after the fire was ignited on a resampled grid of size $64x64x48$. The respective results for our test run (backcurve 80) can be seen in Fig. 20. Our CVAE is able to predict a simulation run for other ridge smoothness parameters, which we demonstrate for backcurve 200 (cf. Fig. 8).

2.8 Analytical Workflow

We propose a top-down analysis, where we start with an overview of the fire's spread and the correlations among the fields using the methods presented in Sections 2.1-2.3 , giving us an idea of the temporal evolution of the fire, the occurrence of VLS and the relations between the different fields. Based on these plots, we can proceed with an interactive analysis of the data. More detailed visualizations of the wind, its vorticity, and volume renderings of scalar fields (as described in Sections 2.4-2.6) are shown in multiple coordinated views that can be used to investigate the conditions that lead to the observed fire spread or compare different ensemble members.

3 Implementation

Our application extends the open source rapid application development framework Voreen [3]. The overview approaches for analyzing the spread of the fire were implemented using Python. The given ensemble was converted from VTS to the HDF5 format by resampling the data into a regular grid of resolution $600 \times 500 \times 60$ (cf. Fig. 21), leading to a reduction in the required disk storage and an accelerated loading of the dataset. To achieve voxel-wise comparability between all simulation runs, we create a second version of the ensemble dataset, parameterized by the height from the ground, which we use to compute the field similarity for the similarity plots.

4 Results

Fire Spread. We begin the analysis by looking at the temporal evolution of the rhof_1 field at ground level to get an overview of the total spread of the fire and when the fire starts to spread (Figs. 3, 1, 2). Additionally, this gives a good intuition for the beginning of the lateral spread on the mountain ridge. We notice that the lateral spread is larger for steeper mountains and does not occur for a roundness of 320. It begins between timesteps 30 and 35 in all the simulations and only occurs when the fire reached the lee side of the mountain ridge. Finally, we compare the streamwise speed of the fire and observe no difference in the headcurve runs, whilst the backcurve speeds are indirectly influenced by the radius of the mountain curve. The latter is due to the fire's movement driven mainly by the lee side upwind.

Comparison of Simulation Runs. The similarity plots reveal two groups (headcurve vs. backcurve) of runs with similar evolution for the rhof_1 (Fig. 5(a)), rhowatervapor (Fig. 6(a)) and theta (Fig. 7(a)) field, while three groups (w.r.t. ridge smoothness) can be observed for the wind speed magnitude field (Fig. 8).

Analysis of Environment. For analyzing the environment, we present results for the run backcurve_40 but similar phenomena can be seen for other runs. The temperature and water vapor field are highly correlated while oxygen and temperature show a strong anticorrelation after an initial transition phase, see Fig. 11. To analyze the direct influence in the burning region, we define a mask covering regions of burned fuel and a margin of one voxel surrounding this region. The parallel coordinates for this region are shown in Figs. 9 and 10. We see that strong wind prevents the fuel from burning. The distribution of fuel shown in Fig. 12 shows that voxels that are burned tend to have a higher density.

Flow Analysis. Expanding on the previous analysis, streamlines allow us to analyze the upwind on the lee side of the mountain (see Figs. 22, 23) by visualizing the flow of particles. Fig. 23 shows how the winds stemming from the windward streamwise flow evolve into updraft winds in the opposite direction on the lee side of the mountain. These winds eventually push the fire towards the mountain ridge. We observe similar vorticity and updraft in the valley run, thus making this simulation behave more like a backcurve in this regard. Additionally, path surfaces help in gaining an intuition for the ambient and pyrogenic atmospheric turbulence (Fig. 24). Turbulence independent of variations caused by the fire is categorized as ambient turbulence (i.e. caused by the mountain, friction with the ground level terrain). In contrast to that is the pyrogenic turbulence originating from the fire's interaction with the wind by heat transfer or pressure changes (Fig. 19).

Vorticity. To analyze the cause and effects of the observed turbulence, especially regarding VLS, we visualize the vorticity field of the wind flow. A volume rendering of the ambient vorticity (AV) ω_v (see Fig. 17) shows a vorticity sheet oriented mostly in the $+Y$ direction on the lee side of the mountain. The vorticity sheet's height and density above the mountain ridge increases with decreasing mountain curvature. In the volume rendering of the categorized pyrogenic vorticity (PV) (see Fig. 18), we observe two distinct regions of high vorticity on the two opposing flanks of the fire. The interaction of this PV with the AV sheet on the mountain's lee side creates the possibility for lateral spread near the ridge [2].

5 Conclusion

We proposed an interactive tool for the analysis and visualization of various phenomena emerging from wildfires. Overview plots showing, e.g., the expansion of fire over time allow immediate comparison between simulation runs. Stream and pathlines (as well as their surface variants) can be used to visualize wind flow in a single or multiple timesteps. Combined with visualizations of vorticity and divergence, these tools provide an intuitive understanding of the VLS phenomenon and the influence that the fire and surrounding terrain have on it. This understanding might help firefighters to better estimate the dangers and future spread of the fire.

References

[1] A. Fofonov and L. Linsen. Projected field similarity for comparative visualization of multi-run multi-field time-varying spatial data. In *Computer Graphics Forum*, vol. 38, pp. 286–299. Wiley Online Library, 2019.

[2] J. Sharples, A. Kiss, J. Raposo, D. Viegas, and C. Simpson. Pyrogenic vorticity from windward and lee slope fires. *Int. Congr. Model. Simul., Gold Coast, Aust. 29 Nov.–4 Dec*, pp. 291–97, 2015.

[3] Voreen. voreen.uni-muenster.de.

[4] X. Yan, J. Yang, K. Sohn, and H. Lee. Attribute2image: Conditional image generation from visual attributes. In *European Conference on Computer Vision*, pp. 776–791. Springer, 2016.

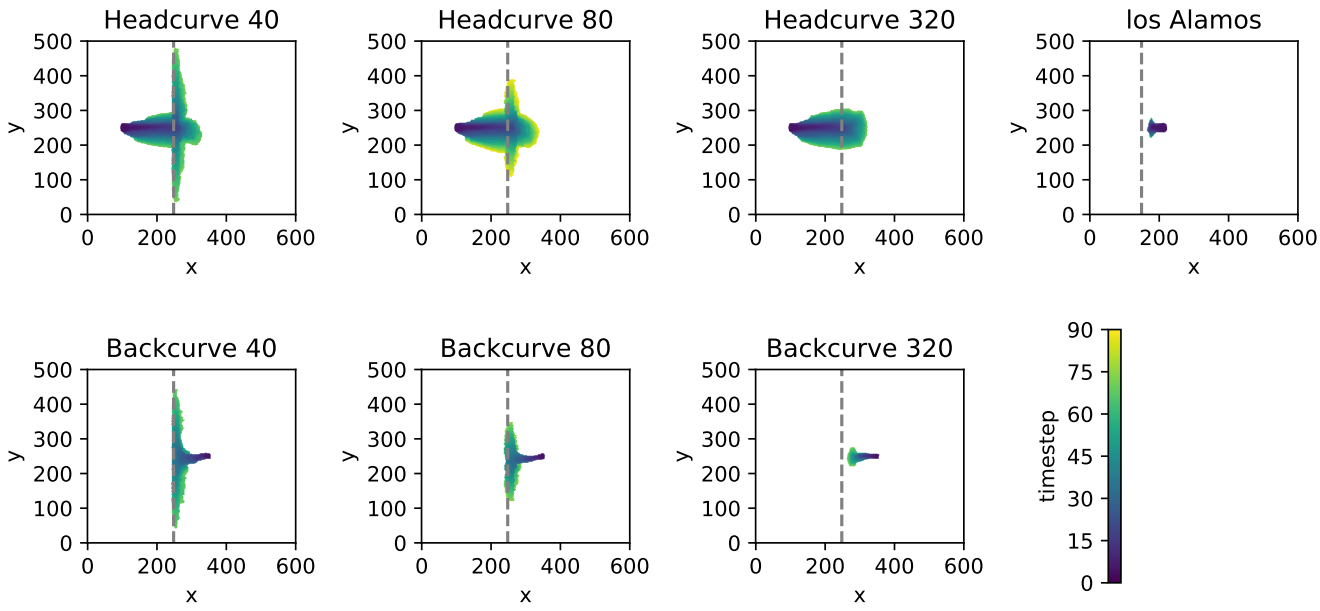

Figure 1: Fire expansion over time for different runs looking from the top. The color encodes the timestep in which a change in the rhof_1 field occurs. We see differences in the expansion speed of the fire. For the valley (los Alamos), we see that the fire not only spreads very little but the spread slows down in the progression of the simulation until it stops completely.

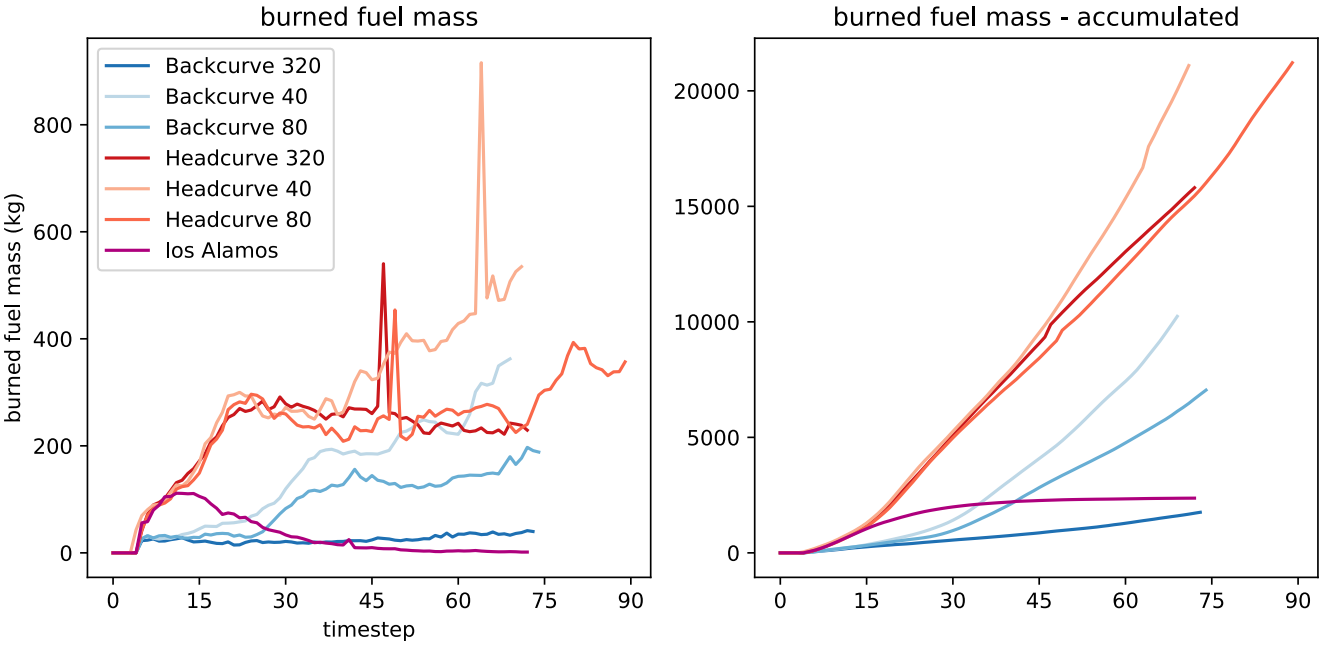

Figure 2: Multi-run plot of the burned dry fuel rhof_1 field per timestep (left) and accumulated (right). These plots confirm our observations in Fig. 1 regarding the different speeds of fire expansion. It is also noticible that the valley run (los Alamos) is the only simulation run for which the burned fuel mass decreases.

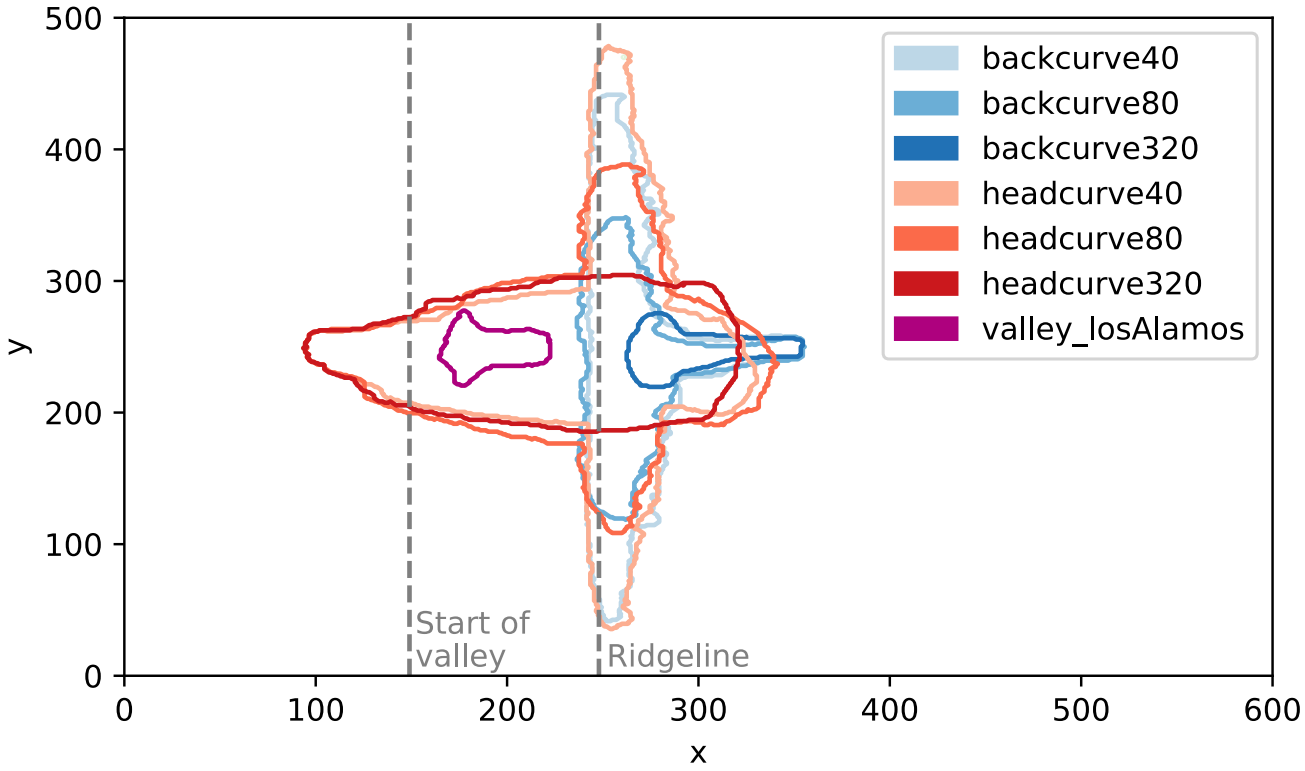

Figure 3: Overview of the fire expansion for different runs looking from the top. The different colors represent the different ensemble members. We observe a stronger lateral spread for higher curvatures of the mountain ridge. For a curvature of 320 we see no lateral spread. In case of the valley, we see a relatively small lateral spread whose shape looks similar to the spread of backcurve_40.

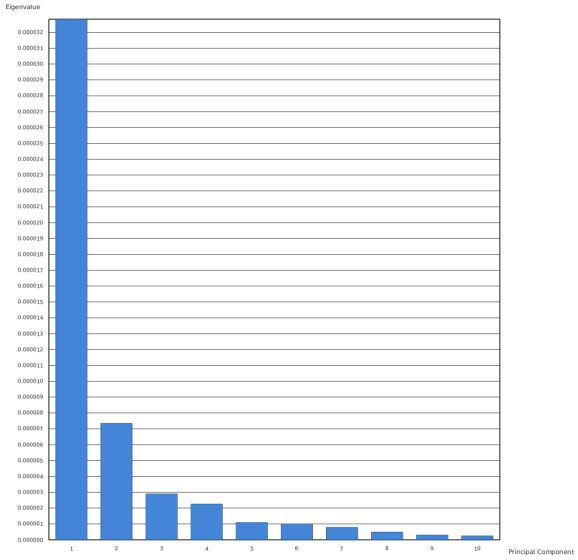

(a) Bulk density of dry fuel

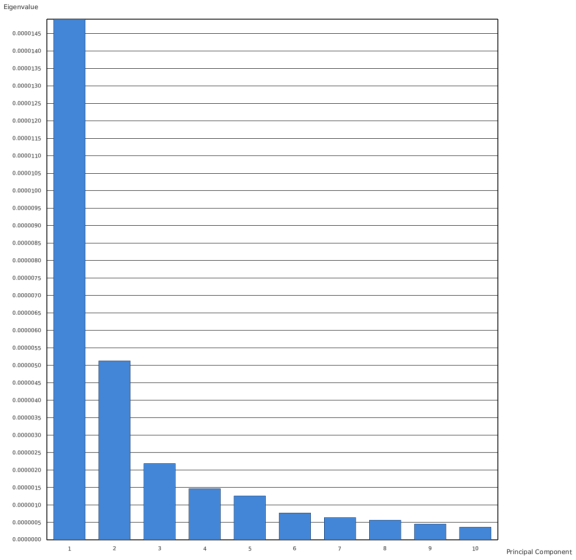

(b) Bulk density of moisture

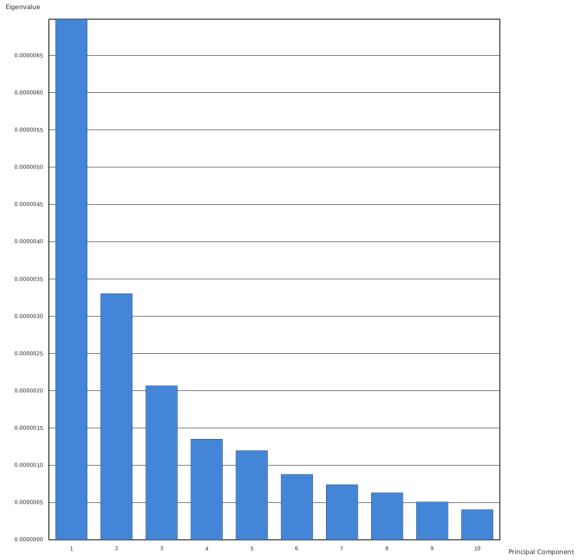

(c) Potential temperature

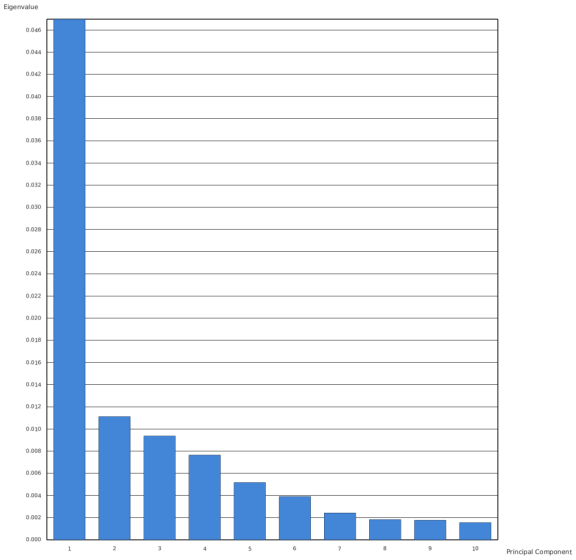

(d) Wind speed magnitude

Figure 4: Eigenvalue plots for the embeddings used for the similarity plots shown in Figs. 5, 6, 7 and 8, respectively, indicate, that the first principal component might not capture most of the variance. However, it is dominant in all cases for the selected fields. We hence make use of 1D similarity plots, where time is used as second axis.

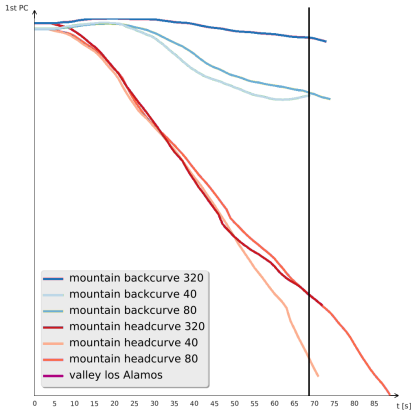

(a) Similarity plot for bulk density of dry fuel. Time step 69 has been selected. The valley run has been removed to better analyze the two groups.

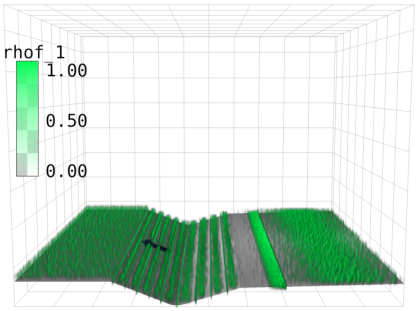

(b) valley_losAlamos

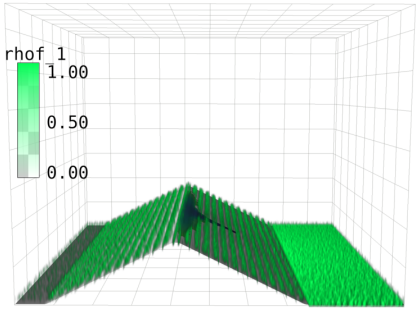

(c) mountain_backcurve40

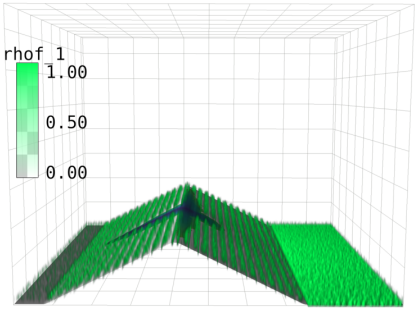

(d) mountain_headcurve40

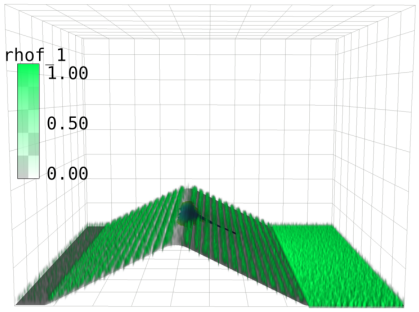

(e) mountain_backcurve80

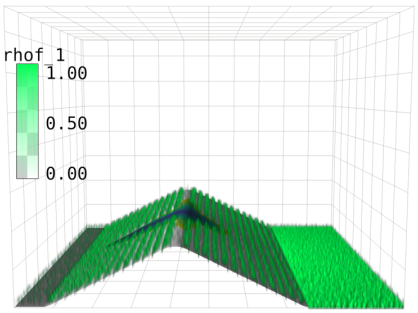

(f) mountain_headcurve80

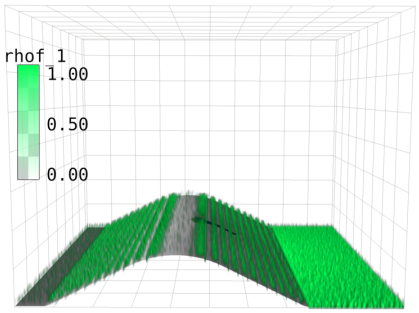

(g) mountain_backcurve320

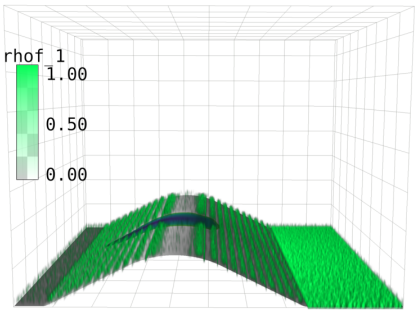

(h) mountain_headcurve320

Figure 5: In the similarity plot of all simulation runs for the bulk density of dry fuel field, we observe two groups. The group using red colors is represented by the headcurve simulation runs (d), (f), and (h). The group using blue colors is represented by the backcurve simulation runs (c), (e), and (g). In the respective volume renderings, the extent of the fire has been projected on the surface geometry to highlight the affected area. After the fire has been ignited (t=5), the groups separate over time. The difference can hence be explained by the amount of voxels burned, that is higher in the headcurve runs (cf. Fig. 2).

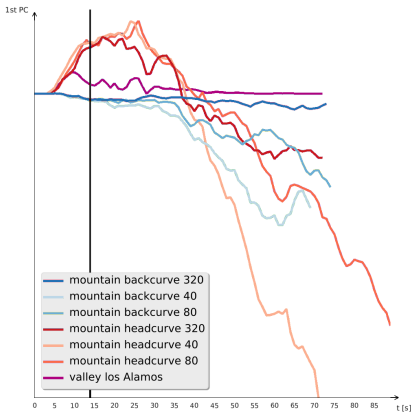

(a) Similarity plot for bulk density of moisture. Time step 14 has been selected.

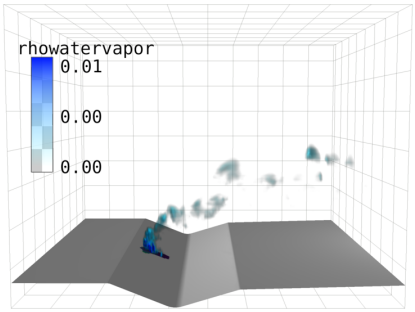

(b) valley_losAlamos

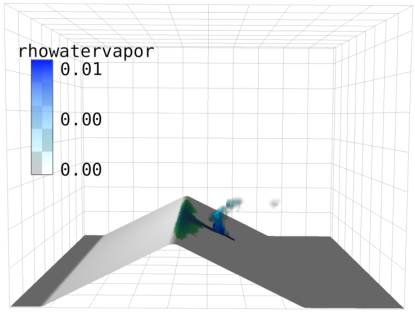

(c) mountain_backcurve40

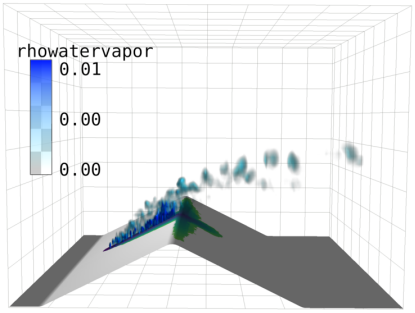

(d) mountain_headcurve40

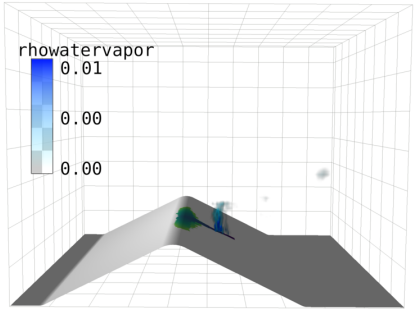

(e) mountain_backcurve80

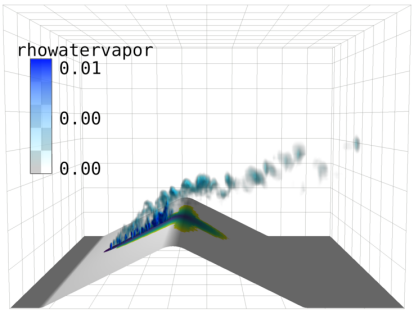

(f) mountain_headcurve80

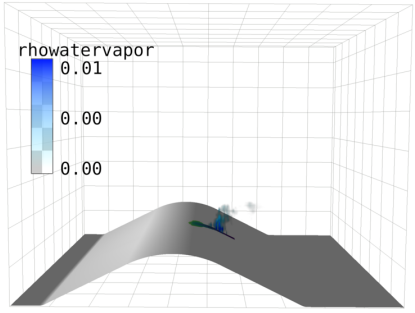

(g) mountain_backcurve320

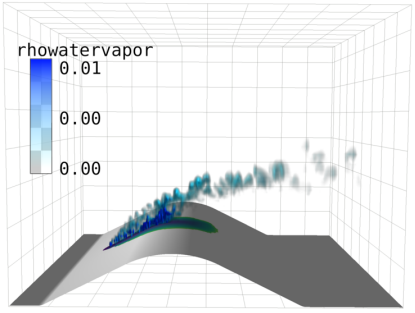

(h) mountain_headcurve320

Figure 6: In the similarity plot of all simulation runs for the bulk density of moisture field, we observe two groups up to a certain point in time. The group using red colors is represented by the headcurve simulation runs (d), (f), and (h). The group using blue colors is represented by the backcurve simulation runs (c), (e), and (g). The valley run (purple, (b)) evolves differently from the two groups, but more similar to the backcurve runs. In the respective volume renderings, the extent of the fire has been projected on the surface geometry to highlight the extent of the fire for the selected time frame. After the fire has been ignited (t=5), the runs separate over time until the groups vanish ($t \approx 40$). For the selected time step, we observe in the volume renderings that the fire in the headcurve runs yields much more water vapor to the athmosphere than the backcurve runs, causing the separation in the similarity plot. The valley run, for both volume rendering and similarity plot, is situated in between the two groups.

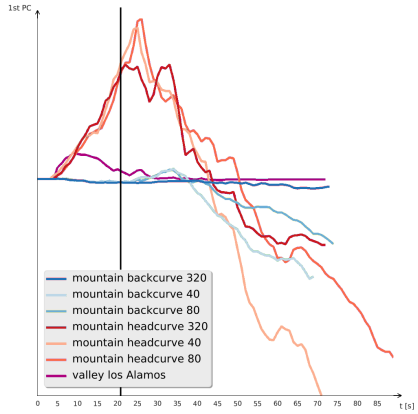

(a) Similarity plot for potential temperature. Time step 21 has been selected.

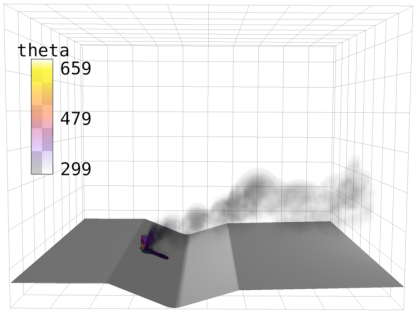

(b) valley_losAlamos

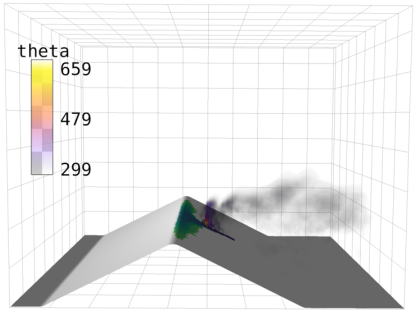

(c) mountain_backcurve40

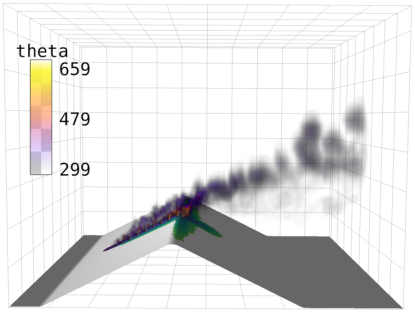

(d) mountain_headcurve40

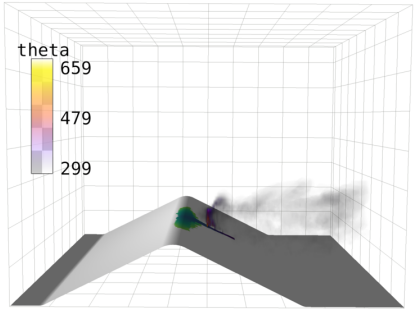

(e) mountain_backcurve80

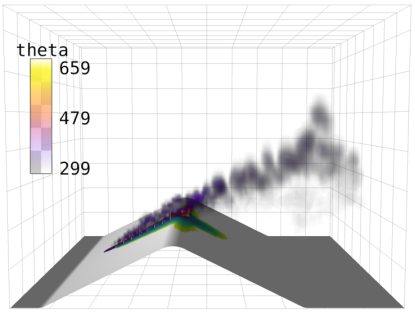

(f) mountain_headcurve80

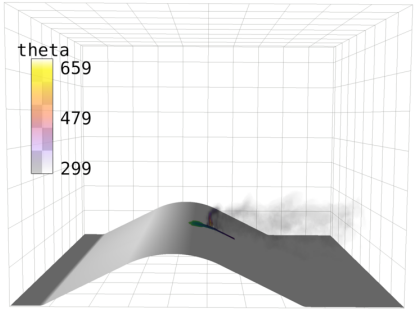

(g) mountain_backcurve320

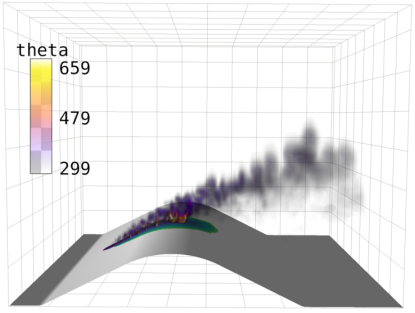

(h) mountain_headcurve320

Figure 7: In the similarity plot of all simulation runs for the potential temperature field, we observe two groups up to a certain point in time. The group using red colors is represented by the headcurve simulation runs (d), (f), and (h). The group using blue colors is represented by the backcurve simulation runs (c), (e), and (g). The valley run (purple, (b)) evolves differently from the two groups, but more similar to the backcurve runs, eventually resembling the backcurve 320 run. In the respective volume renderings, the extent of the fire has been projected on the surface geometry to highlight the extent of the fire for the selected time frame. After the fire has been ignited (t=5), the runs separate over time until the groups vanish ($t \approx 40$). For the selected time step, we observe in the volume renderings that the overall potential temperature is much higher than in the backcurve runs, causing the separation in the similarity plot. The valley run, for both volume rendering and similarity plot, resembles more the backcurve runs, i.e., the purple curve is close to the backcurve group.

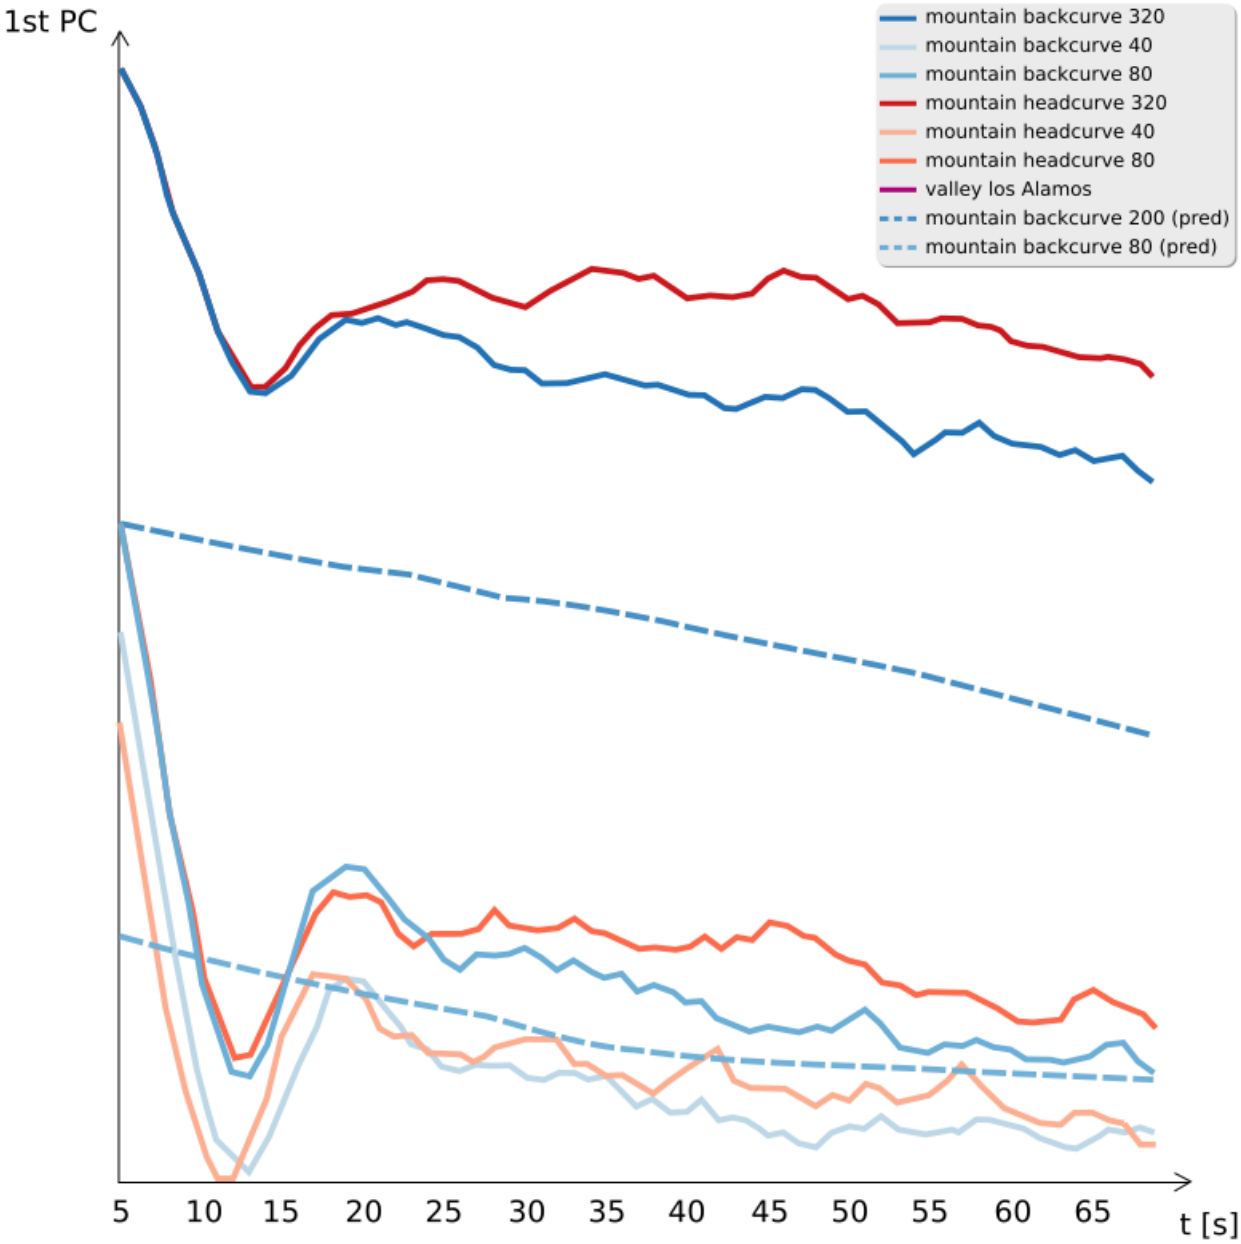

Figure 8: Similarity plot of all simulation runs for the magnitude of the uvw-vector field. We observe three groups that are characterized by the roundness of the mountain ridge where, intuitively, the more similar values (40 and 80) correspond to more similar vector fields. The valley run was projected very differently and was hence removed from the embedding to achieve better comparison of the other runs. Two additional runs have been added (dashed lines) that have been generated using a prediction with a conditional variational auto-encoder neural network. For the training process, only the time range from $t = 5$ to $t = 70$ has been considered, i.e., time after the fire has been ignited and before the first run ends. The predicted run for backcurve 80 is in the proximity of the ground truth data (of course, the ground truth was excluded from the training data). Additionally, a backcurve 200 run was predicted using our neural network, leading to a vector field in between the backcurve 80 and backcurve 320 run. For both predicted runs, the high frequencies could not be predicted, however, the overall trend could be captured. With our architecture it is, in principle, possible to predict a simulation run for any value of the smoothness parameter of the mountain ridge. Of course, a larger training data set would be helpful.

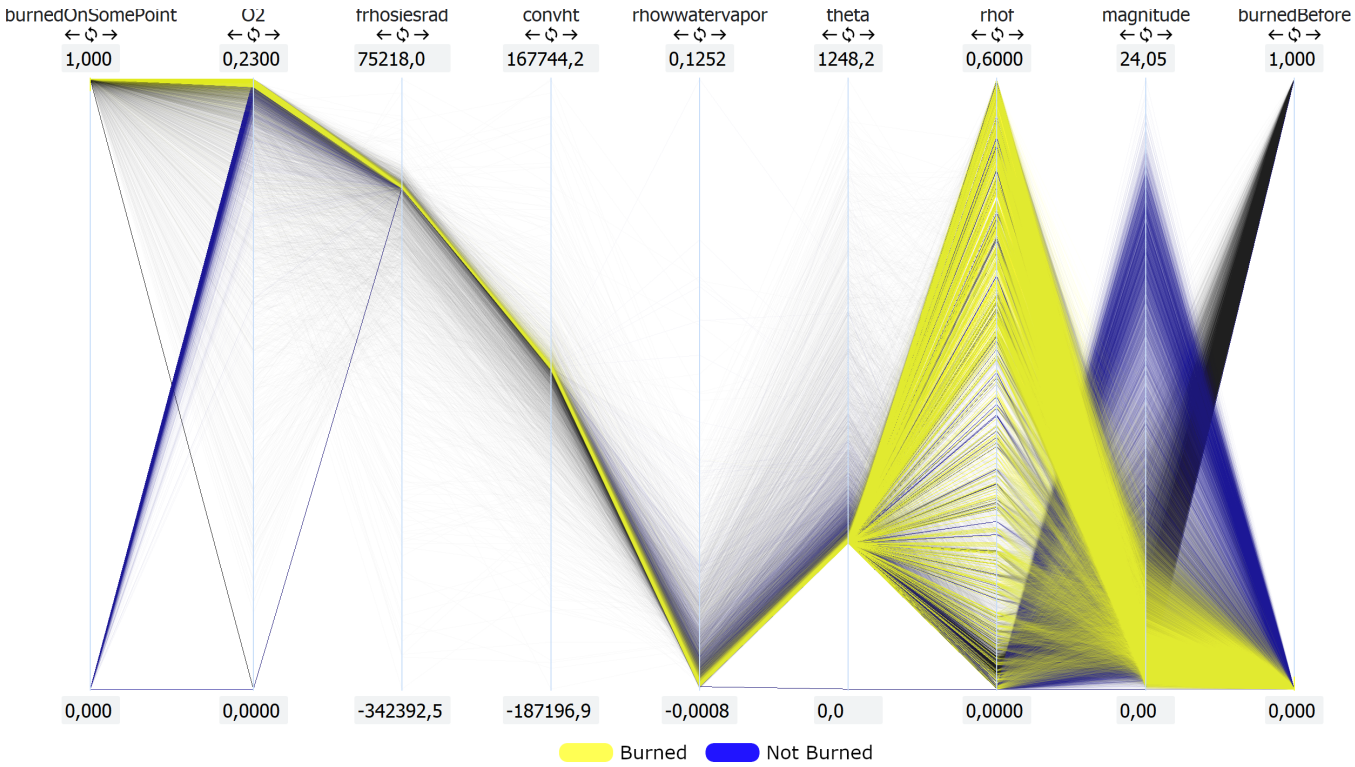

Figure 9: Parallel coordinates for the run backcurve_40 where those samples are selected where the dry fuel did not decrease compared to the previous timestep (no fuel burned). We only consider those voxels, where fuel gets burned on some point in time together with a margin of one voxel in each direction. Here, we see an anticorrelation between the density of dry fuel and the wind magnitude which can be explained by the vegetation stopping the wind. Voxels with strong wind (high values on the magnitude axis) will not burn in the future. We also observe several voxels with high temperature (axis theta) and high moistures (axis rhowatervapor) that will not burn in the future. For further investigations on these voxels, see Fig. 10.

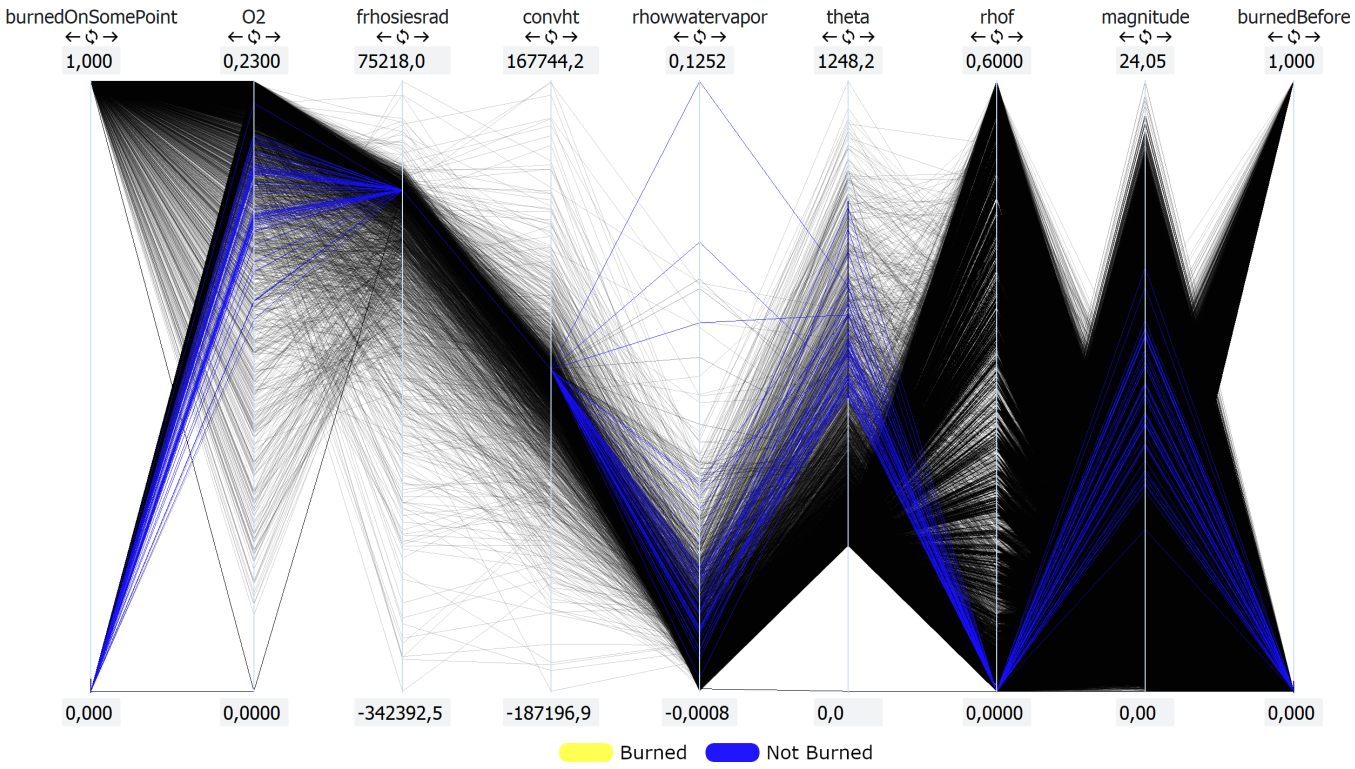

Figure 10: When selecting the voxels with high temperature (between $600\,\mathrm{K}$ and $1000\,\mathrm{K}$) in Fig. 9, we see that those, that do not burn also contain very low densities of dry fuels. This can be explained by considering that we also sample a margin of one around the regions where fuel gets burned. Thus, these voxels correspond to the voxels on the upper part of the margin where very little to no fuel is present.

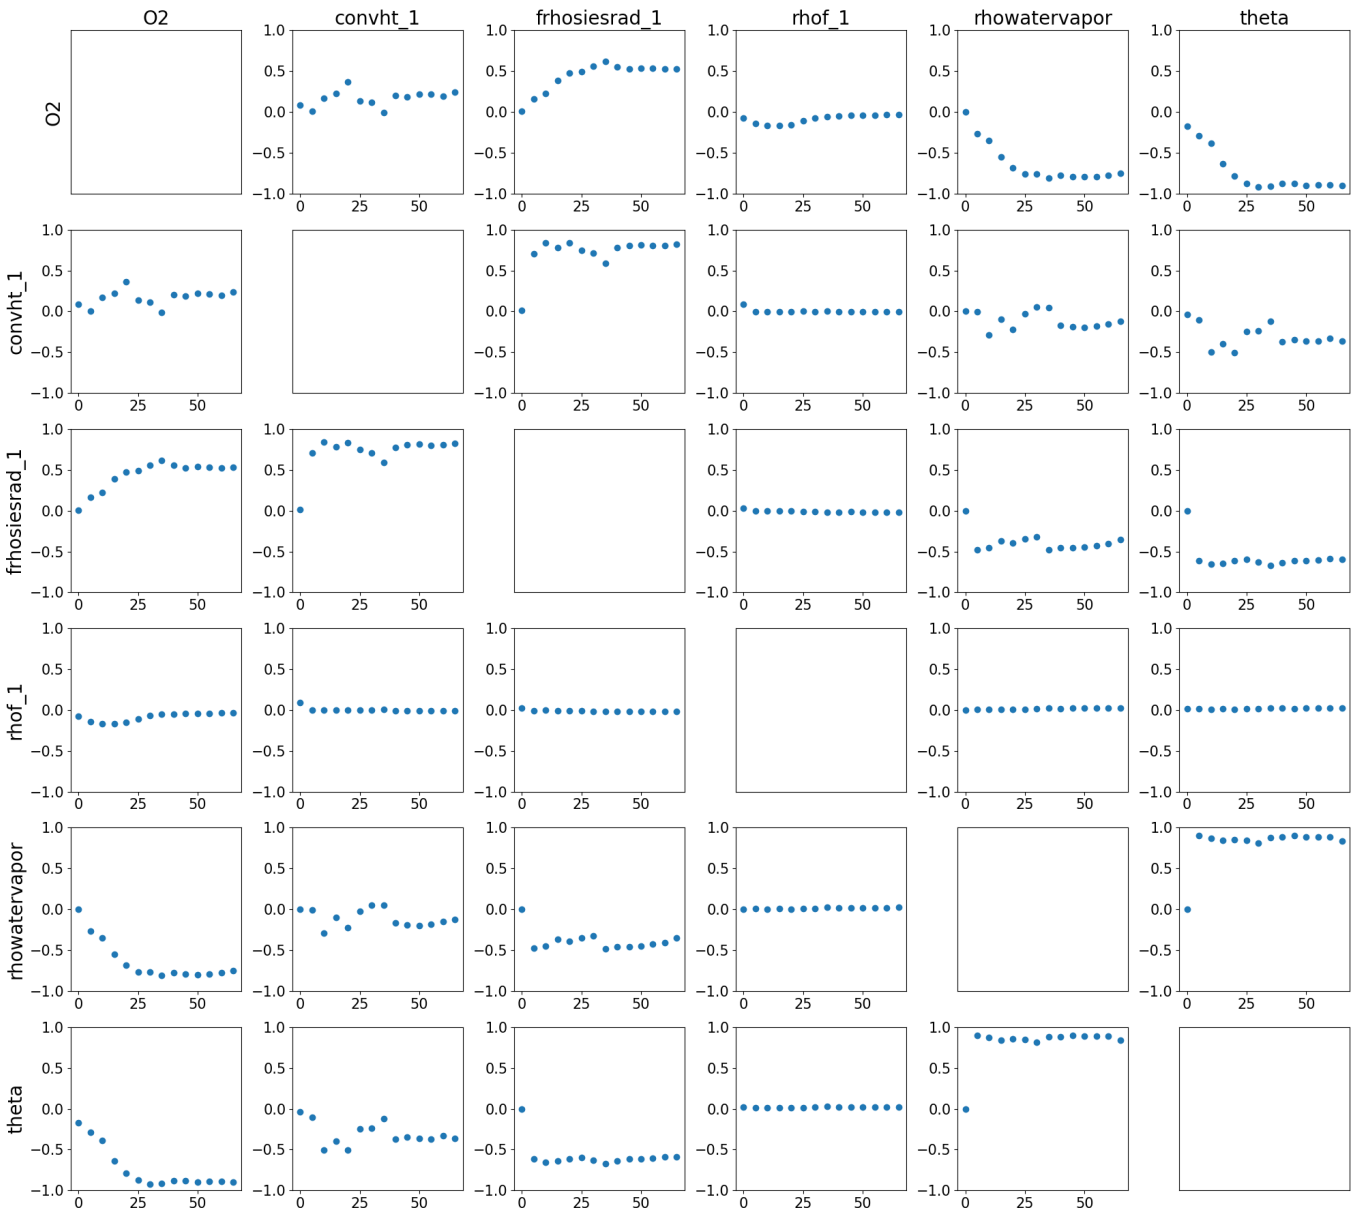

Figure 11: Correlations between all scalar fields of simulation run backcurve 40, computed every five timesteps. As expected, the density of released moisture (rhowatervapor) and temperature (theta) as well as the convective (convht_1) and radiative heat transfer (frhosiesrad_1) show a strong positive correlation, while oxygen (O2) and temperature (and consequently oxygen and the density of released moisture) show a strong negative correlation. The density of dry fuel (rhof_1) shows no correlations with other fields.

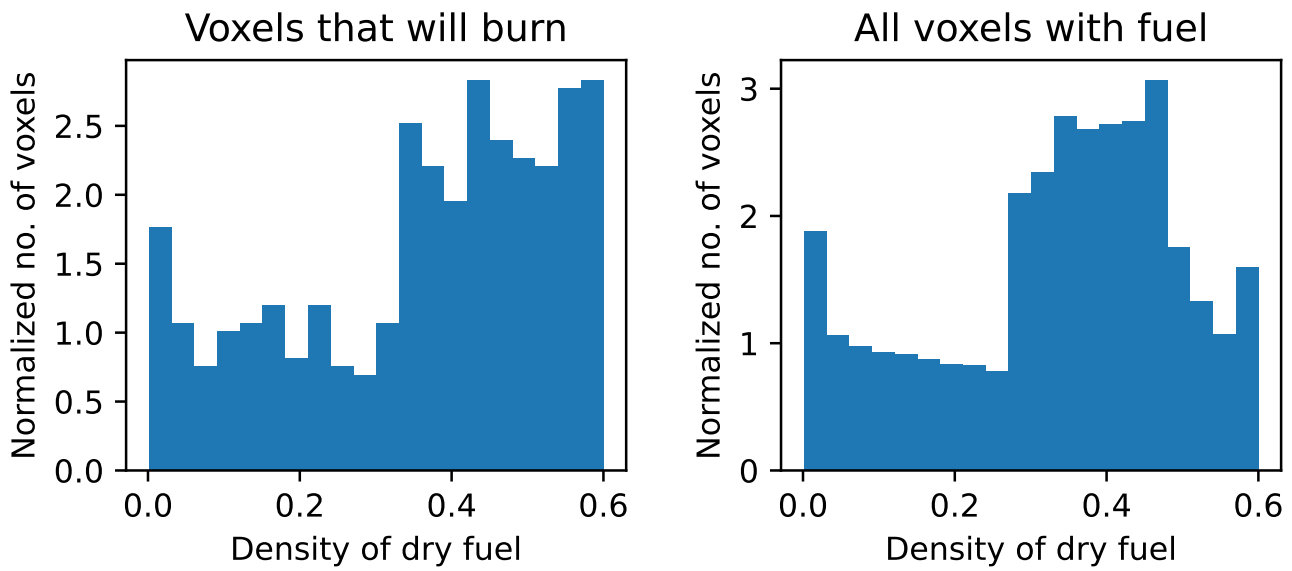

Figure 12: Histograms of the initial fuel distribution for the run backcurve_40 of those voxels that will burn on the left and of all voxels on the right. For both histograms, we set a lower threshhold of $0.001\,\mathrm{kg\,m^{-3}}$ for the fuel density to exclude the voxels that do not contain a significant density of fuel. To make both histograms comparable, we normalize the total density (area below the plots) to 1. Even though the distributions are similar, there are more voxels with high density among those that will start burning compared to all voxels with fuel.

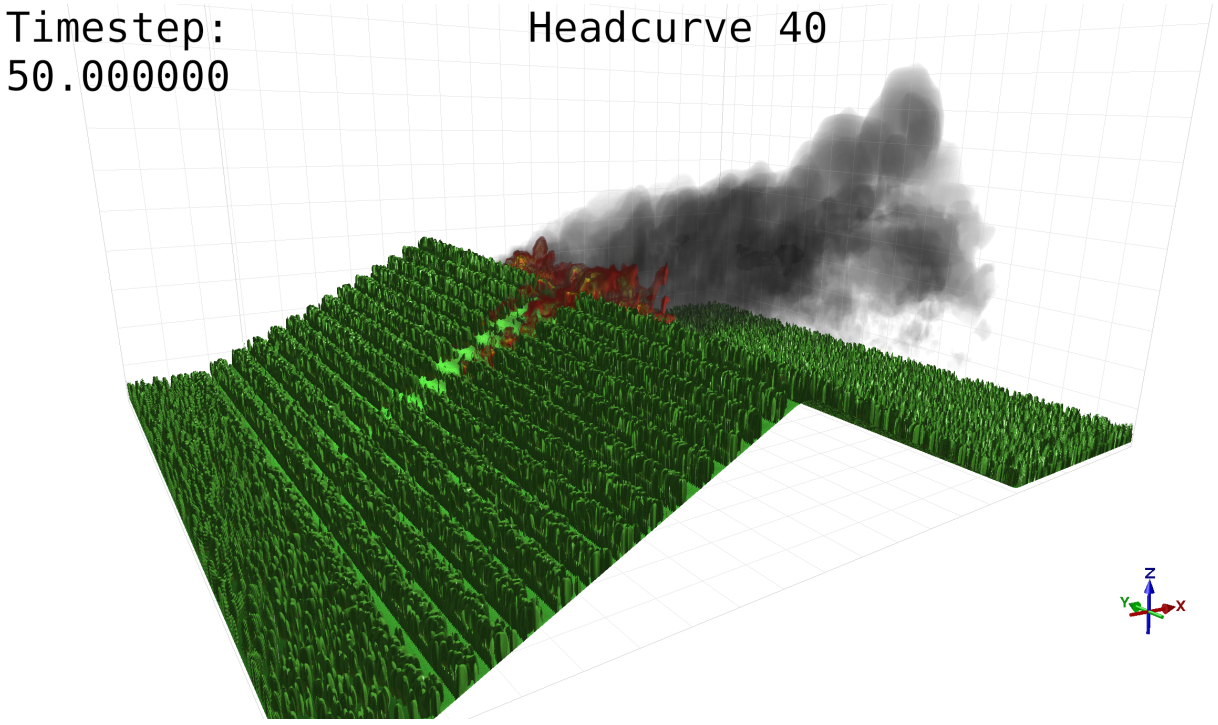

Figure 13: Multi-field volume rendering of the dry-fuel (rhof_1), the temperature (theta) and O2 scalar fields of the headcurve_40 run. The visualization shows the propagation of the fire and the consumption of the dry-fuel at timestep 50. The dry-fuel is colored in green for voxels with values larger than 0.03 $kg\ m^{-3}$. The theta field is colored yellow or red for temperatures above 400K and grey for temperatures between 300K and 330K. The O2 field is colored yellow as well for values greater than 0.21 and lower than 0.23.

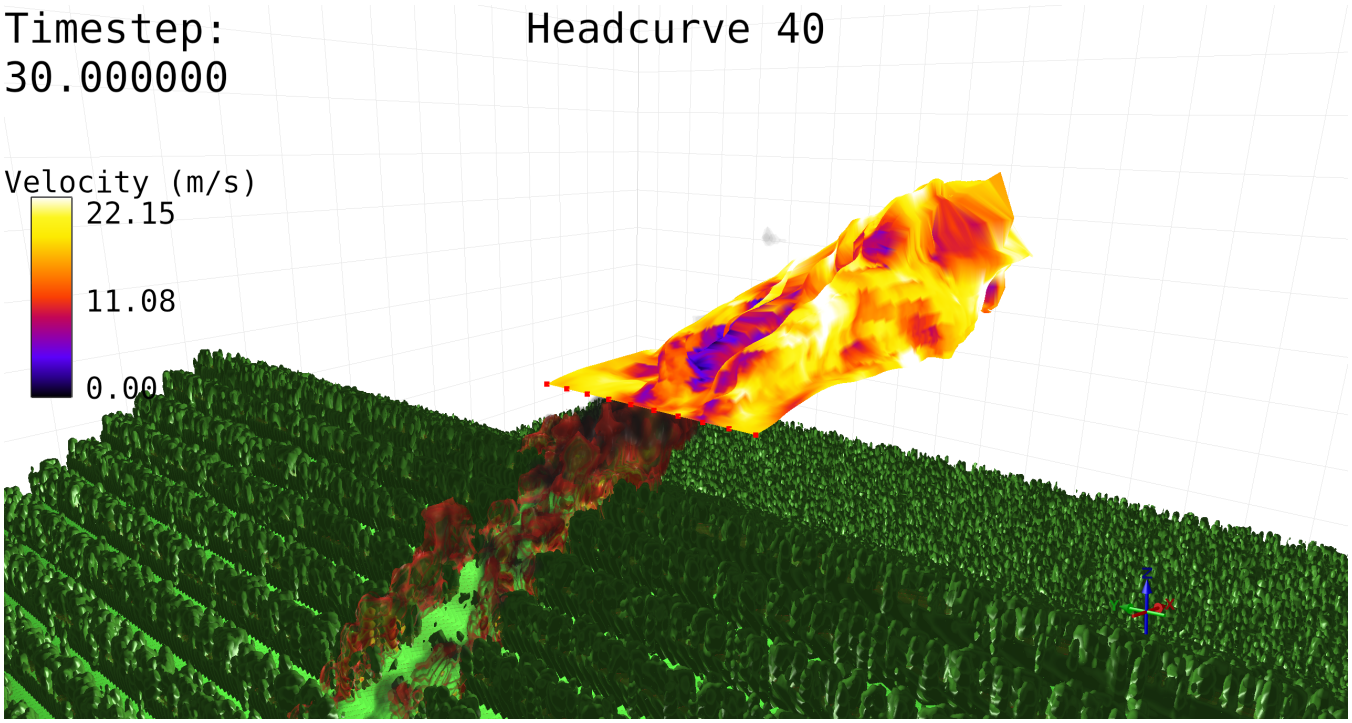

Figure 14: Stream surface visualization of the wind vector field (u, v, w) at timestep 30 of the run headcurve_40. The seeding curve (red dots) is set behind the fire with about twice the length of the fire width. The stream surfaces are color-coded by the velocity of the wind field.

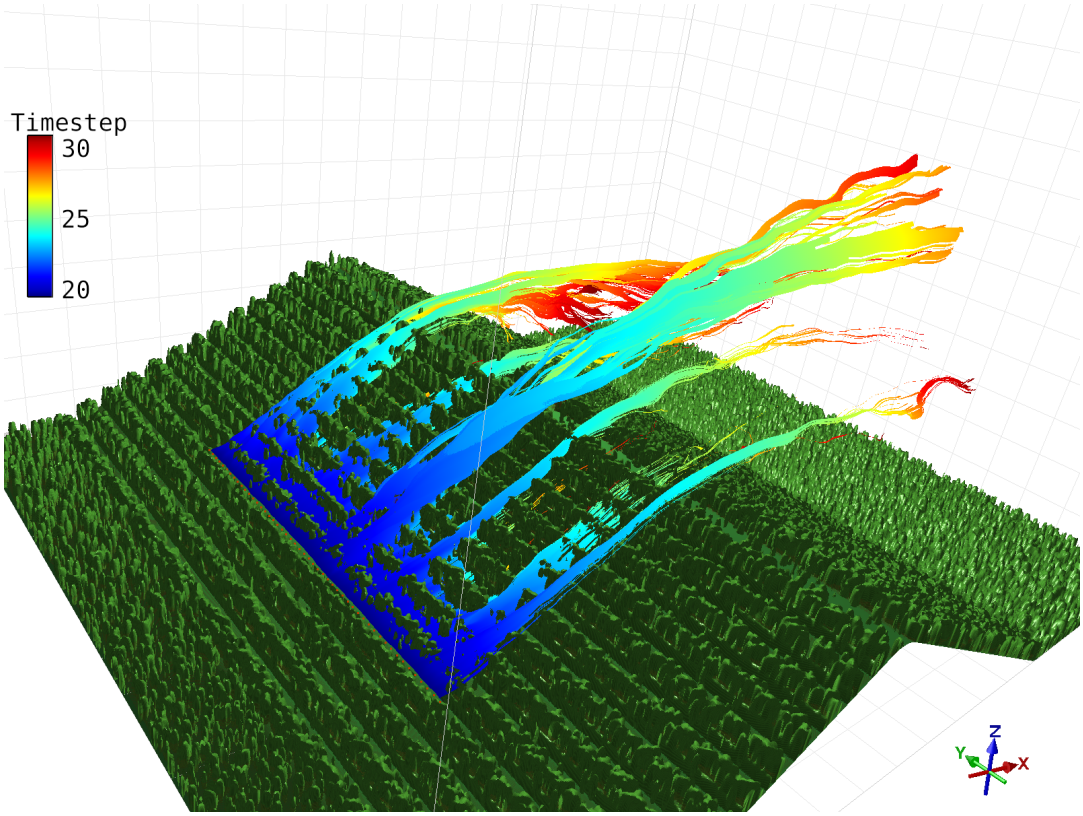

(a) Opaque solid path surface

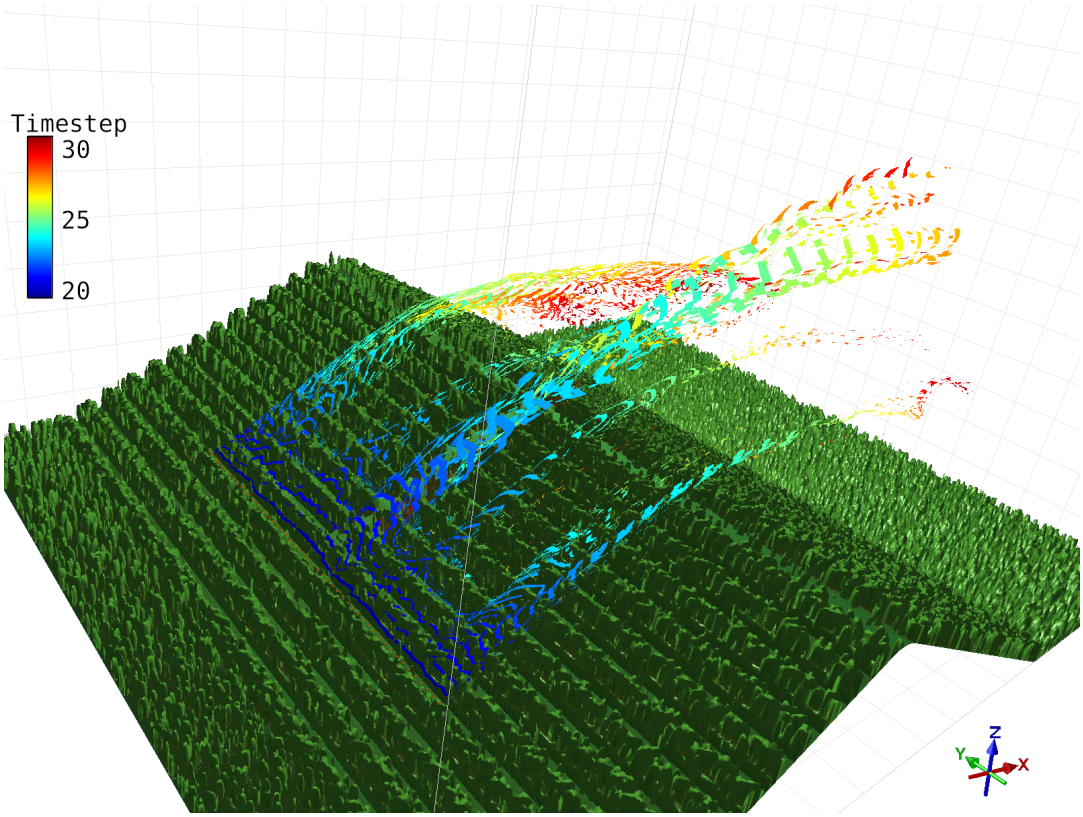

(b) Timeline visualization of the same path surface for better background visibility

Figure 15: (a) Path surface of the run headcurve_40 seeded at the beginning of the windward side of the mountain and integrated from timestep 20 to timestep 30. (b) The same path surface with a texture that highlights the evolution of timelines.

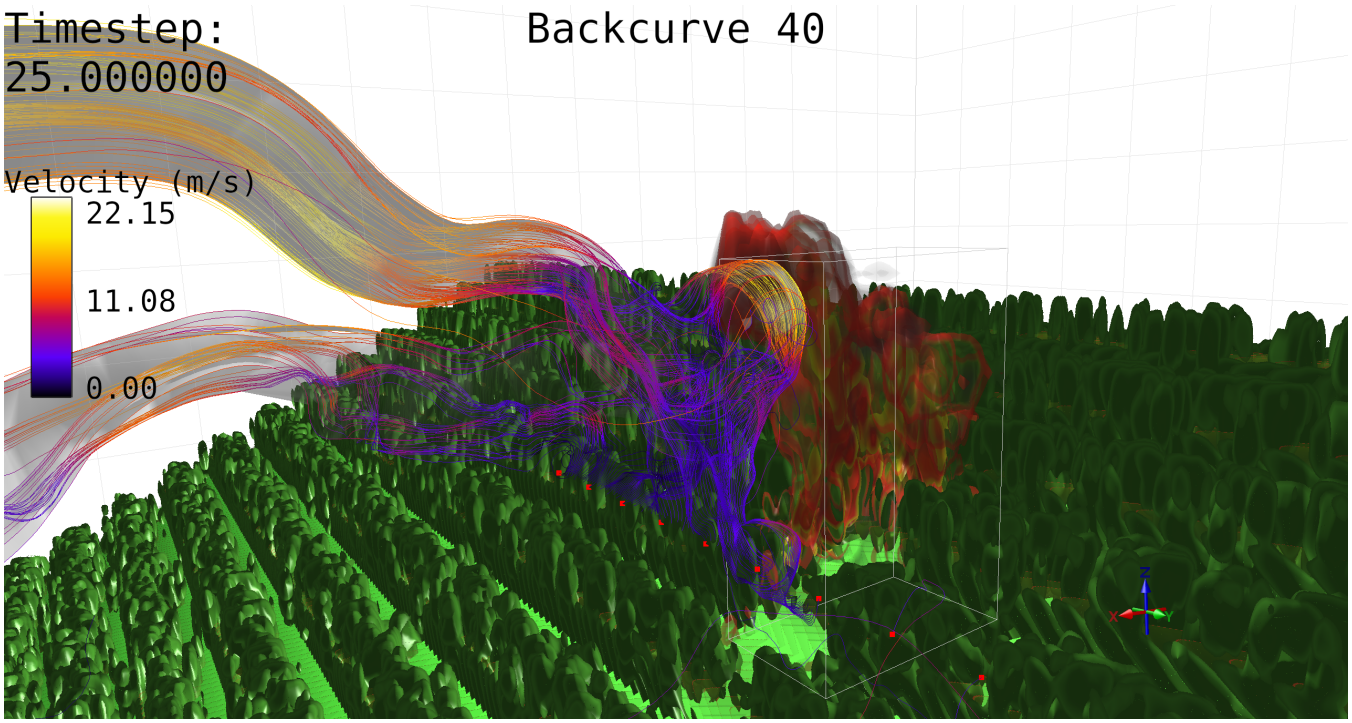

Figure 16: Stream surface visualization of the wind vector field (u, v, w) at timestep 25 of the run backcurve_40 with embedded streamlines. The seeding curve (red dots) is automatically placed a set distance in front of the region of interest given by the fire. The region of interest is visualized with a bounding box. The streamlines are color-coded by the velocity of the wind field. We observe slow winds on the lee side of the mountain raising in altitude while accelerating before moving windward at higher velocities. Furthermore, we observe that the fire disturbs the wind directly in front of it, causing vortices and also achieving higher vertical velocities.

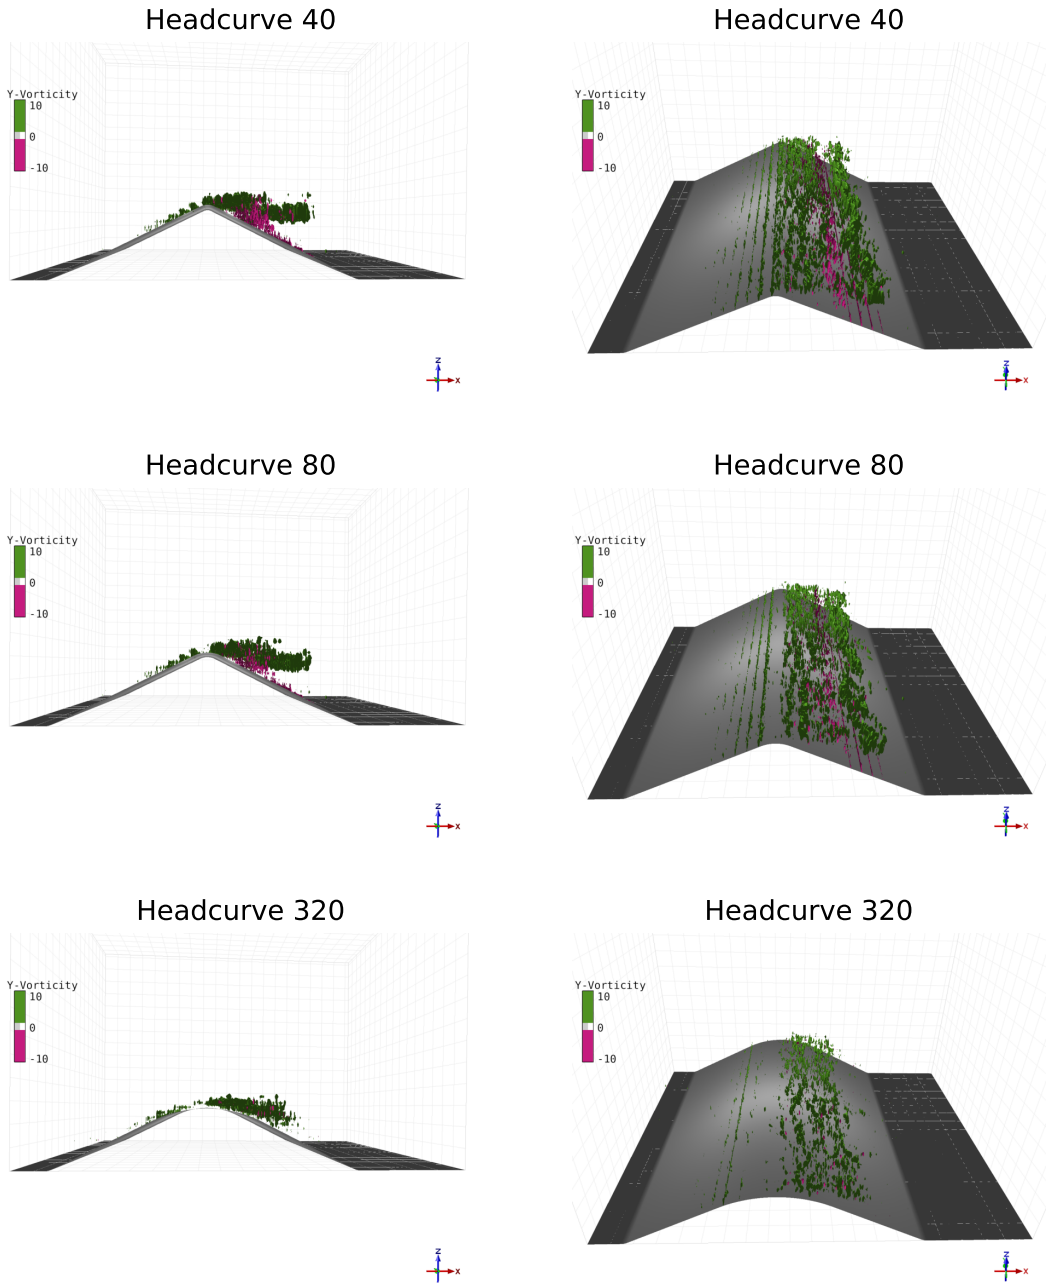

Figure 17: The y-component of the vorticity field for timestep 4 (thus, before the ignition) in the headcurve runs. An increase in y-vorticity (or ambient vorticity) and height of the vorticity sheet on the mountain's leeward side can be observed in steeper mountains. The vorticity sheet is stable and can be observed during all timesteps.

Headcurve 40

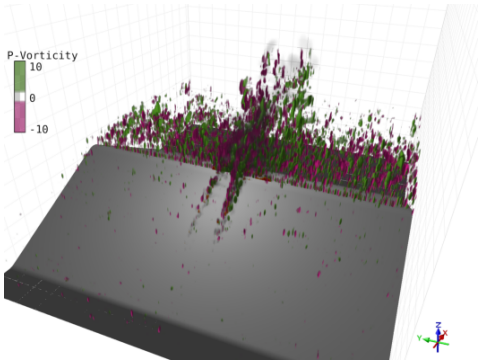

Headcurve 80

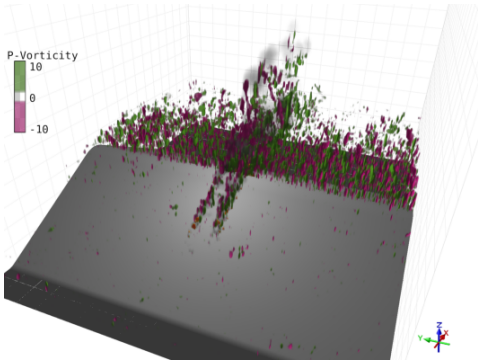

Headcurve 320

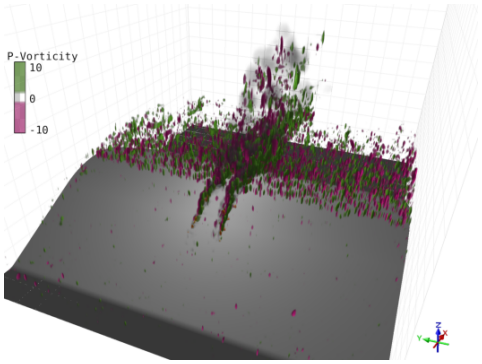

Figure 18: Visualization of the pyrogenic vorticity for timestep 35 in the headcurve runs. In purple, regions where both the x- and z-component are greater than zero and in green, regions where both the x- and z-component are less than zero. Of special interest are the regions at the fire's flanks. The noisy region on the mountain's lee side is caused by atmospheric turbulence.

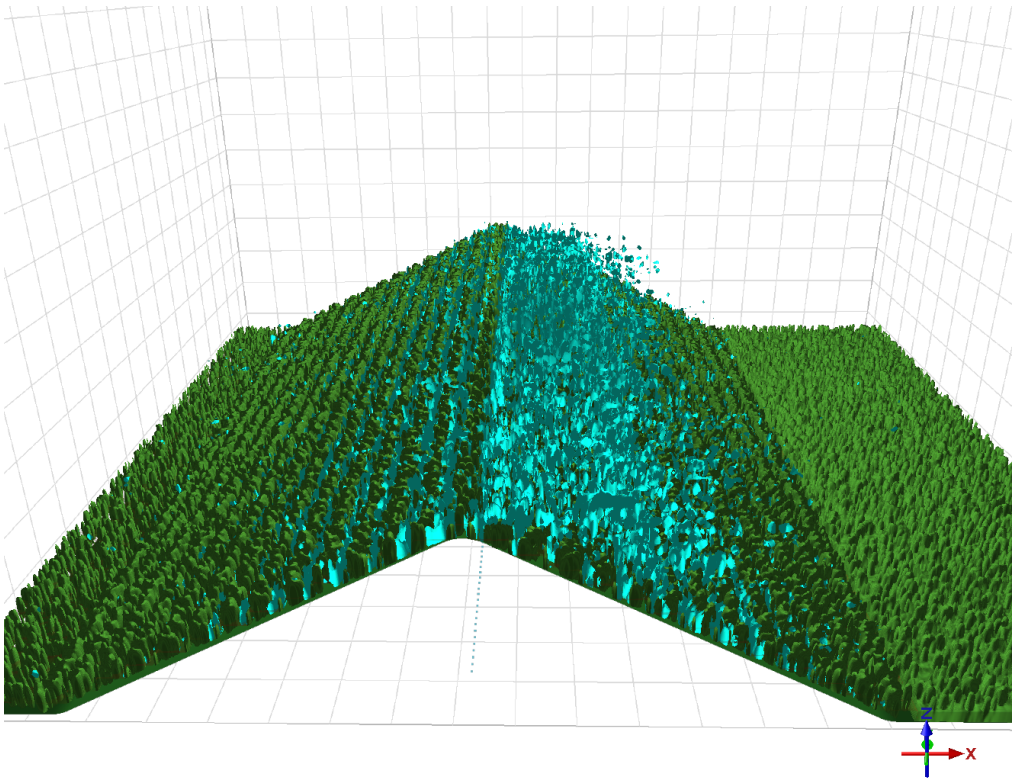

(a) Divergence of the flow field for timestep 4.

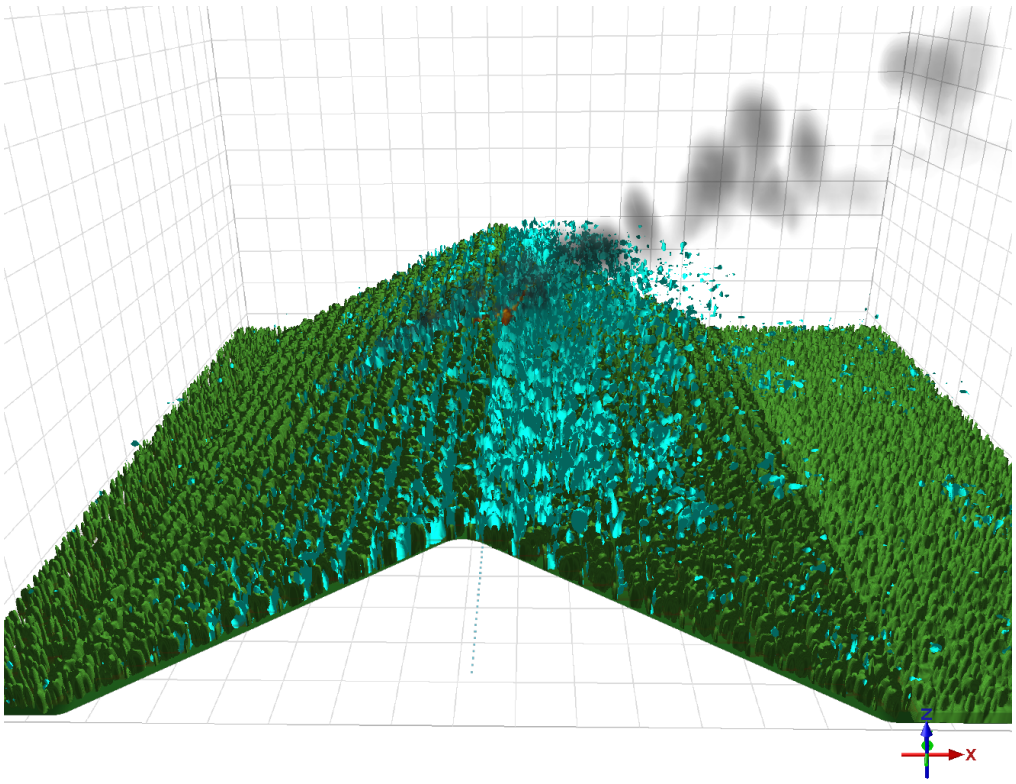

(b) Divergence of the flow field for timestep 20.

Figure 19: Two images comparing the existence of divergence in the timesteps 4 and 20 for the headcurve 40 run. The presence of a divergence $|\mathrm{div}(V)| > 0.4$ is color coded in blue. We can observe the fire causing a divergence in its vicinity.

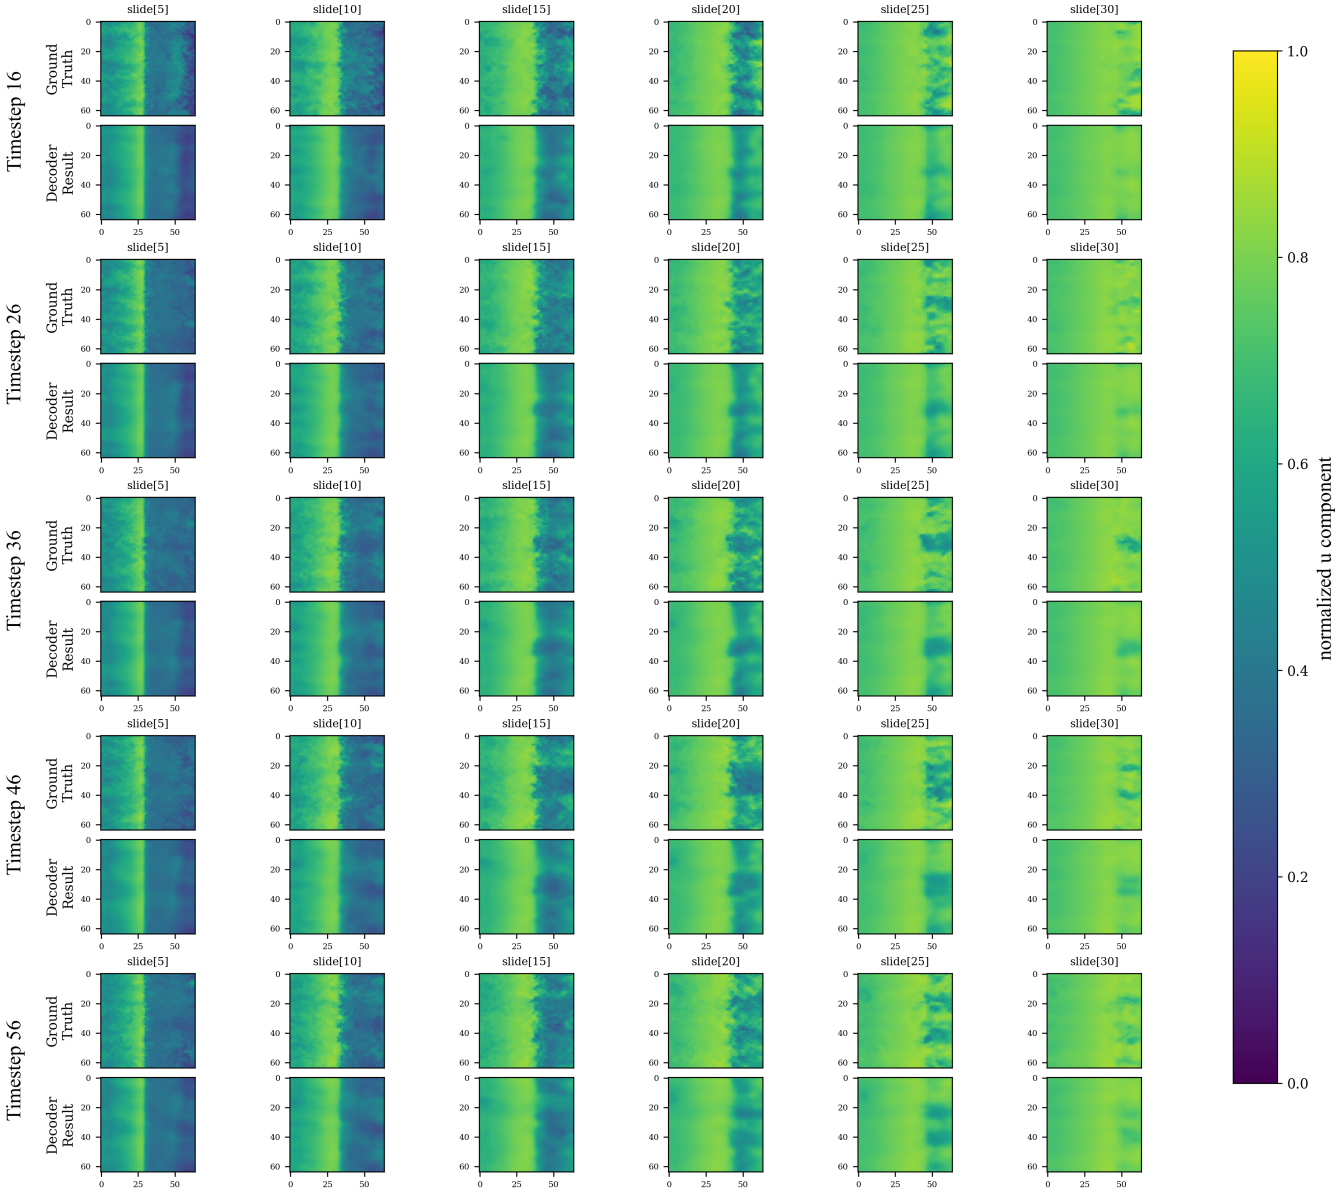

Figure 20: We trained a conditional variational auto-encoder on the u, v and w component of the vector fields of backcurve 40, backcurve 320, headcurve 40, headcurve 80 and headcurve 320 and predicted all timesteps of backcurve 80. Here we show only the u component as the results for the v and w components yielded similar results. The x, y and z domain was resampled onto a regular grid of size $64x64x48$, allowing to efficiently train the model. In this figure, we compare five of the predicted time steps to their ground truth, namely 16, 26, 36, 46 and 56 (from top to bottom). For each of these time steps, we select 6 xy-slices, starting from the ground (5, 10, 15, 20, 25, 30) from the aforementioned resampled z-dimension of size 48.

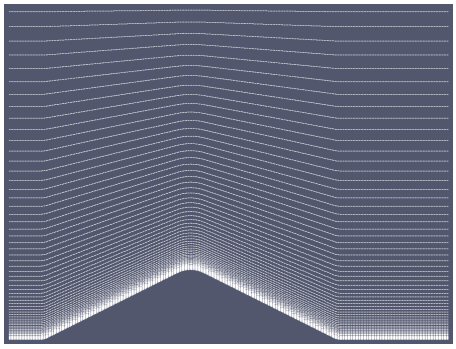

(a) The adaptive grid used for the simulation. The resolution towards the ground level is increased.

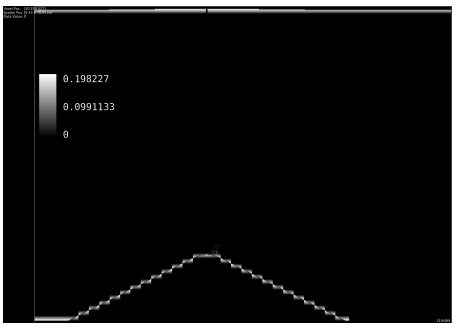

(b) Absolute error for O2 field at time step 40 after resampling onto a regular grid. The error is high at the slopes due to stair-case artifacts. This error is removed after the parametrization.

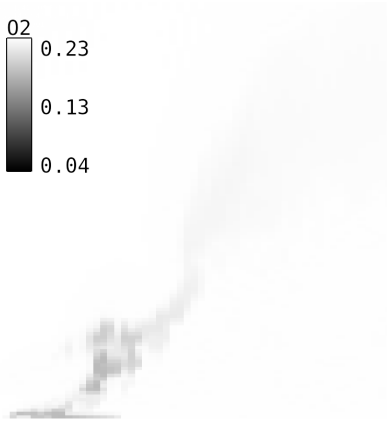

(c) Values of the oxygen field for time step 40 near the surface at the full resolution

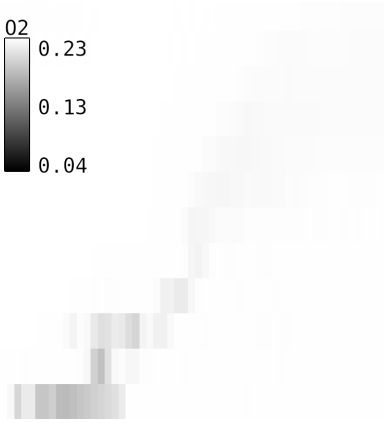

(d) Values of the oxygen field for time step 40 near the surface at the resampled resolution

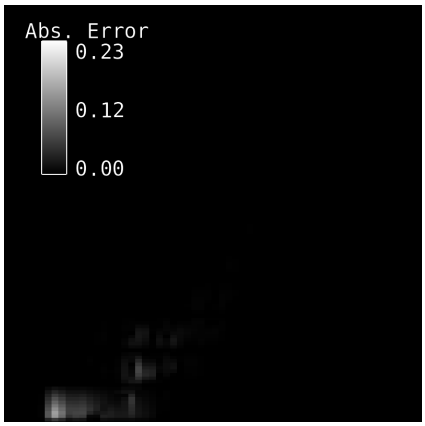

(e) After removing the stair-case artifacts, the relative error can be high (up to 86%) in single voxels near the surface.

Figure 21: For our data analysis, we resample the adaptive grid to achieve a regular grid, for which most of our algorithms have been implemented efficiently. This introduces an error, especially in the lower layers near the surface. For the whole fields, we obtain the following root-mean-square errors (calculated exemplarily for each field of time step 40 of the backcurve 40 run): O2: 0.016; convht_1: $121.1\,\mathrm{W\,m^{-3}}$; frhosierad_1: $151.5\,\mathrm{W\,m^{-3}}$; rhof_1: $0.024\,\mathrm{kg\,m^{-3}}$; rhowatervapor: $0.000056\,\mathrm{kg\,m^{-3}}$; theta: $20.99\,\mathrm{K}$; u: $0.797\,\mathrm{m\,s^{-1}}$; v: $0.316\,\mathrm{m\,s^{-1}}$; w: $0.278\,\mathrm{m\,s^{-1}}$. For our purposes, however, this error is negligible, since the error only is prominent in isolated voxels and because we sample the domain using an eighth of the resolution for each spatial dimension such that those are unlikely to be selected. Hence, the decrease in resolution did not effect the projection of our multi-run similarity plots significantly. For the non-parametrized data, we only analyze features that appear in upper layers (e.g. streamlines), where the error vanishes.

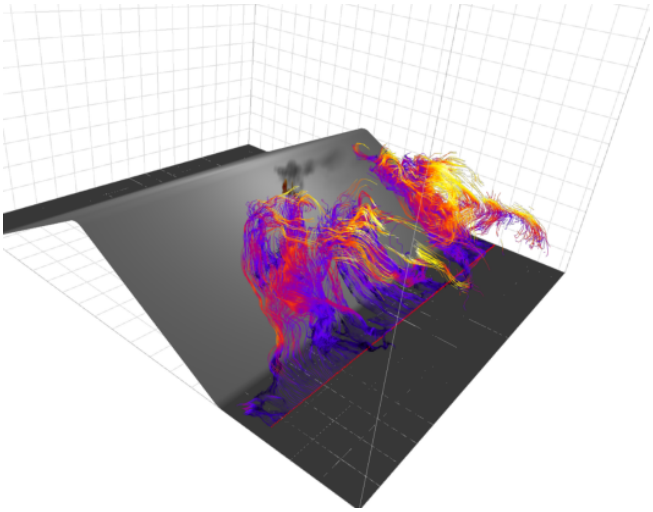

(a) Backcurve 40

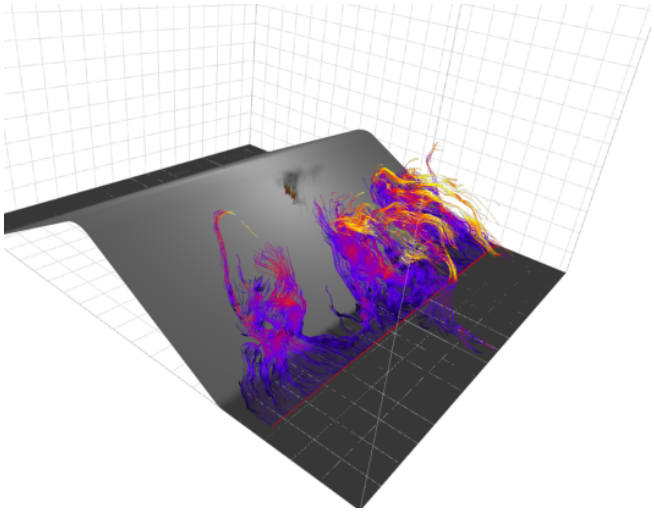

(b) Backcurve 80

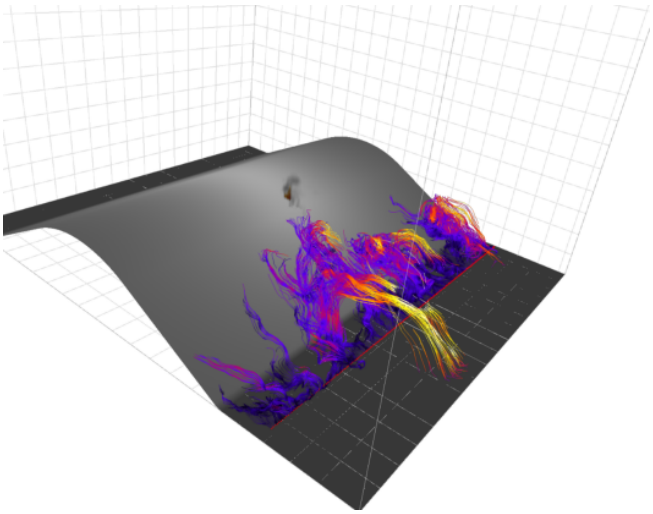

(c) Backcurve 320

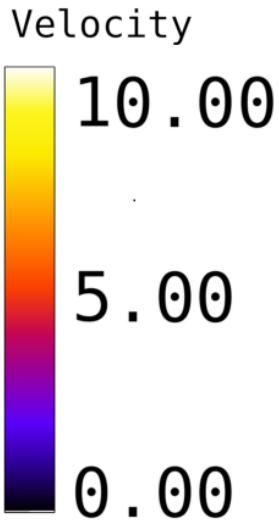

Figure 22: Streamlines for timestep 20 of the lee side upwards draft in the backcurve simulation runs. These upwinds are mainly responsible for the fire's movement towards the mountain ridge in the backcurve simulations. The color of the streamlines is mapped to the magnitude of the wind. A white color indicates a magnitude exceeding $10\ m\ s^{-1}$. Of special interest are the segments of the streamlines in the vicinity of the ground. We observe that in those regions, the density of higher velocity winds increases with the steepness of the mountain.

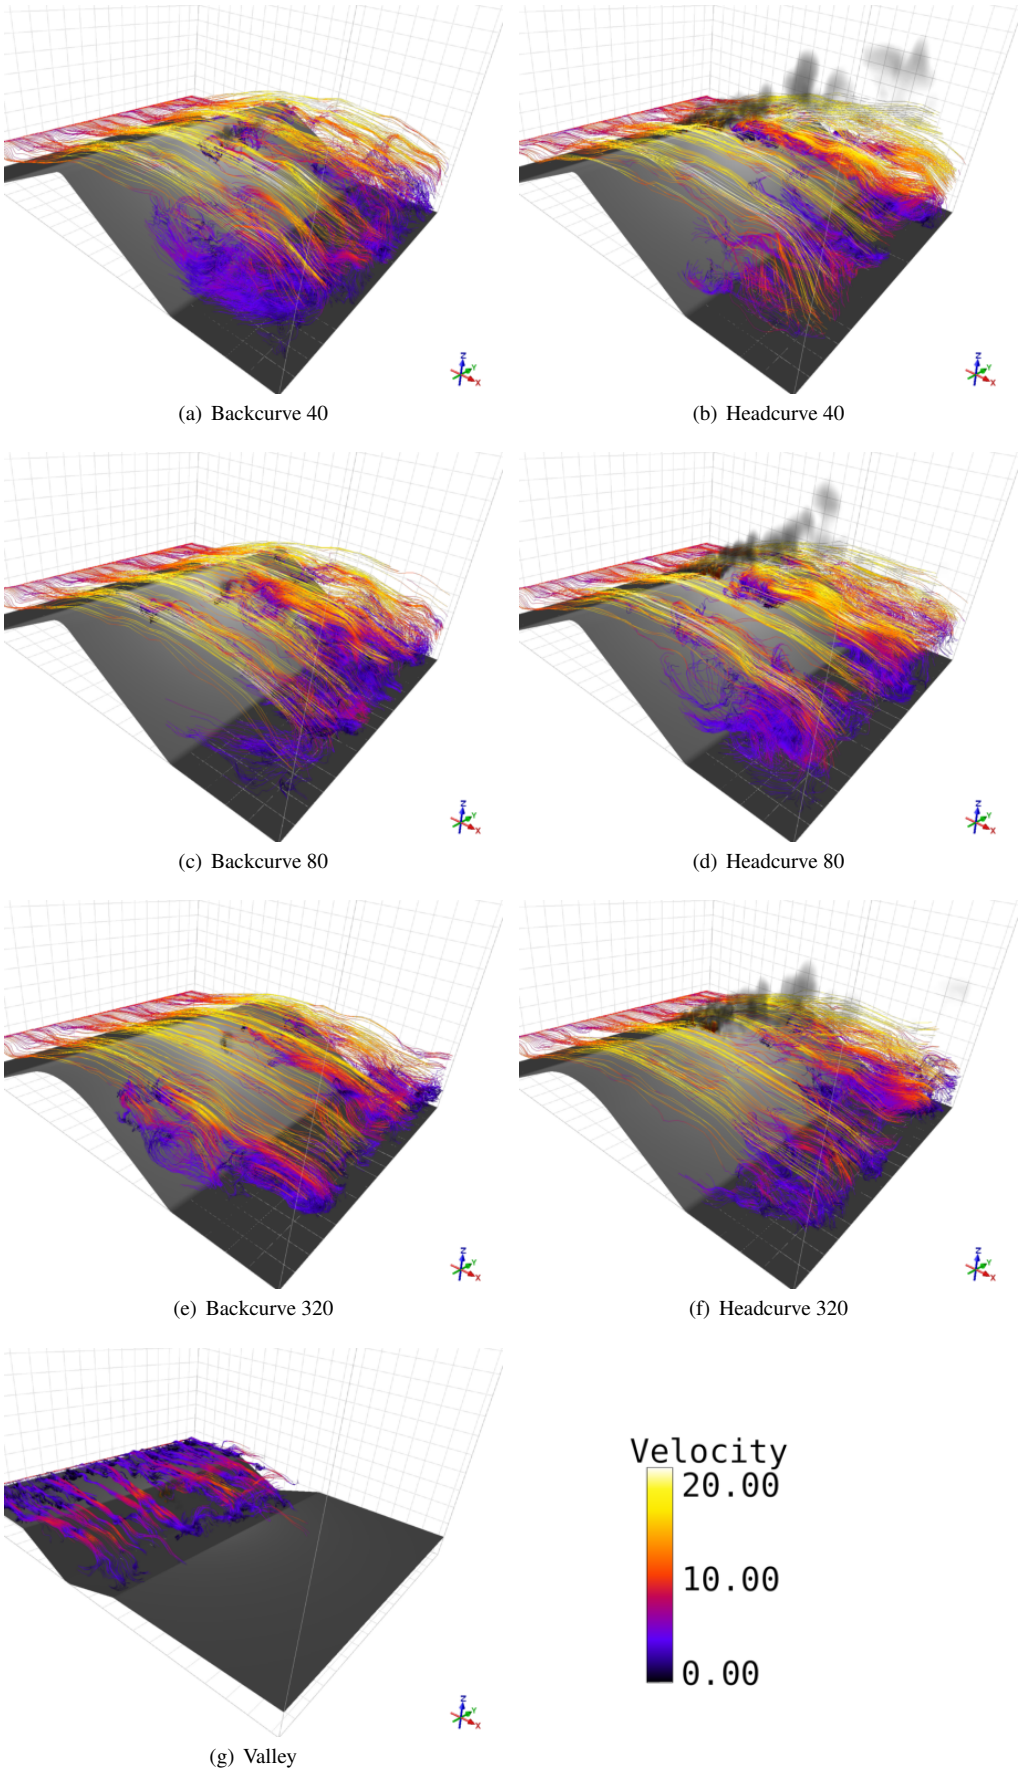

(a) Backcurve 40

(b) Headcurve 40

(c) Backcurve 80

(d) Headcurve 80

(e) Backcurve 320

(f) Headcurve 320

(g) Valley

Figure 23: Streamlines for timestep 20 seeded at the windward side of the mountain. The headcurve visualizations show how the streamwise wind evolves into the updraft at the lee side of the mountain, which - in the backcurve runs - pushes the fire up the hill. As can be seen in the valley visualization, the fire is also pushed up the hill by these winds, thus behaving more like a backcurve simulation in this regard.

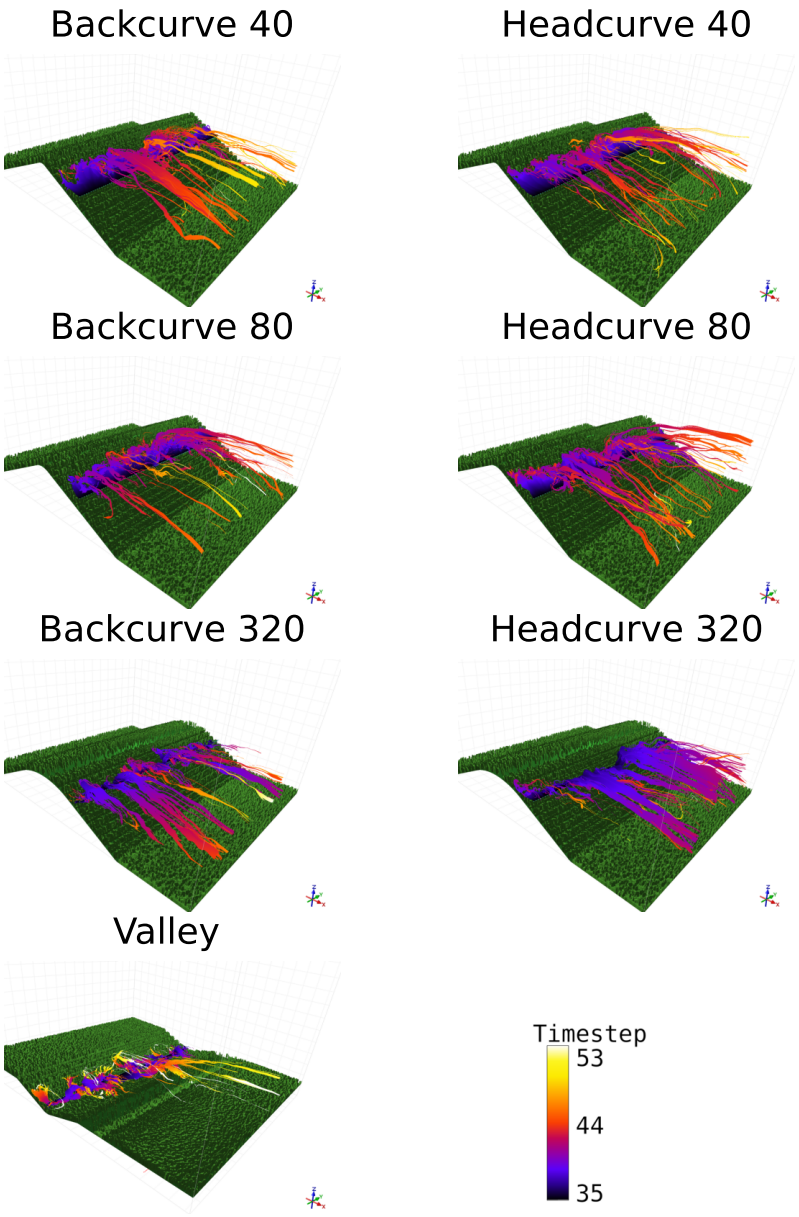

Figure 24: Pathsurfaces for the different runs, integrated from timestep 35 to 55, showing the effects of atmospheric turbulence on the lee side of the mountain. A notable decrease in vertical velocity can be observed in the runs that are less steep. Hence, in these less steep runs the particles follow a trajectory that is more aligned with the general wind direction.