

A Novel Framework for Visual Detection and Exploration of Performance Bottlenecks in Organic Photovoltaic Solar Cell Materials

A. Aboulhassan,¹ D. Baum,² O. Wodo,³ B. Ganapathysubramanian,⁴ A. Amassian,¹ and M. Hadwiger¹

¹KAUST ²Zuse Institute Berlin (ZIB) ³SUNY at Buffalo ⁴Iowa State University

Abstract

Current characterization methods of the so-called Bulk Heterojunction (BHJ), which is the main material of Organic Photovoltaic (OPV) solar cells, are limited to the analysis of global fabrication parameters. This reduces the efficiency of the BHJ design process, since it misses critical information about the local performance bottlenecks in the morphology of the material. In this paper, we propose a novel framework that fills this gap through visual characterization and exploration of local structure-performance correlations. We also propose a formula that correlates the structural features with the performance bottlenecks. Since research into BHJ materials is highly multidisciplinary, our framework enables a visual feedback strategy that allows scientists to build intuition about the best choices of fabrication parameters. We evaluate the usefulness of our proposed system by obtaining new BHJ characterizations. Furthermore, we show that our approach could substantially reduce the turnaround time.

1. Introduction

Organic photovoltaic solar cells (OPV) represent a promising low-cost, low-weight, and flexible alternative for harnessing solar energy. An OPV is a device composed of three main parts: the anode, the cathode, and the so-called Bulk Heterojunction (BHJ) that is sandwiched in between the electrodes (anode and cathode) [WG12], as shown in Fig. 1(a). The BHJ is a blend of two materials, called donor and acceptor, which are separated by an “interface.” The BHJ has a very complex intermixed composition with hierarchical structures spanning several spatial scales.

The photovoltaic process occurs in a sequence of stages: exciton generation, exciton diffusion, charge separation, charge transport, and charge collection. This is illustrated in Fig. 1(a). At each stage of the photovoltaic process, its performance is critically affected by the morphology of the BHJ. The objective of the BHJ design is to maximize the generated photoelectric current. This requires the charges (holes, electrons) to reach the electrodes as fast as possible. To achieve this, the paths of the charges should be as wide and as straight as possible. However, these design criteria conflict with another requirement: increasing the area of the interface surface. In order to harvest more excitons by the interface, the neighboring parts of the interface need to be as curly and as close to each other as possible, in order to increase the probability of excitons reaching the interface.

Even though OPV solar cells bear a great potential, before they will be able to compete with other solar technologies,

some challenges have to be addressed. The most important of these are low efficiency and short life time. Several approaches exist that have resulted in varying degrees of success. One promising approach is to control the BHJ morphology during fabrication. Current BHJ exploration techniques mainly depend on expensive and time-consuming lab tools. These traditional tools deal with the morphologies as black boxes with no knowledge of the photoelectric current within. Their workflow therefore depends on trial-and-error and does not efficiently characterize complex BHJ morphologies with respect to many critical local properties. Accordingly, scientists in OPV research are still lacking a sufficient understanding of the best BHJ material design.

In this paper, we propose a novel framework for exploring one of the critical features of OPV solar cells, called charge path bottlenecks. So far, scientists intuitively refer to bottlenecks as the parts in the BHJ routes that cause contention of charges and hence delay. However, they cannot detect and analyze these features since their tools lack access to the geometric features underlying this phenomenon. Furthermore, the detection of the design structures that reduce bottlenecks is complicated since it involves conflicting design requirements. In our framework, we solve this problem through the following contributions:

- A geometric model that formalizes a previously only intuitive charge bottleneck definition.
- The extraction of new structural features of the morphology that can be correlated with these charge bottlenecks.

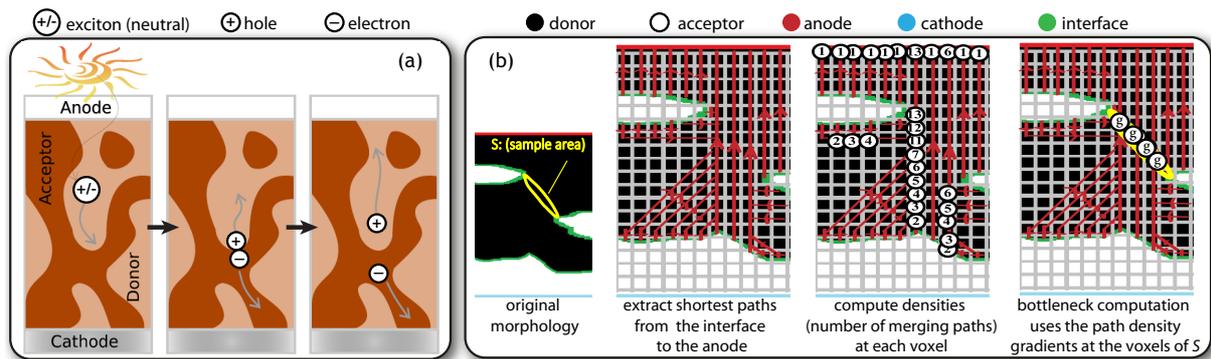


Figure 1: (a) A 2D illustration of the underlying physics of the photoelectric current generation process. (b) A 2D illustration of the bottleneck computation steps (for the donor part). The figure illustrates the computation steps over one sample cross-sectional area (S). We only refer to the donor in the current paper, since the acceptor bottlenecks can be analogously extracted.

- The generation of an abstraction for the BHJ morphology, which we call the BHJ backbone, that visualizes the topology of the structural features of interest. In this way, visual clutter is removed, enabling spatial analysis of the BHJ morphology. Moreover, this abstraction allows for multivariate analysis by sampling the morphology into a minimal set of features that influences the correlation analysis.

However, we see the main contribution of our work in the description of a complete framework for the analysis of BHJ morphologies given as three-dimensional scalar fields.

We demonstrate the validity of our approach by showing how our framework generates new BHJ characterizations. In order to evaluate the results, we have obtained user feedback from domain experts. We also show that our proposed system has enabled dramatic time savings in the exploration process of OPV data, which paves the way for faster exploration of OPV materials in the future.

2. Related Work

This section discusses the most relevant related work grouped into three different categories.

2.1. Analysis of Charge Paths in Organic Photovoltaics

Domain scientists need to detect parts in the BHJ morphology with high charge densities and understand the structural features that cause this problem. To support this goal, scientists have designed a simulation of this phenomenon for sinusoidal structures [GN07], and have correlated the sine width to the charge density. This simulation is suitable for regular structures but not random ones, such as the BHJ structures. Hence, this simulation was later extended for BHJ [KG12]. However, it was still not possible to correlate structural features with charge densities. As a result, domain scientists started to move in the direction of studying the geometric features of charge paths rather than the behavior of the charges themselves. For this purpose, scientists have developed an approximation model [WTCG12b] that extracts

a representative set of charge paths that reflects the physical intuition. However, that work depends only on statistical analysis and is therefore not suitable for exploring geometric features such as bottlenecks. Furthermore, it is unable to explore the interplay between conflicting design parameters. We address these limitations in the current work.

2.2. Morphology Abstraction and Feature Extraction

Our bottleneck model requires the extraction of local domain features as well as the measurement of local properties such as their size. A Voronoi-like decomposition of the pore-space of porous materials has been proposed [HBP*12] to aid the determination of the pore space skeleton. In our approach, we employ a hierarchical watershed algorithm [Beu94, CB97] on the distance map [JBS06] with persistence-based simplification [ELZ02]. This enables the decomposition of the BHJ morphology at potential bottlenecks. In addition, we propose an abstract model that simplifies the BHJ structure. This model is based on a conventional thinning algorithm [Pud98]. Examples of other techniques to achieve simplified representations include Reeb graphs [PSBM07], distance field-based methods [GDN*07], extremum graphs [CLB11], and other topological methods [BGG09, HBP*12]. Geometric path computation is also important to analyze molecular structures [KFR*11, LBH11], but these methods do not apply to our application.

2.3. Geometric and Visual Path Analysis

In our work, we use graph-based models to enable knowledge-based exploration of the underlying data. Previous work applied knowledge-based visualization of charge paths to molecular data [ALK*12]. In this work, we exclusively focus on nanoscale data. One requirement for our method is to visualize charge paths in relation to the surrounding geometry. Similar ideas have been used before [JM10, LGD*05], but employed vector fields resulting from simulations. In our work, we instead employ a set of representative paths that are extracted based on the intuition of domain scientists about the physics of OPV materials.

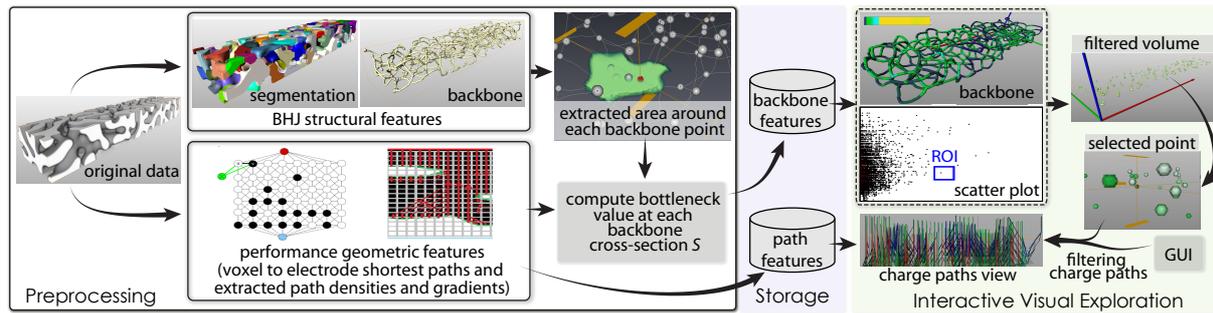


Figure 2: An overview of our framework comprising preprocessing, storage, and interactive visual exploration. The workflow illustrates the dependency among different system modules as well as typical steps performed by users for visual exploration.

In general, abstractions are often needed to be able to focus on important features, both for visualization and for comparative and statistical analysis. Path abstraction models include the ones proposed in [TvW99, LCMH09, HBP*12]. These abstractions can be rendered using simple lines or triangulated tubes, but also using more advanced rendering methods [KLG*13, KGPS13].

3. Overview

This section provides an overview of our framework, describing the type of data we are dealing with, the crucial abstraction of charge paths, and the overall workflow.

3.1. BHJ Data

The BHJ morphologies that we are analyzing in this work result from computer simulations (see Sec. 9 for details). Each morphology is given as discrete scalar field, where each voxel is assigned an acceptor volume fraction value ϕ . These fraction values are between 0 and 1. By tracing the distribution of these variables, individual phases can be identified, i.e., $\phi = 0$ corresponds to pure donor, while $\phi = 1$ corresponds to pure acceptor, respectively. In the regions separating individual phases, the volume fraction changes smoothly across the thin interface. By reconstructing the iso-contour corresponding to the iso-value $\phi = 0.5$, the interface can be identified. We use this interface in our subsequent analysis.

3.2. Charge Path Approximation Model

A charge path is the trajectory of a charge from the point of its creation to the corresponding electrode (either anode or cathode). In this work, we study paths for two types of charges: excitons and holes. Fig. 1(a) illustrates that (1) the path of an exciton starts from the point of creation in the donor and ends at an interface, and (2) the path of a hole starts from the point of creation on the interface and ends at some point on the anode. Note that electrons can be handled in the same way as holes. The only difference is that they travel through the acceptor to the cathode. We are interested in the shape of the charge paths rather than tracking the charges themselves. For this purpose, we use the set of

shortest paths based on a model that has been proved to provide sufficient information about the geometry of the whole charge paths vector field [WG12].

3.3. Workflow

We summarize the proposed system in Fig. 2. Our workflow is divided into three main parts: *preprocessing*, *storage*, and *interactive visual exploration*.

In the preprocessing step, we compute the bottleneck indicator at certain points and the geometric features of each charge path. The computation of the bottlenecks requires extracting cross-sectional areas between neighboring parts of the interface. Before extracting these areas, we first simplify the morphology into its backbone, defined as the medial axis of the morphology. The backbone provides a reduced view of the morphology supporting correlation analysis as well as spatial exploration of the cluttered parts. Then, we extract the areas around each voxel on the backbone only. Moreover, we compute the set of shortest paths from the morphology as well as features of interest of these paths.

In the storage step, we cache the data resulting from preprocessing to eliminate unnecessary computations in the subsequent interactive visual exploration step.

For interactive visual exploration, we provide a variety of views: the backbone view, the charge paths view, scatter plots, and volume rendering. The backbone visualizes relevant information via user-defined color codings. Scatter plots visualize data derived from the backbone. Users can explore multivariate correlations via the scatter plots as well as filter a volume with respect to a certain range of parameters. For spatial analysis, users select one point in the filtered volume, and retrieve the charge paths around this point to explore their features. Users can select regions of interest from the whole set of charge paths using simple GUI widgets.

4. Morphology Simplification

In order to be able to analyze the BHJ morphology more effectively in terms of charge paths, we first need to simplify

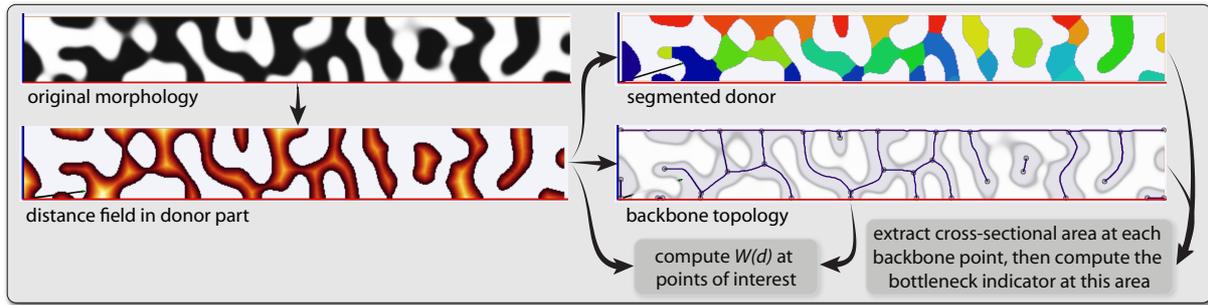


Figure 3: A 2D illustration of the individual steps of extracting the structural features that are used for computing the bottleneck indicators and exciton diffusion probabilities. Note that in our framework these steps are performed entirely in 3D.

it. For this, we first compute the backbone of the BHI morphology. We then compute cross-sections of the morphology, which is important for the calculation of the charge path bottleneck indicator that we will introduce in Section 5.1.

4.1. Backbone Computation

As a first step to analyze the morphology, we compute its backbone to provide a less cluttered visual representation as well as to make the computations more efficient. The backbone is defined as the medial axis of the morphology (see Fig. 3). To compute the medial axis, we apply thinning [Pud98] based on the Euclidean distance field of the donor part with respect to the interface. The thinning is performed by removing voxel by voxel from the segmented object until only a string of connected voxels (the skeleton) remains. The voxel skeleton is then converted into a spatial graph that passes through the medial axis of the donor. The Euclidean distance to the nearest boundary is stored at every point in the spatial graph. This structure simplifies the multivariate morphology analysis as discussed in the subsequent sections. Furthermore, it allows us to compute cross-sections of the morphology along the backbone.

4.2. Extraction of Cross-Sectional Areas

In order to identify potential bottlenecks in the BHI morphology, we need to segment the whole morphology into areas that reflect the gradual change in the routes' thickness from wide to narrow regions. To achieve this, we start with a 3D segmentation of the whole morphology that decomposes the donor morphology at the constrictions (Fig. 3).

For this, we first compute the signed Euclidean distance map, starting from all interface voxels, such that the distance map inside the donor has negative values. We then apply a watershed algorithm [CB97] on the distance map and subsequently apply a persistence-based [ELZ02] merging step to create larger regions. To do so, we compare the scalar minima of two regions to be merged with the scalar value at the potential merge point. If the difference between one of the minima and the scalar value at the merge point is below a user-defined threshold, we merge the two regions. Otherwise, the regions are not merged.

In the following steps, we are only interested in the cross-sectional areas around the points on the morphology backbone. To determine these cross-sections, we compute the directional vectors (tangents) of the backbone in each point of the backbone. The point on the backbone together with its directional vector determine the plane of the cross-section.

Now, we can easily compute the 2D cross-sectional area around each backbone point from the 3D segmented volume. This is done by intersecting the plane given by the backbone point and the directional vector with the 3D volume segmentation, yielding a 2D cross-sectional area as shown in Fig. 4. The cross-sectional area includes all voxels in the plane that belong to the same region as the point on the backbone.

5. Intuitive Geometric Models for Performance

This section introduces two geometric indicators related to the performance of a given BHI material (morphology).

5.1. Geometric Indicators for Bottlenecks

Intuitively, if a charge path s at any point on the path does not overlap with any other charge path, then the charge that travels along the path s will be able to travel at maximum speed. Conversely, the more paths share the same voxels, the larger the possibility that charge contention happens and hence a corresponding delay in charge transport is incurred.

Paths merge when they move from a wider area into a narrower one, similar to an hour glass. Based on this observation, we define the bottleneck indicator $K(S)$ at each cross-sectional area S inside the morphology as follows.

Let $j: \mathbb{R}^3 \rightarrow \mathbb{R}$ be a scalar volume of path density. Considering a cross-sectional surface S (as illustrated in Fig. 1(b)), we obtain the corresponding total bottleneck value $K(S)$ as the surface integral of the magnitude of the path density gradients, normalized by the area A of S :

$$K(S) = \frac{1}{A} \int_S \|\nabla j\| dS. \quad (1)$$

Areas in the morphology with large $K(S)$ indicate the charge path bottlenecks that the domain scientists are interested in.

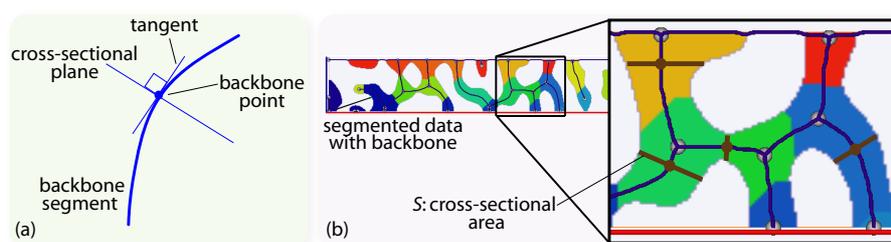


Figure 4: A 2D illustration of the cross-sectional area extraction guided by the 3D segmentation and the backbone. The cross-sectional area at a backbone point is the intersection between the plane perpendicular to the tangent at this point and the current segment of the 3D segmentation. (a) illustrates how the plane of the cross-sectional area is determined. (b) illustrates different examples of cross-sectional areas after intersection with 3D segments. Note that we perform these computations entirely in 3D.

5.2. Geometric Indicators for Exciton Diffusion

Another important measure for the effectiveness of the BHJ is the probability of an exciton actually reaching the interface. This *exciton diffusion probability* can be computed via the following equation [SRS08, WTCG12a]:

$$W(d) = e^{-d/L_d}, \quad (2)$$

where d is the shortest distance from a point in the donor to the interface, and L_d is a material-specific constant: the *exciton diffusion length*. Thus, our geometric model for exciton diffusion is simply the distance field that was computed inside the donor with respect to the interface (see Fig. 3). Hence, we can directly use the value of the distance field as parameter d in Eq. 2 to obtain $W(d)$ at any desired point.

6. Feature Extraction

This section describes how the BHJ features required for subsequent analysis are extracted in our framework.

6.1. Size of the Cross-Sectional Areas

Our domain science collaborators want to explore the correlation of the sizes of cross-sectional areas to the bottleneck indicator defined in Eq. 1. We compute this size as the number of voxels intersecting the cross-section, since this is also the smallest unit used for the features of charge paths.

6.2. Distance between the Interface and the Backbone

$W(d)$ (Eq. 2) is an important measure for the domain scientists, because correlating it with d allows them to estimate the effectiveness of the BHJ. We propose to compute d at the backbone points only, since they provide the worst case for $W(d)$, i.e., the farthest distance from the neighboring parts of the interface.

Moreover, by extracting the bottleneck at the same backbone point as shown in Figs. 3 and 4, we get an effective minimal set of indicators for the trade-off between $W(d)$ and $K(S)$. For the computation of $W(d)$, in this paper we use an exciton diffusion length of $L_d = 10$ nm.

6.3. Charge Path Features

The bottleneck analysis discussed so far summarizes the behavior throughout the whole morphology. However, scientists still need to explore the reasons behind this behavior in detail. In order to support this, we allow scientists to visualize the charge paths around each point on the backbone. Scientists can then explore details about the bottleneck such as the corresponding path density (see Fig. 1).

One other feature of interest is the *tortuosity*. Bottlenecks are not the only source of delay in the charge transport; the path length may also play an important role. We allow scientists to explore this feature through the tortuosity indicator. This indicator was also used by domain scientists in previous statistical contexts [WTCG12b]. To compute the tortuosity, we determine the length L of the shortest path from any point in the donor to the electrode and relate it to the ideal path length C , i.e., the length of a straight line between the ends of the path without constraints. The tortuosity τ then is

$$\tau = \frac{L}{C}. \quad (3)$$

7. Visual Exploration

This section describes the capabilities for visualization and interactive exploration that comprise the interactive visual exploration step of our framework, as illustrated in Fig. 2.

7.1. Backbone Visualization

We visualize the topology of the backbone by rendering trace lines between all consecutive nodes (see Fig. 5).

Exploration of a variety of scalar attributes on the backbone is enabled via color-coding (1D transfer functions). This is illustrated in Fig. 5. In order to facilitate the visualization of attributes, for each point on the backbone the preprocessing step has pre-computed the bottleneck value, the area extracted by segmentation around this point, and its shortest distance to the interface. From the latter we also compute and visualize the exciton diffusion probability (Eq. 2).

7.2. Scatter Plots

We create scatter plots by mapping each point on the backbone to one point in the scatter plot. The user can choose any

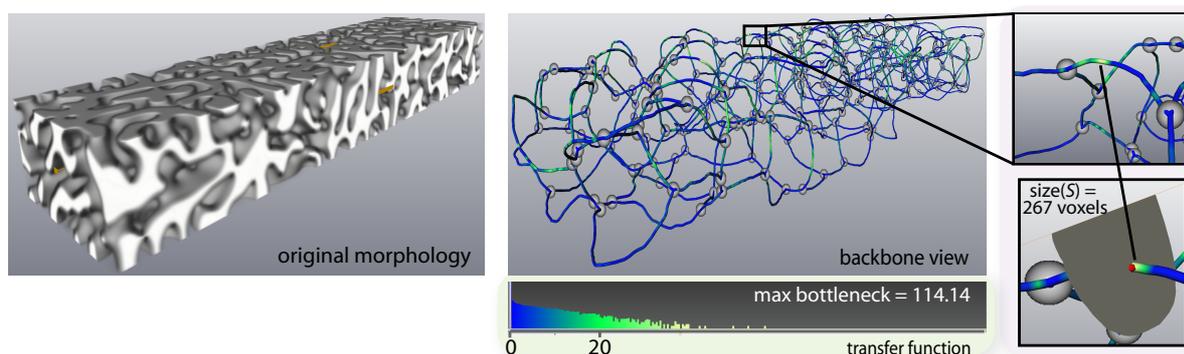


Figure 5: An illustration of spatial exploration of the backbone: The backbone is color-coded via a user-defined transfer function that highlights values of interest for the user ($K(S) = 20$ in this example). The user can then move a point probe to regions with a high bottleneck value (here: yellow regions), and explore the surrounding area to ascertain its shape and size. This information can guide experts in enhancing the morphology, e.g. by increasing the sizes of cross-sectional areas.

two backbone attributes as x and y dimensions of the scatter plot: bottleneck value, area size, distance to interface, and exciton diffusion probability (see Fig. 6).

To reflect the number of voxels that are mapped to a single pixel in the scatter plot, we use a heat map; red represents a large number, black a small one. We further enhance the visualization of the path density by binning.

Our scatter plots serve two main purposes: (1) exploring correlations between structural and performance features, and (2) filtering by brushing in the attribute domain.

Brushing in any scatter plot enables users to select regions according to attributes to be investigated further using the spatial views, e.g., volume rendering. After a brushing operation in a scatter plot, only voxels with values in the specified range will be displayed and rendered (see Fig. 6).

7.3. Charge Path Visualization

Charge paths are rendered as trace lines to allow for studying their topology. These lines are color-coded according to scalar attributes, such as tortuosity (Fig. 7), or path density.

Interaction. We allow users to select a spatial region of interest to reduce visual clutter in the path visualization. In order to explore the paths around a specific location, the user needs to select this location in the spatial domain.

For this interaction, we use a point probe: a ball attached to three orthogonal lines parallel to the x, y, and z dimensions, respectively. The user can interactively move this probe to any point of interest inside the bounding box of the morphology. Then, the system retrieves all paths that pass through a region of interest around the ball.

We also support further filtering via GUI widgets, such as the maximum size of path bundles and the maximum path length in each bundle. This is illustrated in Fig. 7.

8. Implementation Details

To extract charge paths, we use the software GraSPI [WTCG12b], which has been used successfully in previous OPV research. GraSPI is based on the Boost graph library. Details can be found in [WTCG12b]. We run the charge path computation process offline, generate the corresponding topology, and store it in text files. Then, we load these data into our visualization software. Our visualization approach is implemented in the Avizo framework with both computation and interaction modules. The ZIB version of Amira [Ami] on the other hand is used for generating the segmentation. GraSPI uses an equivalence between voxel-wise data and a graph to effectively characterize the morphology.

By translating the discrete morphology into a graph, GraSPI can use standard graph algorithms to find the shortest paths and connected components. The graph is constructed by considering each voxel in the morphology as a node. Each node (voxel) gets a label: black for donors, white for acceptors, green for interface voxels, red for anode, and blue for cathode. An edge is created between each voxel and its 26 neighbors. Each edge is given a weight according to the distance between the two corresponding voxels (i.e., $1, \sqrt{2}, \sqrt{3}$). GraSPI then uses the standard Dijkstra algorithm to extract the set of shortest paths from this graph.

9. Evaluation

As case studies, we employ two different simulated data sets that represent two main BHJ lab synthesis techniques, which are called *Solvent-Based Fabrication* [WG12], and *Thermal Annealing* [WTCG12b], respectively.

The former is used to create morphologies with different patterns and connectivity by manipulating physical parameters such as pressure, donor/acceptor percentage, etc. The latter is used to enhance the morphology performance by

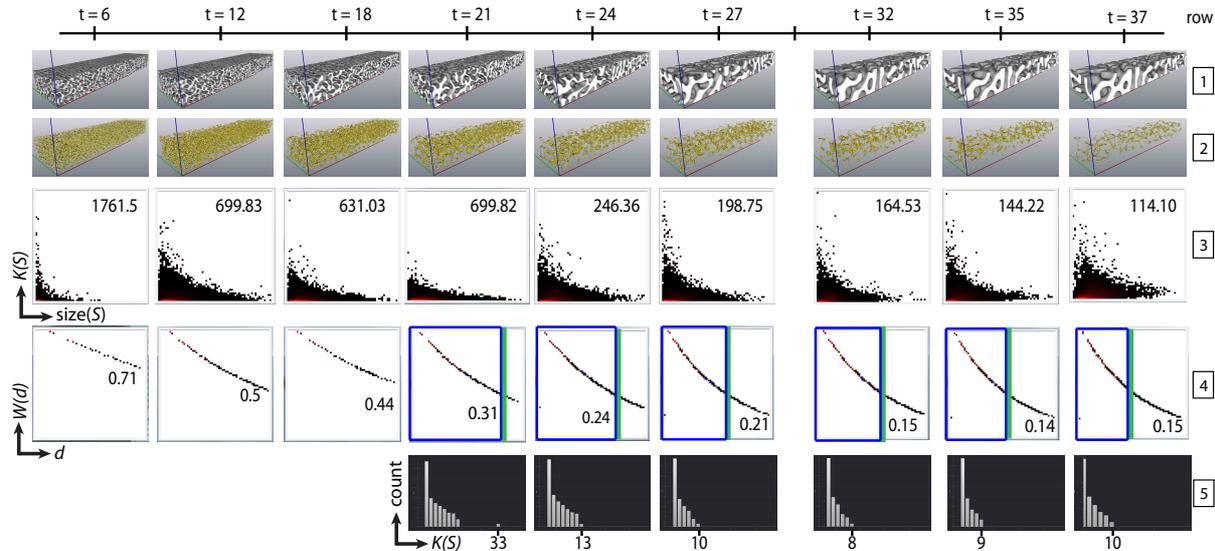


Figure 6: A comparison between different time steps of thermal annealing of Morphology A. Row 1 depicts the shape of the morphology at each time step. Row 2 illustrates the corresponding backbone. Row 3 shows the correlation between the sizes of cross-sectional areas and bottleneck values $K(S)$ using scatter plots. The numbers indicate the highest bottleneck value. Row 4 shows correlations between distances d and exciton diffusion probabilities $W(d)$. The green lines show the value at $d = 10$ nm, while the blue rectangle shows the interactive selection of all the points with $d < 10$ (this is displayed only for morphologies that exhibit a distance of 10). The numbers indicate the lowest probability value. Row 5 illustrates histograms of the bottleneck value distributions of the morphologies after filtering the corresponding backbone in Row 2 via the blue rectangles in Row 4.

successively coarsening it. This is mainly achieved by exposing the morphology to a certain temperature for a certain period of time. Scientists need to make decisions for the parameters to use in these experiments.

The main analysis task performed by material scientists is *characterization*. The aim of characterization is finding correlations between the structural features and performance features in order to decide how to design the BHJ morphology. By using the previous tools and workflow for our case study, scientists characterize fractions of the material with respect to a certain performance metric. For example, scientists need to know which excitons (created in the donor) will possibly recombine before reaching the interface. They already have knowledge that recombination will happen if the distance is longer than $L_d = 10$ nm [WTCG12a]. Similarly, scientists need to study other features such as the tortuosity that also should be less than a value of $\tau = 1.1$ [WTCG12b].

Accordingly, they need to know how thermal annealing influences these fractions in order to learn their ideal values. Lab experiments cannot enable these types of characterization since they provide no access to this level of detail. A successful step towards obtaining this characterization is through studying charge paths, as shown in [WG12]. However, this earlier work only used standard statistical methods, which limits it with respect to two major considerations: (1) It is unable to detect bottlenecks, and (2) it can only study fractions as independent parameters without the critical in-

terplay between them. We show in our evaluation how our new proposed model succeeds in removing these limitations.

The two data sets that we have used in our evaluation are summarized in Table 1.

9.1. Feedback of Domain Experts

This section discusses feedback of our domain science collaborators on (1) producing new BHJ characterizations, and (2) the value of the provided visual analysis capabilities.

9.1.1. Producing New Characterizations

Fig. 6 demonstrates how our system enables scientists to extract novel BHJ characterizations. These characterizations guide scientists to strategies for tailoring morphology structures (in the lab) that have fewer bottlenecks and better path quality. The results show that the maximum value of the bottleneck indicator $K(S)$ continuously decreases as the thermal annealing proceeds until the last time step, while for the

Data Set	# Time Steps
Morphology A	37
Morphology B	39
Dimensions	$561 \times 141 \times 71$ voxels

Table 1: The data sets used in our evaluation. We have used two 3D morphologies, each of which consists of several time steps (one volume each) computed via thermal annealing.

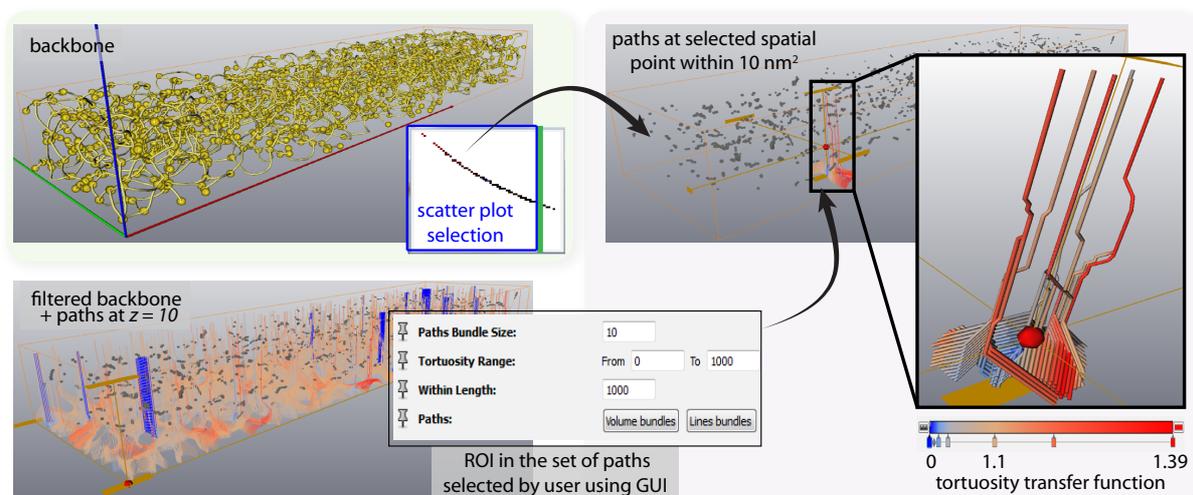


Figure 7: Charge path exploration for a sample morphology: Time step 21 of Morphology A. Typically, visualizing all charge paths results in too cluttered visualizations. Hence, the user needs to navigate to a region of interest for which the paths are then displayed. In the right column, the backbone of the data set as well as a big subset of the charge paths are shown to provide a feeling for their cluttered nature. In the middle, a user-selected region in a scatter plot filters the backbone down to only a few points. Then, the user selects a region of interest (10 by 10 voxels), whose center is a point interactively probed in the filtered volume, using the GUI shown in the center. Finally, the final selected paths are color-coded according to tortuosity.

exciton probability diffusion $W(d)$ the minimum is already reached before the last time step. Domain experts have commented that these correlations match their intuition and their simulation results. The experts then selected the distances less than 10 nm (similar to [WTCG12a]). The results show that - until time step 21 - $W(d)$ is 100% optimal while $K(S)$ has the highest values.

Using this combination of $K(S)$ with the $W(d)$ analysis, the scientists could observe that the most favorable structures do not necessarily exist at the coarsest morphology but earlier than the last time step (37). The determination of the most favorable time steps can be made by using the multi-view brushing feature of our system. The experts increased the size of the blue rectangle such that it encloses all points with smaller distance than 10 nm (Fig. 6, fourth row). They then generated filtered bottleneck volumes at the backbone points corresponding to the brushed value. The histograms of the filtered values are shown in the fifth row of Fig. 6.

Scientists concluded that most favorable time steps lie somewhere between time steps 21 and 27 because $K(S)$ values are in general not high while more material fraction is distributed within the 10 nm distance. To make this decision, they further explored the time steps in detail and compared the results. We demonstrate the scenario of this exploration through time step 24.

Fig. 5 illustrates how spatial analysis is used to provide a detailed level of analysis. The scientists explored if it is possible to make time step 21 match time step 24 more closely by reducing the bottlenecks. First, they explored the whole backbone for points with bottlenecks more than $K(S) = 20$

(an adequate value for comparison). Then, they displayed the area around each point to ascertain its shape and size and to explore strategies for editing it.

Fig. 7 illustrates the usefulness of the charge path visualization. After filtering the volume to include only the useful material, the experts selected the lowest points, since they have potential for higher tortuosity. In Fig. 7, they can see the charge paths with a tortuosity higher than 1.1 and identify the route of these paths. The domain scientists have commented on this approach: “Having these regions identified, it is possible to explore various ways for removing or mitigating the effect of bottlenecks, e.g., by increasing the local cross-sectional area. This step is particularly important considering the multi-step nature of the photovoltaic in organic solar cells, as an improvement of one performance indicator can result in the deterioration of others.”

9.1.2. Feedback on the Visual Analysis

Besides deriving the novel characterizations mentioned in the previous section, the scientists have also commented on the visual analysis framework that we provide.

First, they have made the following general comment on the framework: “The tool provides means to develop intuition regarding linking morphology with performance. Ultimately, we envision this tool to enable design of fabrication that leads to desired morphologies with improved properties. An understanding how to improve performance by locally modifying morphology is a very crucial step.”

On the other hand, the scientists have requested the following additions. First, they found it helpful to enrich the backbone visualization. They showed particular interest in

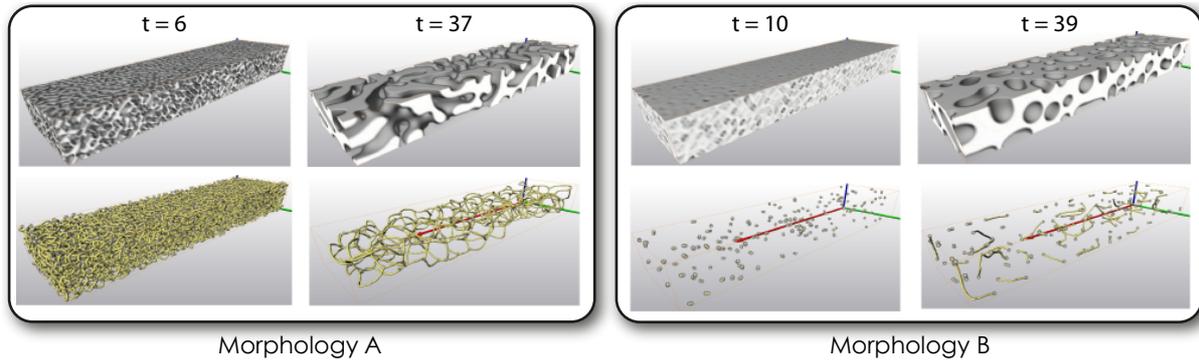


Figure 8: A comparison between two different data sets (Morphology A, and Morphology B). The backbone visualization reveals that almost all parts are well-connected in Morphology A. However, Morphology B has many obvious islands (disconnected parts) for all time steps of thermal annealing. The backbone helps with exploring this feature even in the cluttered parts of the morphology. Scientists critically require this information, because a higher number of islands leads to reduced charge transport.

the branching vertices, since they could quickly infer the potential for bottlenecks from them. Moreover, they commented on the great benefit that the backbone provides, besides the analysis introduced in this paper, since it reveals connectivity. It shows all the routes from one point in the donor to the anode. Furthermore, it can instantaneously detect islands as shown in Fig. 8. Connectivity is important because disconnected parts (islands) will trap the charges rather than transport them. By a quick visual comparison, scientists can see that Morphology B has a lot of disconnected parts vs. Morphology A which indicates less charge transport. The scientists also showed interest in displaying the segmentation as well, since similar conclusions can be made from the boundaries between the segments.

One limitation of our framework pointed out by the domain scientists is the tortuosity analysis. The current framework could visualize, for the first time, the local tortuosity. However, it still depends on exploring each point in an exhaustive manner. It would be useful to have a quick summary first, similar to the analysis in Fig. 6, which requires creating a new model for this property.

9.2. Performance Analysis

The GraSPI run time for a typical 3D morphology (with 5.5M voxels) is 15 min on a typical work station (Intel Xeon Quad 2 GHz, 12 GB of RAM). The visualization is run on Intel Xeon X5550, 2.67 GHz processors, 24 GB RAM. The morphology generation is run for 20 hours on a 160-nodes cluster, each node with dual quad core AMD Barcelona 2.2 GHz and with 8 GB RAM.

Action	Avg. CPU Time [s]
Preprocessing	99
Stored Data Loading	6.83
Volume filtering via scatter plots	0.56
Selected lines: 1713068 lines	27.47
Selected lines: 3800 lines	0.75

Table 2: Running times for all steps of our framework.

The visualization time taken by the main interaction tasks is illustrated in Table 2. We display the average time taken for Morphology A, time step 6, since it includes the largest interface surface and the most complex structure. We notice that the most time-consuming step is the preprocessing. However, once the data are generated and stored, the rest of the actions are quite interactive except in case of too many lines selected for visualization. However, usually users avoid a too large number of lines, since the corresponding visualizations become too cluttered.

These results show that the time required to perform analysis tasks in general are dramatically reduced, since the lab experiments that comprised our collaborators' previous workflow can take days to generate a single sample, and it can therefore take months to reach conclusions.

10. Conclusions and Future Work

We have proposed the first framework for visual detection and analysis of performance bottlenecks in OPV materials based on geometric features of charge paths. To visualize the complex BHJ morphologies, we use a novel visual representation called the backbone, which provides a suitable geometric abstraction. We have shown how this abstraction enables efficient multivariate analysis. Our framework has helped domain scientists to produce novel characterizations, while at the same time drastically reducing the analysis time.

In the near future, we plan to extend the analysis to include more variables. Moreover, we plan to develop tools for editing the BHJ morphology to further accelerate the design of improved OPV materials.

Furthermore, our novel framework could pave the way for analyzing similar complex material morphologies, such as porous media, critical to other fields of science.

Acknowledgments

This work was supported in part by King Abdullah University of Science and Technology (KAUST). The third and fourth author were supported in part by Global Collaborative Research, KAUST: CRG-1-2012-THO-015-ISU.

References

- [ALK*12] ABOULHASSAN A., LI R., KNOX C., AMASSIAN A., HADWIGER M.: Crystalexplorer: An interactive knowledge-assisted system for visual design of solar cell crystal structures. *Proceedings of Eurovis 2012 Short Papers* (2012), 31–35. 2
- [Ami] <http://amira.zib.de>. 6
- [Beu94] BEUCHER S.: Watershed, hierarchical segmentation and waterfall algorithm. In *Mathematical morphology and its applications to image processing*. Springer, 1994, pp. 69–76. 2
- [BGG09] BAJAJ C., GILLETTE A., GOSWAMI S.: Topology based selection and curation of level sets. In *Topology-Based Methods in Visualization II*. Springer, 2009, pp. 45–58. 2
- [CB97] COUPRIE M., BERTRAND G.: Topological gray-scale watershed transformation. *Proc. SPIE 3168*, 6 (1997), 136–146. 2, 4
- [CLB11] CORREA C., LINDSTROM P., BREMER P.-T.: Topological spines: A structure-preserving visual representation of scalar fields. *Visualization and Computer Graphics, IEEE Transactions on 17*, 12 (2011), 1842–1851. 2
- [ELZ02] EDELSBRUNNER H., LETSCHER D., ZOMORODIAN A.: Topological persistence and simplification. *Discrete and Computational Geometry* 28, 4 (2002), 511–533. 2, 4
- [GDN*07] GYULASSY A., DUCHAINEAU M., NATARAJAN V., PASCUCCI V., BRINGA E., HIGGINBOTHAM A., HAMANN B.: Topologically clean distance fields. *Visualization and Computer Graphics, IEEE Transactions on 13*, 6 (Nov 2007), 1432–1439. 2
- [GN07] GAVIN A., NIGEL C.: Computer simulation of polymer solar cells. *Modell. Simul. Mater. Sci. Eng.* 15, 2 (Dec. 2007). 2
- [HBP*12] HOMBERG U., BAUM D., PROHASKA S., KALBE U., WITT K.: Automatic extraction and analysis of realistic pore structures from μCT data for pore space characterization of graded soil. In *Proceedings of the 6th International Conference on Scour and Erosion (ICSE-6)* (2012), pp. 345–352. 2, 3
- [JBS06] JONES M., BAERENTZEN J. A., SRAMEK M.: 3D distance fields: A survey of techniques and applications. *Visualization and Computer Graphics, IEEE Transactions on 12*, 4 (2006), 581–599. 2
- [JM10] JONES C., MA K.-L.: Visualizing flow trajectories using locality-based rendering and warped curve plots. *Visualization and Computer Graphics, IEEE Transactions on 16*, 6 (Nov 2010), 1587–1594. 2
- [KFR*11] KRONE M., FALK M., REHM S., PLEISS J., ERTL T.: Interactive exploration of protein cavities. In *Proceedings of the 13th Eurographics / IEEE - VGTC Conference on Visualization* (Aire-la-Ville, Switzerland, Switzerland, 2011), EuroVis'11, Eurographics Association, pp. 673–682. 2
- [KG12] KODALI H., GANAPATHYSUBRAMANIAN B.: Computer simulation of heterogeneous polymer photovoltaic devices. *Modelling Simul. Mater. Sci. Eng.* 20 (2012). 2
- [KGPS13] KRETSCHMER J., GODENSWAGER C., PREIM B., STAMMINGER M.: Interactive patient-specific vascular modeling with sweep surfaces. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2828–2837. 3
- [KLG*13] KUHN A., LINDOW N., GÜNTHER T., WIEBEL A., THEISEL H., HEGE H.-C.: Trajectory density projection for vector field visualization. In *EuroVis - Short Papers 2013* (2013), pp. 31–35. 3
- [LBH11] LINDOW N., BAUM D., HEGE H.-C.: Voronoi-based extraction and visualization of molecular paths. *Visualization and Computer Graphics, IEEE Transactions on 17*, 12 (2011), 2025–2034. 2
- [LCMH09] LAMPE O., CORREA C., MA K.-L., HAUSER H.: Curve-centric volume reformation for comparative visualization. *Visualization and Computer Graphics, IEEE Transactions on 15*, 6 (Nov 2009), 1235–1242. 3
- [LGD*05] LARAMEE R. S., GARTH C., DOLEISCH H., SCHNEIDER J., HAUSER H., HAGEN H.: Visual analysis and exploration of fluid flow in a cooling jacket. In *Proceedings IEEE Visualization 2005* (2005), pp. 623–630. 2
- [PSBM07] PASCUCCI V., SCORZELLI G., BREMER P.-T., MAS-CARENHAS A.: Robust on-line computation of Reeb graphs: Simplicity and speed. *ACM Trans. Graph.* 26, 3 (July 2007). 2
- [Pud98] PUDNEY C.: Distance-ordered homotopic thinning: A skeletonization algorithm for 3D digital images. *Comput. Vis. Image Underst.* 72, 3 (Dec. 1998), 404–413. 2, 4
- [SRS08] SHAW P. E., RUSECKAS A., SAMUEL I. D. W.: Exciton diffusion measurements in poly(3-hexylthiophene). *Advanced Materials* 20, 18 (2008), 3516–3520. 5
- [TvW99] TELEA A., VAN WIJK J. J.: Simplified representation of vector fields. In *Proceedings of the Conference on Visualization '99: Celebrating Ten Years* (1999), VIS '99, IEEE Computer Society Press, pp. 35–42. 3
- [WG12] WODO O., GANAPATHYSUBRAMANIAN B.: Modeling morphology evolution during solvent-based fabrication of organic solar cells. *Computational Materials Science* 55 (April 2012). 1, 3, 6, 7
- [WTCG12a] WODO O., TIRTHAPURA S., CHAUDHARY S., GANAPATHYSUBRAMANIAN B.: Computational characterization of bulk heterojunction nanomorphology. *Journal of Applied Physics* 112, 6 (2012). 5, 7, 8
- [WTCG12b] WODO O., TIRTHAPURA S., CHAUDHARY S., GANAPATHYSUBRAMANIAN B.: A graph-based formulation for computational characterization of bulk heterojunction morphology. *Organic Electronics* 13, 6 (June 2012), 1105–1113. 2, 5, 6, 7