Simulated Annealing Reconstruction and Effective Transport Properties of PEM Fuel Cell Catalyst Layers

by

Pascal Ruzzante

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Mechanical and Mechatronics Engineering

Waterloo, Ontario, Canada, 2023

© Pascal Ruzzante 2023

Author’s Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

A proton exchange membrane fuel cell (PEMFC) is an energy-conversion device that can be used in a wide variety of applications and is considered a promising alternative to gas and diesel engines for vehicular purposes. During operation, a significant source of energy losses stem from the catalyst layer (CL) of the PEMFC, where resistance to transport and reaction mechanisms under high current density can severely reduce the power output of the device. Computational reconstructions of the CL microstructure are therefore very important in order to study the pore-scale phenomena that take place there. Conventional reconstruction techniques typically use either experimental images, which is expensive and time-consuming or use stochastic reconstruction techniques such as simulated annealing, which strive to generate a statistically accurate representation of the structure, and therefore depend on the selection of appropriate statistical correlations to optimize. Most studies to date have exclusively optimized the two-point correlation in simulated annealing despite other studies indicating it does not adequately characterize a unique microstructure by itself.

In this thesis, a hybrid stochastic reconstruction technique is presented, combining experimental images with simulated annealing. Two additional statistical correlations are considered for optimization, namely the lineal path distribution and a 2D analog of the pore-size distribution. The reconstruction technique reconstructs a CL from 2D scanning electron microscope experimental images, combining a sphere-packing approach with sphere-based simulated annealing and simultaneously optimizes the microstructure using the porosity, the two-point correlation, the lineal path distributions and the pore size distributions. Each reconstruction is analyzed by comparing its porosity, surface area, mean chord lengths and pore size distributions with those of the reference image. The study found that the most accurate reconstructions incorporate both the lineal path distribution and the pore size distributions. Pore-scale transport simulations are conducted on the reconstructions to obtain the effective thermal, ionic, and electrical conductivities in addition to the effective diffusivity, in which Knudsen effects are accounted for by using a random walk method. Estimates for the effective ionic and thermal conductivities are in good agreement with those obtained from reference image. However, effective diffusivity and effective electrical conductivity are greatly over and underestimated, respectively. This is attributed to the high sphericity of the sphere-based reconstructions. Narrow connections between sphere contribute to the low electrical conductivity, and the pore phase lacks the long, dead-end pores present in the reference image, which causes the mass diffusivity to be overestimated.

Acknowledgements

I would like to express my deepest appreciation and gratitude to the following individuals who have contributed significantly to the completion of this thesis:

First and foremost, I am very grateful to my supervisor, Dr. Xianguo Li for his support and guidance throughout my MASc degree. His expertise, insightful feedback and dedication have been instrumental in shaping this thesis and expanding my knowledge of fuel cells. I am honored to have had the privilege of working under his supervision.

I am also indebted to the thesis readers, Dr. XiaoYu Wu and Dr. Jeff Gostick, for their time, effort and meticulous review of this thesis. I am grateful for their expertise and for taking the time to provide me with invaluable feedback.

Big thanks are owed to Yixiang ‘Adam’ Zhang for his earlier work on the topic and for bringing me up to speed on what he had done so that he could pass this project on to me.

Furthermore, I would like to thank my wife, Sarah Peverill, for all her loving support and encouragement while I was deep in the writing process. I couldn’t have done it without her help.

Lastly, I wanted to thank my parents for their unwavering support and belief in my abilities, even when I doubted myself.

Table of Contents

Author's Declaration ii
Abstract iii
Acknowledgements iv
List of Figures vii
List of Tables x
List of Symbols xiii
Chapter 1 Introduction 1
1.1 Proton Exchange Membrane Fuel Cells Background 1
1.2 Catalyst Layer Structure 6
1.3 Thesis Objectives 8
1.4 Scope and Outline 9
Chapter 2 Literature Review 10
2.1 Properties of a Porous Structure and Statistical Characterization 11
2.1.1 Porosity and Volume Fraction 11
2.1.2 Two-point Correlation 12
2.1.3 Lineal Path Distribution 13
2.1.4 Chord Length Distribution 14
2.1.5 Pore Size Distribution 15
2.1.6 Specific Surface Area 17
2.2 Reconstruction Methods 18
2.2.1 Direct Reconstruction (Experimental) 19
2.2.2 Fabrication-based Reconstruction 21
2.2.3 Simulated Annealing 23
2.2.4 Deep Learning Reconstruction 25
2.2.5 Ionomer Segmentation 26
2.3 Fuel Cell Transport Simulation 27
2.3.1 Governing Equations 28
2.3.2 Ionic Conductivity 31
2.3.3 Electrical Conductivity 33

2.3.4 Thermal Conductivity 35
2.3.5 Mass Diffusion in Porous Media 37
2.3.6 Other Modeling Frameworks 43
2.4 Summary of Chapter 45
Chapter 3 Methodology 46
3.1 Methodology: Stochastic Reconstruction and Characterization 46
3.1.1 Numerical Representation of Structure 46
3.1.2 Correlations 47
3.1.3 Stochastic Reconstruction 67
3.1.4 Numerical Implementation 81
3.2 Methodology: Transport Simulation and Effective Properties 82
3.2.1 Implementation of Solvers for Thermal, Ionic and Electrical Conductivity 82
3.2.2 Random Walk Model for Knudsen Diffusion 83
3.3 Chapter Summary 94
Chapter 4 Results and Discussion 95
4.1 Reconstruction 95
4.1.1 Part 1: Lineal-Path Distribution Reconstructions 95
4.1.2 Part 2: Analysis of PSD-Based Reconstructions 114
4.2 Transport Simulation of Reconstructions 134
4.2.1 Effective Diffusion and Tortuosity 135
4.2.2 Effective Thermal Conductivity 139
4.2.3 Effective Ionic Conductivity 142
4.2.4 Effective Electrical Conductivity 144
4.2.5 Summary of Conductivity Data 146
4.2.6 Chapter Summary 148
Chapter 5 Conclusions and Recommendations for Future Work 149
5.1 Conclusions 149
5.2 Future Work 151
References 153

List of Figures

Figure 1.1: PEMFC diagram [2]. 2
Figure 1.2: A typical fuel cell polarization curve. 5
Figure 2.1: Examples of the two-point correlation ($\boldsymbol{S2}$) and the lineal path function ($\boldsymbol{L2}$). 13
Figure 2.2: Cross-section of chord-length measurements for an infinitely long line intersecting with the porous solid/void interfaces. 15
Figure 2.3: Methods for reconstruction of porous media. 19
Figure 3.1: Two-point correlation of a random packing of uniform radius spheres. The blue circles indicate the two-point correlation extracted directly from the 3D structure, and the orange circles indicate the two-point correlation extracted from all of the 2D slices of the domain and average together. 49
Figure 3.2: Images of distinct structures with the same porosity that have a similar two-point correlation. 51
Figure 3.3: Two-point correlations for each of the images in Figure 3.2. 51
Figure 3.4: An example image of the lineal distance transform applied to the solid phase in the horizontal direction on an SEM image of the catalyst layer. Brighter colors indicate a longer uninterrupted lineal path distance. Original SEM image obtained from [17], image shown is modified to show lineal path distance transform. 53
Figure 3.5: Depiction of lineal path distribution. 54
Figure 3.6: An example image of the chords filter applied to a CL SEM image. Red and yellow represent the solid and void phases, respectively. The blue lines represent chords drawn in the pore phase (separated from each other for visualization purposes). Chords that would touch the image boundary are artificially shortened and so are removed in order to avoid skewing the data. Original SEM image obtained from [17], image shown is modified to show chords. 55
Figure 3.7: An example illustration of the local thickness filter applied to the pore phase of a CL SEM image. Brighter colours indicate larger pore sizes. Original SEM image obtained from [17], image shown is modified to show local thickness filter. 57
Figure 3.8: Medial Axis Implementation and Progression Depiction. A) Binary Image of Porous Medium (artificially generated). B) Euclidean distance transform on image from (A); C) a skeleton of the pore phase; D) Medial Axis Ridge (skeleton with values from distance transform); E) Medial axis ridge overlaid on image of original pore space; F) Local Thickness image generated by adding spheres to each location on ridge in (E). 60
Figure 3.9: Artificial image generated of a sphere pack of overlapping spheres with a uniform radius of 5 voxels and a porosity of 0.50 in a domain of size 250 x 250 x 250 voxels. A 2D slice is shown of this 3D domain. 63
Figure 3.10: Two-point correlation validation of empirical function using the analytical expression for a pack of overlapping spheres. 64
Figure 3.11: Lineal path distribution validation of numerical implementation using the analytical expression for a pack of overlapping spheres. 65
Figure 3.12: 2D analogue of the pore size distribution calculated by the medial axis method the distance transform method in *porespy*. 66
Figure 3.13: Reference image two-point correlation, and solid and pore phase lineal path distributions. Arrows point to where all three correlations have reached at least 99% of their theoretical total decay. 68
Figure 3.14: 2D depictions of sphere representation. 70

Figure 3.15: Sphere-based simulated annealing algorithm flow control. 76
Figure 3.16: An example of the segmented reconstruction. Note that due to the 3D nature of the reconstruction, solids that appear to be floating and pores that appear blind in the 2D depiction are not actually disconnected from the rest of their respective phases, as they are connected in the 3rd dimension. 81
Figure 3.17: Computation of transition regime effective diffusivity for a domain of overlapping spheres of 25 nm radius at varying porosity, and three different computation methods. 94
Figure 4.1: Distribution of spheres inserted into the reconstructions for each of the reconstructions. CR0 and VR0 are shown in blue and in each plot because they are the starting structure for the reconstructions using Methods 1 and 2. 97
Figure 4.2: Sectioned (30th) slice of each of the reconstructions. A) full reference FIB-SEM image [18]; B, C, D) Reconstructions using a variable radius initial sphere pack; E, F, G) Reconstructions starting from a constant radius sphere pack. Black and orange colors represent the solid and void phases, respectively. 98
Figure 4.3: Reconstruction 3D visualization of the solid phase for VR2. 99
Figure 4.4: Two-point correlation, lineal path distribution void phase and lineal path distribution solid phase for variable radius reconstructions. VR1 is pictured on the left (A, C, E) and VR2 is pictured on the right (B, D, F). 102
Figure 4.5: Two-point correlation, lineal path distribution void phase and lineal path distribution solid phase for constant initial radius reconstructions. CR1 is pictured on the left (A, C, E) and CR2 is pictured on the right (B, D, F). 103
Figure 4.6: Pore Size Distributions for each reconstruction. Reconstructions are shown in green and the reference is shown in purple. The squared error is represented by the dotted black line and the sum of squared errors is shown in the legend. 106
Figure 4.7: Porosities of the reference image and each reconstruction. 107
Figure 4.8: Specific surface area for the reference image and each of the reconstructions, computed using the marching cubes algorithm implemented in porespy [55]. 109
Figure 4.9: Mean pore diameter for the reference and reconstructions. 110
Figure 4.10: Mean particle diameter for the reference and reconstructions. 111
Figure 4.11: Mean chord length for the pore phase of the reference and reconstructions, computed using the chord image filter. 112
Figure 4.12: Mean chord length for the solid phase of the reference and reconstructions, computed using the chord image filter. 113
Figure 4.13: Normalized configuration energy as a function of reconstruction iteration number. The energy is normalized by the initial configuration starting energy for each method. 117
Figure 4.14: Annealing Temperature as a function of iteration number. 118
Figure 4.15: Two-point correlation for each of the reconstructions. The squared error is represented as a black dashed line, and the integral of the squared error is stated in the legend. 121
Figure 4.16: Lineal path distribution in the solid phase for the reconstructions. The squared error is represented as a black dashed line, and the integral of the squared error is stated in the legend. 122
Figure 4.17: Lineal path distribution in the pore phase for the reconstructions. The squared error is represented as a black dashed line, and the integral of the squared error is stated in the legend. 123

Figure 4.18: Pore size distribution of each of the reconstructions. The squared error is represented as a black dashed line, and the integral of the squared error is stated in the legend. 125
Figure 4.19: Particle size distribution for each of the reconstructions. The squared error is represented as a black dashed line, and the integral of the squared error is stated in the legend. 126
Figure 4.20: Sum of squared errors of each correlation for each reconstruction. 127
Figure 4.21: Porosity of reconstructions as compared to the reference. 128
Figure 4.22: Specific Surface Area for reconstructions M0-M5. Sample error is represented by the standard deviation and shown as a black bar. Relative error compared to the reference is shown in red text. . 129
Figure 4.23: Mean pore diameter for the reference and reconstructions. 131
Figure 4.24: Mean Particle Diameter for the reference and reconstructions 132
Figure 4.25: Mean pore phase chord lengths for the reference and reconstructions. 133
Figure 4.26: Mean particle phase chord length for the reference and reconstruction 134
Figure 4.27: Views of the reference image after zeroth order interpolation so that the resolution is 5 nm in all three cartesian directions 135
Figure 4.28: Effective Oxygen Diffusivity for each of the reconstruction methods. Simulations were conducted with a temperature of 298K and pressure of 1 atm. 136
Figure 4.29: Pore tortuosity of reconstruction methods 137
Figure 4.30: Comparison of pathways for artificially generated sphere-based image (left) and a structure similar to the reference FIB-SEM image generated by applying a Gaussian blur to random noise (right). It is evident that the sphere-based image produces a structure with less tortuous pathways. 138
Figure 4.31: Thermal conductivity of reference and reconstructions. Relative error compared to the reference is indicated in red text. Sample variability is represented as error bars and is the standard deviation. 141
Figure 4.32: Ionic Conductivity for the reference and reconstruction methods. 144
Figure 4.33: Electrical Conductivity of reference and reconstructions 145

List of Tables

Table 3.1: Parameters input to an exponentially modified Gaussian distribution used to generate the sphere radii inserted into reconstruction. 72
Table 3.2: Bulk Phase Conductivities 83
Table 4.1: Notation used for Part 1 reconstructions. 96
Table 4.2: Parameters used in the reconstruction of the PEMFC CL 100
Table 4.3: Notations for reconstruction methods for PSD-based reconstruction analysis. 116
Table 4.4: Simulation computation time for each reconstruction method. 119
Table 4.5: Summary of normalized conductivities measured for each reconstruction method, reported with calculated sample standard deviation from three samples of each reconstruction method. Each conductivity is normalized by its bulk value as given in Table 3.2. 147
Table 4.6: Error of each conductivity measurement relative to the reference value (calculated $RE = k - kref/kref$) where RE is the relative error and k is the generalized conductivity measurement. 147

List of Abbreviations

2D	two-dimensional
3D	three-dimensional
BET	Brunauer-Emmett-Teller
CFD	computational fluid dynamics
CL	catalyst layer
CR	constant radius
DNS	direct numerical simulation
EDC	effective diffusivity coefficient
FFT	fast Fourier transform
FIB-SEM	focused ion beam scanning electron microscopy
FVM	finite volume method
GAN	generative adversarial network
GDL	gas diffusion layer
GPU	graphical processing unit
IFFT	inverse fast Fourier transform
jit	just-in-time
LBM	lattice Boltzmann method
M	reconstruction method
MEA	membrane electrode assembly
MIP	mercury intrusion porosimetry
MPL	micro porous layer
MSP	method of standard porosimetry
MW	molecular weight
PaSD	particle size distribution

PEMFC	proton exchange membrane fuel cell
PNM	pore network modeling
PSD	pore size distribution
PuMA	Porous Microstructure Analysis library
RSD	radial size distribution
RW	random walk
SEM	scanning electron microscope
SSA	sphere-based simulated annealing
SSE	sum of squared errors
VR	variable radius
XCT	X-ray computed tomography

List of Symbols

C	concentration ($\mathrm{mol/L}$)
D	mass diffusivity ($\mathrm{m^2/s}$)
E	energy cost; expected value
G_f	Gibbs' free energy (J)
I	indicator function
J	overlap indicator function
K	permeability ($\mathrm{m^2}$)
Kn	Knudsen number
L	Lineal path distribution
N	diffusive molar flux (mol/s); model domain side length (pixels)
N_{MC}	max number of successful trial configurations in one Markov chain
N_{T}	max number of total trial configurations in one Markov chain
P	probability; pressure (atm)
P_c	capillary pressure (atm)
Q	heat (J)
R	sphere radius (constant); ideal gas constant ($\frac{\mathrm{J}}{\mathrm{mol\cdot K}}$)
S	source term
S_2	two-point correlation function
T	temperature (K); fictitious annealing temperature
V	electric potential (V); union volume of two spheres
c_p	specific heat capacity ($\frac{\mathrm{J}}{\mathrm{kg\cdot K}}$)
d	distance (nm)
k	thermal conductivity ($\frac{\mathrm{W}}{\mathrm{cm\cdot K}}$); Boltzmann constant ($\mathrm{J/K}$)
l_c	mean chord length (nm)

q	heat flux ($\mathrm{W/m^2}$)
r	distance between two points; sphere radius (variable)
s	specific surface area ($\mathrm{m^2/m^3}$); water saturation
t	time (s)
$\vec{u}$	gas velocity ($\mathrm{m/s}$)
$\bar{v}$	mean thermal velocity ($\mathrm{m/s}$)
$\boldsymbol{x}$	Cartesian coordinates in vector form
x	Cartesian coordinate in x-direction
y	Cartesian coordinate in y-direction
z	Cartesian coordinate in z-direction

Greek

ε	porosity
ϵ_v	sphere overlap tolerance
θ	contact angle
λ	mean free path (nm); exponential shape parameter; simulated annealing temperature scaling factor
μ	viscosity ($\mathrm{Pa \cdot s}$); distribution mean
ν	sphere overlap ratio
ρ	mass density ($\mathrm{kg/m^3}$); sphere-packing density
σ	surface tension ($\mathrm{N/m}$); distribution standard deviation
σ_{e^-}	electrical conductivity ($\mathrm{S/cm}$)
$\sigma_{\mathrm{H^+}}$	ionic conductivity ($\mathrm{S/cm}$)
τ	tortuosity
ϕ	phase volume fraction

Subscripts

bulk	bulk property
C	carbon
eff	effective
g	gas
ion	ionomer
K	Knudsen
mfp	mean free path
Pt	platinum
ref	reference
spp	steps per path
w	water
i	chemical species; initial

Superscripts

i	phase (solid or pore)

Chapter 1

Introduction

1.1 Proton Exchange Membrane Fuel Cells Background

Continued heavy reliance on fossil fuels as the global primary energy source has had and will have increasingly severe negative consequences for the climate, the environment, and society at large. Proton exchange membrane fuel cell (PEMFC) is a clean, efficient energy conversion device which produces electricity directly, consuming hydrogen and oxygen electrochemically It is an area of intensive research due to their promising potential for use in many practical applications, including vehicular applications [1].

Typically, a commercial PEMFC stack is composed of up to several hundred individual cells, compressed together to create a device that will produce enough power for a given application. Each individual fuel cell (depicted in Figure 1.1) is composed of two outer bipolar plates, which contain flow channels and are responsible for the distribution of reactant gases in addition to the conduction of electricity and dissipation of heat. Between the bipolar plates lies the membrane electrode assembly (MEA), which is composed of a several layers. Starting from the outside, there is a gas diffusion layer (GDL) next to the bipolar plates. This is a highly porous fibrous material used as a backing for the electrodes, which allows the reactant gases in the flow channels to diffuse through to the catalyst layer. The catalyst layer (CL) is a very thin (~10 μm) layer of carbon support particles, catalyst (typically Platinum) and a proton-conducting polymer, referred to as the 'ionomer'. Lastly, separating the anode CL from the cathode CL is the membrane, which is also an ionically conducting polymer. The membrane allows protons to travel across the cell but is electrically insulating such that electrons are forced to move through an external circuit.

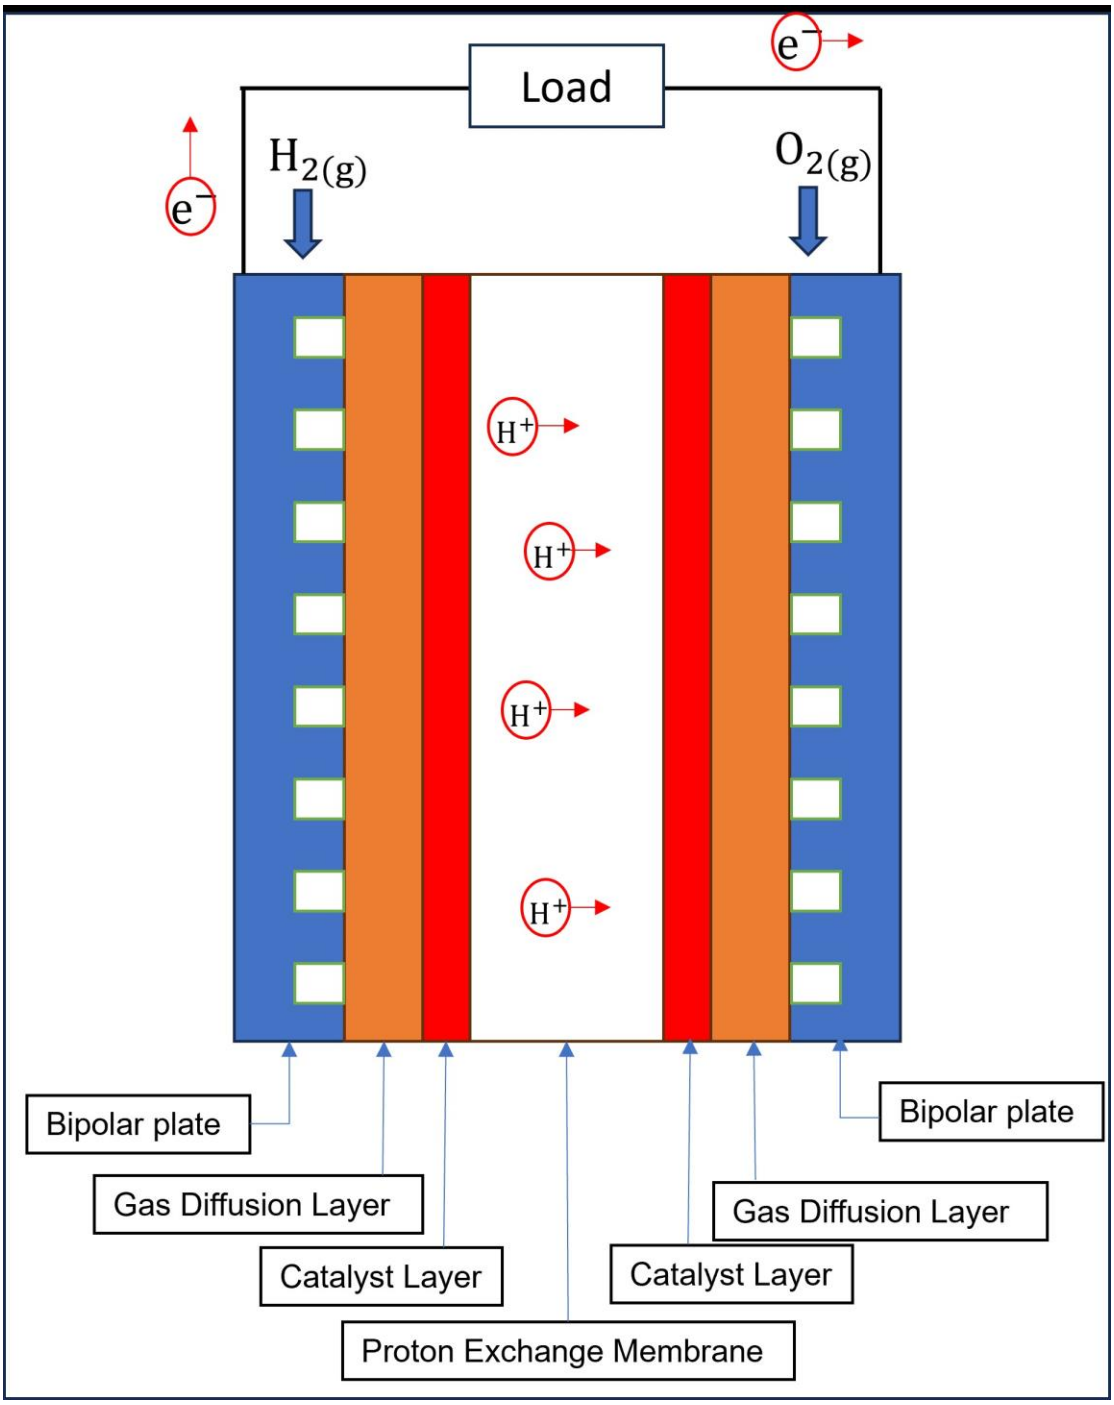


Figure 1.1: PEMFC diagram [2].

On the anode side, hydrogen gas is convected through the flow channels in the bipolar plate, from where it diffuses through the gas diffusion layer and the catalyst layer where the hydrogen splits through a catalyzed reaction into individual protons and electrons according to:

$$H_2 \rightarrow 2H^+ + 2e^- \tag{1.1}$$

The protons are conducted across a thin membrane to the cathode side of the cell, while the electrons are forced to travel through an external circuit (which is how the electrical work is extracted). On the cathode side of the cell, oxygen molecules diffuse through the gas phase of the GDL and CL, where they react at active catalyst sites with the protons and electrons to form water:

$$\frac{1}{2}O_2 + 2e^- + 2H^+ \rightarrow H_2O \tag{1.2}$$

The water formed at the reaction sites in the cathode CL must then diffuse back out through CL, and GDL where it is carried out of the fuel cell by the flow channel.

The overall reaction equation for a fuel cell given at a temperature of 298K and 1 atm pressure is:

$$H_2 + \frac{1}{2}O_2 \rightarrow H_2O_{(l)} + W_e + Q,\ \Delta G_f = -237\ kJ/mol,\ \ Q = -48.6\ kJ/mol \tag{1.3}$$

where ΔG_f is the Gibbs free energy and Q is the heat released by the reaction.

Design of PEMFCs is an area which is defined by compromise and optimization. Three overarching and competing metrics are often referenced: 1) Performance (how effectively it consumes hydrogen and outputs electricity), 2) Durability (how long it lasts before it no longer works) and 3) Cost (how expensive it is to make and operate). While these considerations apply to the entire fuel cell, take the catalyst layer as an example. The typical catalyst used in PEMFCs is Platinum, which is a precious metal and is very expensive. Therefore, in order to reduce costs, some research has been conducted in reducing Platinum particle size, the amount of

Platinum used [3] or using non-precious metals as catalyst instead [4]. However, less Platinum or smaller Platinum particles are more susceptible to degradation or poisoning, so this technique is often reducing cost at the expense of durability. Likewise, using non-precious metal catalysts may be cheaper and may also be durable, but the performance is usually lower than that of PEMFCs with Platinum catalysts. Attempts to improve performance often come at increased costs, and improved durability may reduce performance or increase costs as well. This is all to say that the path forward for research in PEMFC design and manufacturing requires optimization of performance, durability and costs.

Optimization of the PEMFC design also requires understanding of the underlying processes and transport phenomena occurring within the fuel cell. These are best explained using a typical PEMFC polarization curve (graph of voltage vs current density), which is depicted in Figure 1.2. The thermoneutral voltage is the voltage expected if there were no energy losses (reversible and irreversible) and it is independent of the current density (the amount of current drawn from the fuel cell). The reversible potential is the voltage expected in the ideal case where entropic (reversible) losses are considered, but no irreversible losses are present. It is likewise independent of the current density. In the real case, there are of course irreversible losses that must be considered, and these losses increase as a higher current density is drawn from the cell. The PEMFC has a theoretical reversible potential of 1.23V, though in practice this is reduced to 0.6-0.8 V due to a number of sources of irreversible energy losses, sometimes referred to as overpotentials. These losses can be categorized into three primary regions: the activation overpotential, Ohmic overpotential and concentration overpotential, which are also depicted in Figure 1.2. At low current densities, the activation polarization is typically the only significant source of overpotential and is due to energy ‘spent’ to drive up the slow rate of the electrochemical reactions. Ohmic polarization occurs at moderate current densities and is primarily due to losses occurring from the resistance to proton transfer in the membrane and ionomer phase of the catalyst layer. Concentration

polarization occurs at high current densities and is what defines the limiting current density (the maximum current density that can be drawn from the fuel cell before the voltage is reduced to 0). This overpotential primarily arises from the limiting mass transfer of oxygen to, and the accumulation of water at, the reaction sites. Water accumulation at the reaction sites can block the oxygen transfer, further increasing concentration polarization. The catalyst layer pore structure is therefore a critical parameter in determining the effectiveness of oxygen and water transport; a higher porosity will enable better transport, as will ensuring the pore space is well-connected and less tortuous.

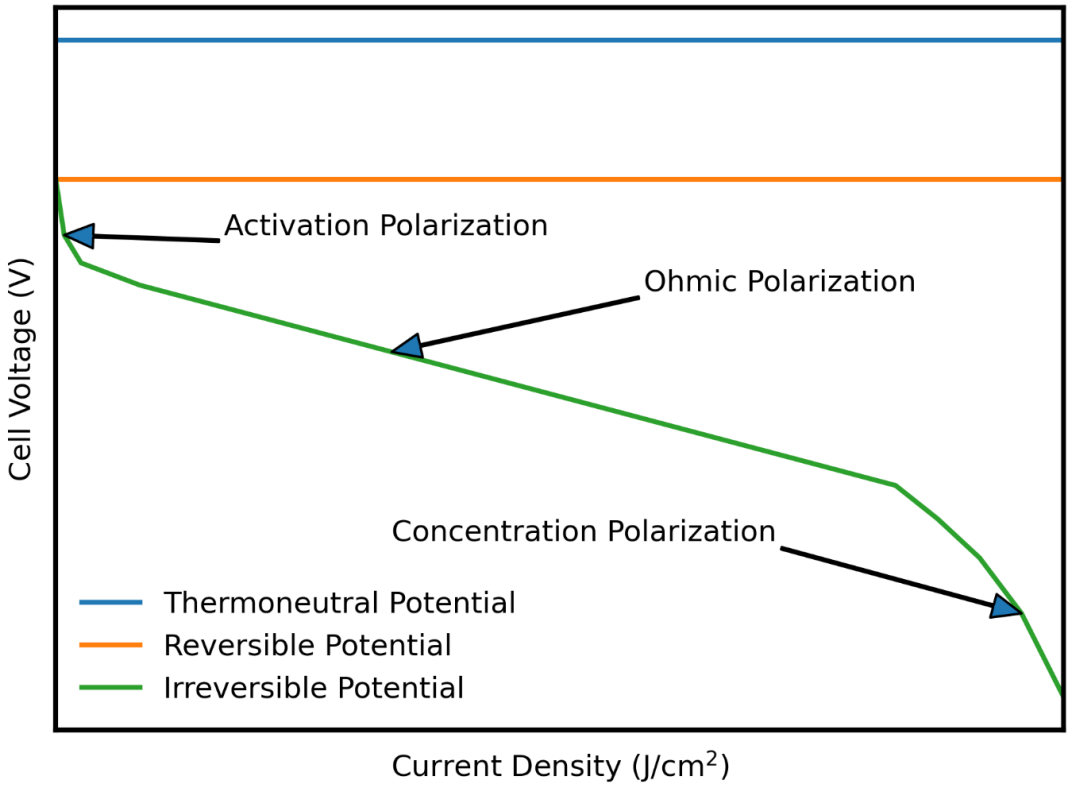


Figure 1.2: A typical fuel cell polarization curve.

The degree of Ohmic polarization and concentration polarization is strongly affected by the catalyst layer microstructure [6]. In particular, the limited diffusion of oxygen through the cathode catalyst layer is one of the main causes for concentration overpotential due to the larger size of the oxygen molecule (compared to water or hydrogen) and due to the small, tortuous pathways available in the pore phase of the catalyst layer structure. The effective diffusivity in the catalyst layer can typically be improved by increasing the porosity, but this then reduces the carbon and ionomer fractions, which causes a reduction in the effective ionic and electronic conductivities, leading to an increase in Ohmic polarization. It is therefore once again evident that some sort of optimized balance must be found between the volume fractions and structures of the pore, ionomer, and carbon phases, while ensuring that each phase is sufficiently interconnected to enable efficient transport for heat, mass, protons and electrons. Simulation is therefore required to find and test optimal structures for the catalyst layer which maximize the effective thermal, mass, ionic and electrical conductivities. In order to conduct simulations, the catalyst layer microstructure must first be digitally reconstructed, something that has been the subject of considerable study [7–11].

1.2 Catalyst Layer Structure

The catalyst layer (CL) of a PEMFC is a thin layer of the electrode which interfaces with the cell membrane. It contains the electrochemical reaction sites, and it must provide a pathway for the delivery of reactants and removal of products through its porous network. A traditional PEMFC CL is in the range of a few microns to about 50 μm thick and is made up of four principal components: 1) the carbon support ‘backbone’, 2) the platinum catalyst particles, supported on the carbon; 3) an ionomer (typically Nafion™), which is an ionically conducting polymer and 4) a void phase throughout the structure which enables transport of gaseous and liquid species. Despite the small thickness of the CL, transport and reaction resistances in the CL of a PEMFC are

dominant contributing factors to cell overpotential and efficiency loss [2]. The following mechanisms must all occur simultaneously in the CL for effective operation of the PEMFC:

A) Oxygen transport to reaction sites via diffusion through the tortuous pore microstructure [5];
B) Proton transport from membrane to reaction sites through the ionomer phase [5];
C) Electronic conduction in the solid carbon phase [5];
D) Water is produced by the reaction and is removed from the CL by capillary action, or else water flooding phenomenon occurs which severely degrades fuel cell output [12,13];
E) Heat transfer away from CL by conduction due to exothermic nature of reaction [14].

These processes are all intricately coupled and heavily dependent on the porous catalyst layer microstructure. Zhao et al [15] used the method of standard porosimetry to confirm that the microstructure of the catalyst layer has a significant effect on the overall transport in the fuel cell electrode. They found that 95%+ of the pore surface area is characterized by pores smaller than 100nm. Additionally, they found that the pore specific surface area, mean pore size and pore size distributions of the electrode were drastically impacted by the structure of the catalyst layer, as well as by the concentration of platinum loading. This demonstrates the need to better understand the microstructure of the PEMFC catalyst layers.

In the literature, various reconstruction methods have been developed and implemented such as direct reconstruction from experimental images [9,16–18], stochastic techniques such as fabrication-based [8,11,19–22] or simulated annealing [7,23–27], or image-based machine learning [28–30], and each has some advantages and disadvantages compared to the others. Direct reconstruction from experimental images is theoretically the most accurate as it is based on objective reality, but imaging and subsequent processing of the images is expensive and time-consuming, making it an impractical method to quickly generate and test

many different structures. Reconstruction through machine learning is able to generate many structures from one dataset once the network is trained but requires a large image dataset to begin with for model training. Stochastic fabrication-based reconstruction is generally quite fast but is dependent on knowing and simulating the process by which the catalyst layer is formed at the particle-scale, which is not necessarily well-known or established. Simulated annealing is a stochastic process by which a structure is optimized to match a set of reference statistical correlations, theoretically obtaining an accurate structure, but this accuracy is dependent on the correlations chosen to represent the target structure and how well those correlations encode the structural information. Simulated annealing is also considerably slower than the fabrication-based process.

1.3 Thesis Objectives

This thesis aims to improve upon existing reconstruction techniques to obtain more accurate structures of the PEMFC catalyst layer. This approach is accomplished by combining stochastic methods such as simulated annealing and sphere-packing with data extracted from experimental images to form a 'hybrid' reconstruction algorithm. In the current literature, simulated annealing is conventionally only used to optimize the two-point correlation as it is relatively simple to implement and quick to compute. The work in this thesis expands the use of simulated annealing to optimize the lineal path distribution and a 2D analog of the pore size distribution. In both methods, the porosity is also optimized, and a variable sphere size distribution is used to generate the spheres inserted into the domain. Spheres change size dynamically during the simulated annealing to allow the algorithm to optimize the sphere size distribution. The reconstruction accuracy is determined from metrics such as the porosity, specific surface area, linear characteristic lengths, and radial characteristic diameters.

The second objective of the thesis is to implement pore-scale transport simulations in order to solve for the effective transport properties of each reconstruction. Finite-volume simulations are used to determine thermal,

ionic and electronic conductivities of each reconstruction. A random walk method is implemented to determine the effective gas diffusivity in the catalyst layer which accounts for the effects of Knudsen diffusion.

1.4 Scope and Outline

Chapter 2 first presents a literature review of the structural correlations and metrics used to characterize porous media in section 2.1. It then proceeds to review the different methods of catalyst layer reconstruction in section 2.2. Section 2.3 presents a literature review of pore-scale transport simulation in reconstructed catalyst layers and the calculation of the effective transport properties. Chapter 3 presents the methodologies used in reconstruction and transport modeling of the catalyst layers. Chapter 4 presents the results and discussion of the study, where the section on reconstruction is split into a Part 1, focused on using the lineal path distribution, and a Part 2, where simulated annealing is conducted by optimizing the pore size distribution. Chapter 4 then also presents the results of the transport simulations on the reconstructions from Part 2. Chapter 5 presents the conclusions of the work and provides some suggestions for future direction.

Chapter 2

Literature Review

The traditional approaches to fuel cell transport modeling are the continuum or volume-averaged methods (VAMs), in which the catalyst layer is treated either as an infinitely thin interface, as a homogenous layer with finite distance, or as a layer with large micron-scale agglomerates[31]. VAMs are able to efficiently perform macroscopic-scale simulations but are not able to resolve pore-scale processes which can have significant impact on performance. In order to implement a volume-averaged approach, effective transport properties are required, which incorporate pore-scale phenomena as a sub-scale (or sub-grid) aggregate suitable for VAM-based analysis, modeling and simulation. Traditionally, these are obtained via empirical correlation equations such as the Bruggeman correlation [32,33], which was developed for pores formed by uniform solid spheres.

With the advent of increasingly better imaging techniques, more information on the microstructure of the CL is available. It follows therefore that models have since been developed that incorporate the geometric features/details of this microstructure, resolved down to approximately the nanometer length scale and are referred to collectively as pore-scale models (PSM). These are able to predict effective transport properties more accurately, but modern computing power is unable to efficiently simulate transport at pore-scale resolution over a macroscopic domain-size most relevant to practical fuel cell engineering design.

A promising new direction of research has emerged known as multi-scale modeling, in which pore-scale models are coupled to macroscopic full cell or stack models. Zhang et al [34] developed a multi-scale model which couples a pore-scale agglomerate model that incorporates electrochemical reaction and oxygen transport with a cell-scale model that simulates non-isothermal conditions and multi-phase flow. Similarly, Ingham et al [35] developed a multi-scale model using the agglomerate approach as their pore-scale model which

demonstrated that small changes in the shape of the agglomerate can have significant effects on the current density. They also reported a high sensitivity of the current density to agglomerate specific surface area. However, at the most basic, it is possible to extract effective transport properties from a pore-scale model, which are then input into a separate continuum-scale simulation.

The multi-scale modeling techniques are promising, but the agglomerate approach only estimates the actual pore structure based on assuming a geometry for the agglomerate. Ideally, the pore-scale modelling would use the actual pore microstructure of the CL. Digital reconstruction of the CL provides a means to do exactly this – the microstructure can be characterized, and the reconstructed pore geometry can be used to model the transport phenomena and reaction processes in order to obtain effective transport properties [36].

2.1 Properties of a Porous Structure and Statistical Characterization

In order to better understand the different reconstruction methods used in catalyst layer reconstruction, it is first useful to understand the different ways in which the catalyst layer is characterized. This section introduces various metrics and statistical correlations used to characterize porous structures such as the PEMFC catalyst layer.

2.1.1 Porosity and Volume Fraction

Porosity is a special name given to the volume fraction of void space in a porous medium. It is defined as the volume of void space per unit volume of porous medium [37]. More generally, the volume fraction of any phase in a porous medium can be defined as the volume of that phase per unit volume of the porous medium. In the catalyst layer, a volume fraction can be calculated for each of the carbon, ionomer and platinum phases, in addition to the porosity.

Porosity is one of the most important metrics for the transport of gases within the catalyst layer of the fuel cell. In fact, the effective diffusivity D_{eff} is often estimated directly from only the porosity, raised to the power n:

$$D_{\text{eff}} = \varepsilon^n \tag{2.1}$$

where reported values for n vary from 1.5 (such as in the Bruggeman correlation [38] for a packing of spheres) all the way up to 6 in some reported reconstructions [39]. Regardless of the exact value of n, it is evident that the porosity has an oversized effect on the effective oxygen diffusivity.

PEMFC catalyst layer porosity usually ranges from 0.2 up to 0.5, while typical values are around 0.35[40].

2.1.2 Two-point Correlation

The two-point correlation S_2 for an isotropic material can be described as the probability that a line of magnitude $|\boldsymbol{r}|$ has endpoints in the phase of interest when randomly tossed into the porous structure. Mathematically, this can be defined as

$$S_2(\boldsymbol{r}) = P(\boldsymbol{x} \text{ in phase } i \text{ and } \boldsymbol{x} + \boldsymbol{r} \text{ in phase } i) \tag{2.2}$$

where $\boldsymbol{x}$ is Cartesian coordinate of one endpoint of the line segment and $\boldsymbol{x} + \boldsymbol{r}$ is the other endpoint**.**

The two-point correlation contains information on short-range periodicity and interfacial surface area but is unable to detect the presence of different characteristic length scales or the degree of phase connectivity, and therefore struggles to account for long-range structural information. One case of the two-point correlation is depicted in Figure 2.1 in the green color. Note that the two-point correlation only looks at the endpoints of the line segment; the line itself can cross phase boundaries.

For isotropic media, the two-point correlation decreases from [41]

$$S_2{}^i(r=0) = \phi^i \quad \text{to} \quad S_2{}^i(r \to \infty) = (\phi^i)^2 \tag{2.3}$$

where ϕ is the volume fraction of phase i.

For a binary medium of solid and pore phases the two-point correlations of pore phase and solid phase are scalar multiples of each other and are therefore linearly dependent [41,42].

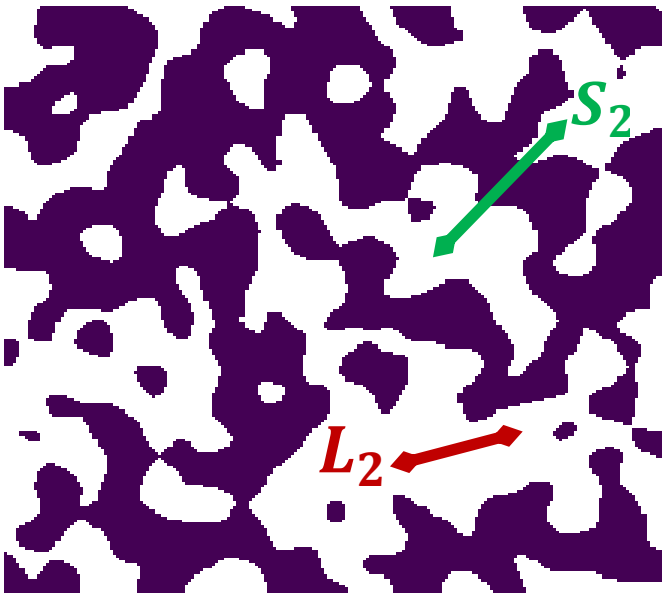


Figure 2.1: Examples of the two-point correlation (S_2) and the lineal path function (L_2).

2.1.3 Lineal Path Distribution

The lineal path distribution function $L^i(\boldsymbol{r})$ for phase i is similar to the two-point correlation, except that it represents the probability that *all* points on a line segment of magnitude $|\boldsymbol{r}|$ between coordinates $\overrightarrow{\boldsymbol{x_1}}$ and $\overrightarrow{\boldsymbol{x_2}}$ are within the phase i [41,43,44]. An example of a case of the lineal path distribution is depicted in Figure 2.1 in the red colour noting that the line does not pass through a phase interface. The lineal path distribution in the void phase is denoted L^v and in the solid phase is denoted L^s. This correlation contains coarse information on

the long-range linear structure and connectedness of the structure. Unlike the two-point correlation, L^v and L^s are linearly independent[43,44], therefore both should be included in reconstruction efforts using simulated annealing. Čapek et al. [45,46] showed that the use of the solid phase lineal path distribution in simulated annealing for structures with porosity $\phi^v < 0.5$ helps to maintain the connectivity of the pore phase. This is because inclusion of the lineal path function tends to cause clustering of the phase of interest, which may cause isolated clusters of the minority phase to form. The lineal path distribution varies from

$$L^i(r \to 0) = \phi^i, \qquad L^i(r \to \infty) = 0 \tag{2.4}$$

2.1.4 Chord Length Distribution

The chord length distribution is defined by Lu and Torquato [47] as the intersections of lines with the interfaces in a porous medium (see Figure 2.2) where the chord is the line segment from a point on one interface to a point on the next interface.

The mean chord length is the first moment of the chord length distribution (sometimes also referred to as the mean intercept length). This is a useful measure of the linear characteristic length of the model which arises in various transport models such as Knudsen diffusion [48] (introduced later in section 2.3.5).

The mean chord length l_c for phase i can be estimated using

$$l_c{}^{(i)} = \frac{-\phi^i}{\left.\frac{dS_2}{dr}\right|_{r=0}} \tag{2.5}$$

where $\left.\frac{dS_2}{dr}\right|_{r=0}$ is the slope of the two-point correlation at $r = 0$ and ϕ^i is the volume fraction of phase i [41]. This is the same definition sometimes used to define the experimentally determined mean pore size such as in [49].

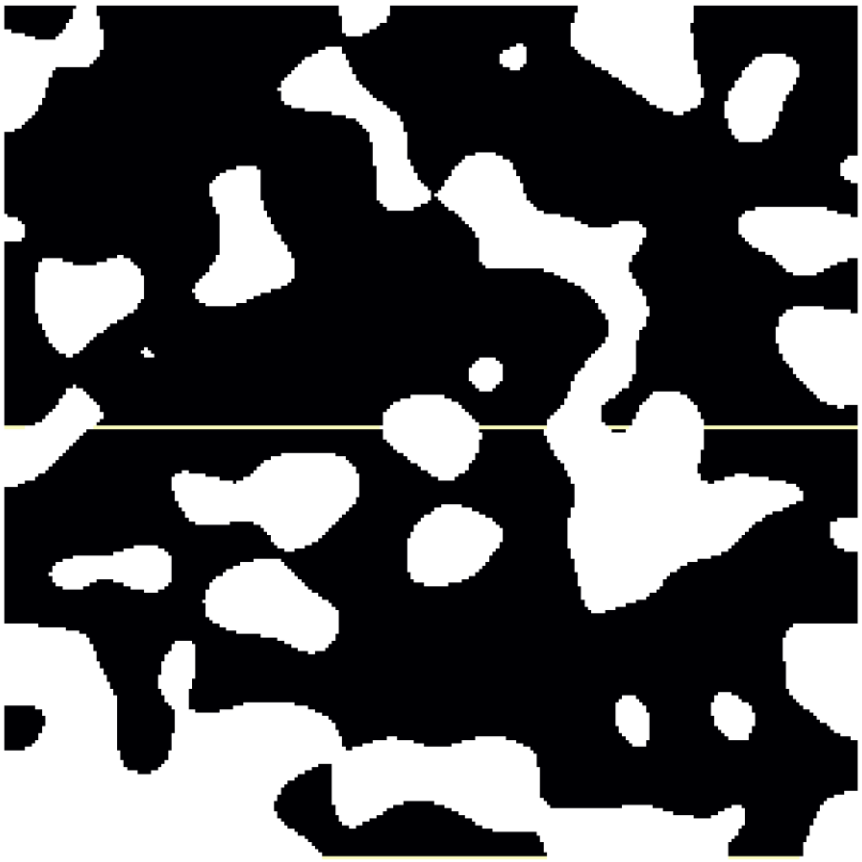


Figure 2.2: Cross-section of chord-length measurements for an infinitely long line intersecting with the porous solid/void interfaces.

2.1.5 Pore Size Distribution

In a PEMFC catalyst layer the pores will not necessarily be all of uniform size and shape and will vary over a given range. The pore size distribution is thus defined as the relative abundance of pore sizes in the porous medium [50].

Experimentally, there are various methods to measure the pore size distribution of a material. Perhaps the most common method is known as mercury intrusion porosimetry (MIP)[51,52] . In MIP, the sample is completely evacuated then immersed in mercury which is highly non-wetting on most materials. The pressure is slowly increased, forcing the mercury into smaller and smaller pores. The volume of pores invaded at each pressure is then related to the pore size by the Washburn equation (Eq. (2.6)) from which it is possible to obtain the pore size distribution [53].

$$P_c = -\frac{2\sigma\cos\theta}{r} \tag{2.6}$$

where σ is the surface tension, θ is the contact angle of the invading fluid (mercury) and r is the pore radius corresponding the pores invaded at the applied capillary pressure P_c.

A second method of experimentally obtaining the pore size distribution is through the method of standard porosimetry (MSP) in which a porous sample is placed between two standard samples, for which the pore size distribution is known. In this method, the sample and standards are initially immersed in a highly wetting liquid such as octane. The samples are slowly dried and weighed intermittently, allowing time in between each weighing to allow the octane to reach a capillary equilibrium between the sample and standards. A relation between pore volume and radius can then be established by comparing to the capillary pressure curves of the standard, from which the pore size distribution can be determined [15,54].

The pore size distribution can also be estimated numerically from images of the porous medium. This typically is done by a method in which spheres of varying radius are fitted into the image and then obtaining a histogram of the sizes. For example, the ImageJ library in Java has a Local Thickness plugin for their graphical user interface (GUI) which fits spheres into a 3D image from which the pore size distribution can be calculated. Likewise, the *porespy* [55] library in python has a local thickness method which inserts spheres

into the image, but this method can also be applied in an access-limited way in which larger pores that are only connected to the surface via smaller pores will only be filled once the smaller pores are invaded. This would more closely resemble the experimental results from a test such as MIP.

In the PEMFC catalyst layer, the pore size distribution is often considered to be bimodal, where a peak exists in the 6-20 nm diameter range, due to pores inside the carbon agglomerates. A second peak is typically in the 20-100 nm range, representing pore spaces between agglomerates [56]. In the PEMFC catalyst layer, it is expected that 95+% of the pores will be less than 100 nm in diameter [57], and the mean pore size will be in the 50-60nm range [18,27].

2.1.6 Specific Surface Area

The specific surface area s of a porous medium is defined as the average interface area between solid and pore phases per unit volume in units of inverse length [58]. Porous media such as the PEMFC catalyst layer can have very high surface areas. Generally, for a given porosity, if the pores are smaller in size on average, the surface area will be higher.

In the PEMFC catalyst layer, more surface area yields more locations for catalyst sites which theoretically increases the reaction rate. However, if an increase in surface area means the average pore size is reduced, then the transport of reactants in the pore space will be hindered, further demonstrating the need to find an optimal structure.

The specific surface area can be estimated in units of inverse length using the slope of the two-point correlation [41] at $r = 0$. In three dimensions,

$$s = -4 * \frac{dS_2}{dr}\Big|_{r=0} \tag{2.7}$$

Experimentally, the most common method of obtaining the specific surface area is the Brunauer-Emmett-Teller (BET) method [59], though it can also be obtained through MSP and MIP as explained in section 2.1.5 on the pore size distribution. In BET, nitrogen gas at 77K is released into an evacuated tube containing the porous sample. Some nitrogen gas will condense onto the pore surfaces, initially forming a monolayer of nitrogen molecules covering the pore surface. The condensation results in a pressure drop from the missing gas, from which it is possible to determine the number of nitrogen molecules that condensed. The area occupied by one nitrogen molecule is known, so it is then possible to calculate the specific surface area of the porous sample occupied by the monolayer.

Numerically, the surface area can be estimated using a marching cubes algorithm first developed by Lorensen and Cline [60], and improved upon by Lewiner et al [61]to ensure watertight mesh topology. This method works by generating a computational mesh from a 3D voxel grid and then calculating the sum of all the surface areas of the mesh.

A typical specific surface area of a PEMFC catalyst layer is reported by Zhao et al [62] to be $8652 \pm 801\ \mathrm{cm^2 cm^{-2}}$ found by MSP and where the units are reported as $\mathrm{cm^2}$ surface area per $\mathrm{cm^2}$ of sample cross-sectional area.

2.2 Reconstruction Methods

There are a variety of methods that have been employed in the reconstruction of the PEMFC CL, some of which are shown in Figure 2.3. In 2016 Shojaeefard et al [10] published a comprehensive review of reconstruction methods that have been employed in PEMFC CL reconstruction for pore-scale simulation;

however, a number of new articles in this field have been published since then. This section summarizes a selection of the reconstruction methods employed in the literature, both before and after 2016 and is organized as follows: Direct reconstruction using experimental imaging techniques is described in section 2.2.1; stochastic fabrication-based (also called process-based) reconstruction methods are described in section 2.2.2; optimization-based simulated annealing methods are described in section 2.2.3; finally a brief mention and description of a relatively new reconstruction method using machine learning and generative adversarial networks is made in section 2.2.4. While each of these methods is different, they all generally result in a binary solid/pore structure (with the possible exception of the machine learning method). In order to conduct effective pore-scale simulations on a CL structure, it is necessary to include the ionomer in the structure. Section 2.2.5 describes various ionomer segmentation processes that have been undertaken in the literature.

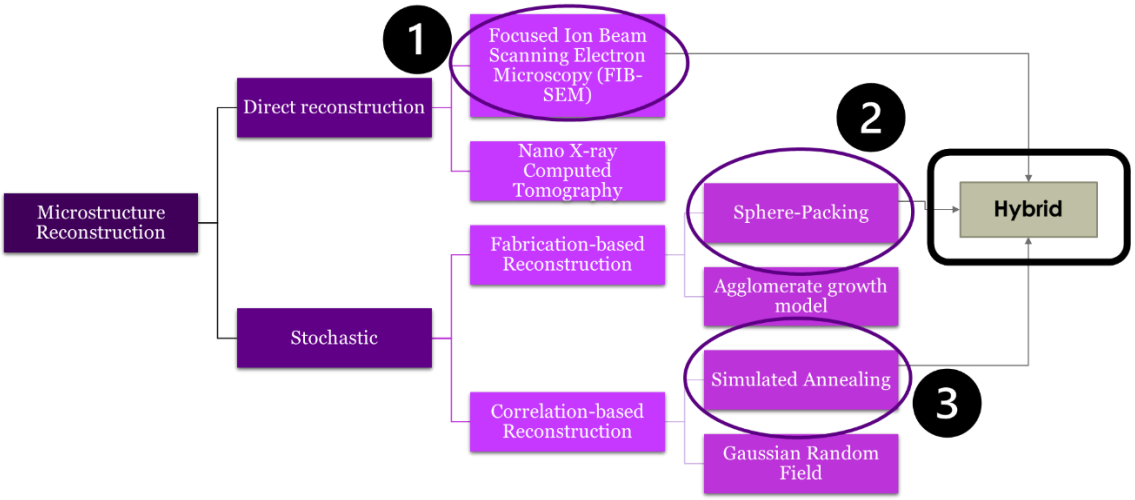


Figure 2.3: Methods for reconstruction of porous media.

2.2.1 Direct Reconstruction (Experimental)

Direct reconstruction from experimental imaging of PEMFC CLs has been attempted various times in the literature, both by focused ion beam scanning electron microscopy (FIB-SEM) [9,17,18,63–68] and by x-ray

computed tomography (XCT) [69]. XCT is a useful and non-destructive technique for reconstruction of porous media, but to best of the author's knowledge, it is limited in resolution to a minimum of approximately 50nm as in [69]. This limitation makes it unsuitable for CL reconstruction, where the majority of the pore volume consists of pores ranging from < 5nm up to 100 nm [15] and which therefore requires an imaging resolution of less than 10 nm to resolve enough of the structure to be meaningful. On the other hand, XCT *is* an effective method of reconstructing porous media with slightly longer length scales, such as in the micro-porous layer (MPL) or gas diffusion layer (GDL) in a PEMFC as published in [70] with a voxel resolution of 650 nm. In [17,18,63] FIB-SEM was used to directly reconstruct a subsection of the PEMFC CL from 113 stacked 2D images with 2.5 nm/pixel resolution in the imaging plane, but only 30 nm/pixel resolution in the through-plane (z) direction, generating a 3D image of dimensions 4127.5 nm × 1752.5 nm × 3390 nm in the x, y and z directions, respectively. The low resolution in the z-direction is on the order of average pore size [71], hence significant three-dimensional structural information is lost. Additionally, this method is very time-consuming, requiring manual, or semi-manual segmentation, filtering and thresholding of the images, which can take up to 3 hours per image slice [63]. In 2015, Vierrath et al. [9] used FIB-SEM alongside atomic layer deposition sample preparation to reconstruct the CL of a PEMFC and achieved a resolution of 1.8 nm in the imaging plane and 5.5 nm in the through plane direction, perhaps one of the highest resolution experimental images of the PEMFC CL to date. Ghosh et al [68] conducted a FIB-SEM reconstruction of the CL but prepared the CL for the FIB-SEM process by filling the pore space with a metalliferous epoxy polymer in order to be able to better differentiate pore foreground pore space from background material phase in the resulting images.

In 2016, Jung et al. [72] reconstructed a dual layered PEMFC CL from FIB-SEM images, achieving an image resolution of 1.5 nm in the imaging plane and 15 nm in the through-plane direction. They confirmed that FIB-SEM was able to accurately predict the pore size distribution of the CL as compared to that obtained

experimentally from mercury intrusion porosimetry. Hutzenlaub et al [66] reconstructed the cathode catalyst layer of a proton exchange membrane fuel cell using FIB-SEM images at a resolution of 3.4nm x 3.4nm x 13.6nm. They used this reconstruction determine the Knudsen diffusivity with a random walk method.

Klingele et al [67] simulated the FIB-SEM imaging process to determine the typical errors caused by the large cutting distance usually inherent in the FIB-SEM process. They accomplished this by generating a reconstruction via a stochastic sphere-packing process (described in 2.2.2.1) and then removing slices from the model to simulate the cutting process. Slices that are kept are repeated, in a typical FIB-SEM processing fashion. They found that significant errors up to 61% due to the cutting distance are incurred when computing effective conductivity properties in the solid phases, while only 4% errors were found in calculating continuum gas phase diffusivity. Knudsen diffusivity calculations showed a larger error than the continuum calculation. Additionally, the cutting process tended to shift the pore size and grain size distributions to lower values.

Based on the reviewed works, it is possible to conclude that while FIB-SEM has been shown to be accurate, it is not a very practical method of reconstructing CLs for design and testing purposes because of the high equipment costs, time-consuming image segmentation process, and destructive nature of the milling process.

2.2.2 Fabrication-based Reconstruction

2.2.2.1 Sphere-packing

Reconstruction has also been attempted using fabrication-based stochastic methods in which parameters such as carbon particle size, ionomer loading, agglomerate size and sample porosity are incorporated into the model. Lange et al [8,19,20] published a series of papers using a sphere-packing approach, in which carbon spheres are randomly inserted into the computational domain until the desired porosity is reached. In [19], the model is extended to consider non-isothermal conditions, water vapor concentration, thermal conduction, Ohmic

heating, and activation losses. Effective thermal, protonic, and electronic conductivities are computed in addition to diffusion coefficients for water vapor and oxygen. When compared to experimental data obtained by GM researchers [73], it is evident that the model greatly overestimates experimental oxygen diffusivity. The authors hypothesize that this could be because the model does not account for liquid water, ionomer swelling, or non-uniform ionomer distribution within the CL.

While this was a comprehensive work using the sphere-packing method, it is evident that the sphere-packing method alone struggles to accurately predict transport in real PEMFC CLs. Berson et al [74] used a gravity sphere packing method where hard (non-overlapping) spheres are 'dropped' into the domain and allowed to roll over other spheres until they either contact the bottom or a stable position. Sabharwal et al [75] reconstructed a catalyst layer using a fully penetrable sphere packing approach adjusting the sphere radius distribution so that the chord length, two-point correlation and pore size distributions match those from a FIB-SEM image. The fully penetrable sphere approach enables the formation of larger clusters and therefore also larger pore sizes. They found that a sphere radius of 40nm gave the best results.

2.2.2.2 Agglomerate-growth

The agglomerate growth reconstruction approach (not to be confused with the agglomerate simulation models) [11,76–79] has also been used to reconstruct PEMFC CLs. In this model, nucleation seeds are dispersed throughout the computational domain initially, and then carbon, platinum and ionomer spheres or voxels are grown from the original nucleation sites until their respective volume fractions are achieved. Since the total volume fraction of each component is fixed, the size of the agglomerates is determined by the number of seeds initially distributed. There exists a trade-off here: A higher number of seeds results in smaller agglomerates, which also indicates a smaller average pore size which will negatively impact oxygen diffusion. On the other

hand, too few seeds, and the agglomerates become very large, which could occur due to poor mixing of components in the fabrication process.

Both sphere-packing and agglomerate – growth fabrication models have achieved some success in qualitatively replicating the porous structure and transport properties of the PEMFC CL but fall short in quantitative analysis such as statistical characterization or calculation of effective transport coefficients.

2.2.3 Simulated Annealing

Simulated annealing is a stochastic approach which has been successfully employed in porous media reconstruction. It is a technique formalized by Yeong and Torquato [80] and adapted from Ingber [81] which works by casting reconstruction as an optimization problem such that a function quantifying the error between the reconstruction and a reference is minimized. The generalized procedure begins by extracting the porosity and statistical correlation(s) from a reference image or trial function. A computational domain of randomly distributed pore and solid voxels with the same porosity as the reference image is then generated and the same statistical correlations are computed on the generated domain. A perturbation is induced in the reconstruction configuration, and the correlations are recomputed. If the relative error is reduced, the configuration is kept. This cycle repeats iteratively until a pre-specified convergence criterium is reached.

2.2.3.1 Conventional Pixel-based Simulated Annealing

In conventional simulated annealing, the perturbation induced at each time step is the swapping of two randomly selected voxels of opposite phase. In theory, this strategy can achieve highly accurate reconstructions because of the enormous quantity of possible configurations. However, it also means that convergence can take an extraordinarily long time as the structure approaches its final structure because the probability of a perturbation resulting in an energy-decreasing move is very small and because each time step only causes a

small change in energy. Pant et al [24,26,82] were able to improve on this method by introducing a different phase neighbour pixel selection process whereby the algorithm preferentially selects voxels that exist on the solid/pore interface for swapping. However, even their optimized method takes 30-50 hours to reconstruct a 300^3 voxel image [25]. While this was able to significantly reduce the convergence time, it is still a computationally expensive process and relies solely on the ability of the chosen statistical function(s) to completely describe the porous microstructure.

2.2.3.2 Sphere-based simulated annealing

Hybrid stochastic reconstruction is a two-phased approach which combines fabrication-based methods with simulated annealing. In doing so, the reconstruction time can be drastically reduced as compared to simulated annealing alone because the fabrication-based methods such as sphere-packing are generally fast. Hybrid stochastic reconstruction using sphere-based simulated annealing (SSA) was first proposed by Kim and Pitsch [7]. Their model was generated by sphere-packing with a uniform sphere size until the desired porosity was met, followed by a simulated annealing process in which the individual spheres were moved within the domain in order to optimize the two-point correlation of the reconstruction configuration to a trial two-point correlation function, defined analytically. Their model was validated by comparing the pore size distribution (PSD) from the reconstructed catalyst layer to that obtained from experimental mercury intrusion porosimetry.

Wu and Jiang [23] extended the approach pioneered by Kim and Pitsch [7] by implementing SSA with a non-uniform sphere distribution. They reconstructed a CL in a computational domain 200 x 200 x 200 nm with a resolution of 2nm per voxel using carbon spheres ranging in size from 10 nm to 30 nm with a mean of 20 nm. They found that this method resulted in a reconstructed microstructure that very closely resembled that of the real CL. They also implemented an LBM transport simulation and determined, similarly to previous studies, that the Bruggeman correlation significantly underestimates the effect of tortuosity in gas transport and

overestimates the oxygen diffusivity in the CL. Despite the findings, a 200 nm cubic domain is very small and unlikely to be a representative elementary volume for practical macroscopic level of analysis.

Kang et al [27] extended the initial work of Kim and Pitsch [7] to a non-uniform carbon particle distribution and introduced a variety of non-uniform ionomer distributions in an attempt to make the model more representative of real CLs. Using the statistical trial function approach, they were able to control the size of carbon agglomerates in the reconstruction by varying a clustering parameter in the function. An interesting result of this is that for the same ionomer, carbon and pore volume fractions, a larger agglomerate size results in thicker ionomer layers, and vice versa. While previous studies have noted that increased agglomerate size generally results in better oxygen diffusion in the gas phase due to the larger pore sizes, the increase in ionomer thickness counteracts this phenomenon by increasing the distance through which it must diffuse in the ionomer phase. On the other hand, an increased agglomerate size and increased ionomer thickness reduces the tortuosity for proton transport, thus improving proton transfer rate.

More recently, Zhang [83] extended the SSA model by incorporating both the two-point correlation and the lineal path function into the reconstruction algorithm, postulating that including the lineal path function would result in a more accurate microstructure because the lineal path function exhibits longer-range structural information and also provides a measure of pore connectivity. Their model was able to achieve reasonable convergence of both the lineal path and two-point correlations and found that the correlations were in good agreement with their respective definitions.

2.2.4 Deep Learning Reconstruction

It is also worth mentioning a novel reconstruction technique that has emerged in the last 6 years or so, which is based on machine learning. This technique uses a type of convolutional neural network called a “generative

adversarial network" (GAN). GANs consist of two separate functions: a discriminator and a generator. The generator creates images, and the discriminator has to decide if they are from the real image data set or if they are 'fake' [28]. The advantage of this method is that once the training has been completed, the network is able to generate many sample images quite quickly. GANs have been used by Mosser et al [28], who implemented them for reconstruction of a bead pack, Berea sandstone and Ketton limestone. Lombardo et al [29] used GANs to reconstruct a lithium-ion battery cathode and a solid oxide fuel cell anode. One advantage of their method is the ability to reconstruct the images in 3 or more phases, whereas conventional reconstruction methods are generally limited to a binary phase structure. Shams et al [30] also used GANs to reconstruct sandstone and carbonate samples from an oilfield. Reconstruction of the PEMFC catalyst layer using GANs has been implemented more recently by Liu et al [84].

This is an interesting and promising approach that should be explored further in future works, but for the purposes of this thesis was not pursued.

2.2.5 Ionomer Segmentation

In all the reconstruction methods mentioned above, a binary structure is generated of solid and pore phase. However, a real CL consists of pore phase, carbon phase, ionomer phase and platinum. Therefore, in order to conduct transport simulations on the reconstruction, it is usually necessary to segment the binary solid and pore phases into 3 or 4 phases (depending on the model, platinum may or may not need to be included).

In their fabrication-based reconstructions, Lange et al [8] tested various ionomer generation methods. These include a) a uniform ionomer thickness covering all surfaces; b) a random ionomer growth on carbon surfaces (uniform thickness not guaranteed); c) random ionomer growth spread uniformly across 4 subdomains; d) random ionomer growth with a gradient such that more ionomer is present closer to the membrane. They found

that the random ionomer growth (case b) resulted in better agreement with experimental results than using the uniform ionomer thickness but did not find significant differences with the other cases, except at very low porosities the gradient based method resulted in lower oxygen diffusion, likely due to the complete filling of some pores with ionomer.

Sousa and Rangel [79] also implemented a segmentation of their agglomerate-growth based reconstruction. In their approach, each carbon particle was assigned an ionomer thickness based on a normal distribution with a mean of 2.5nm and variance of 4nm. However, this approach resulted in some carbon particles that did not get any ionomer coverage once the pre-specified ionomer volume fraction was met.

Finally, So et al [22] used a random walk method to generate the ionomer film on their agglomerate growth-based reconstruction. Ionomer particles perform a random walk until they contact a carbon particle or another fixed ionomer particle at which point, they too become fixed. This approach results in branched agglomerates and more ionomer attached to the tips than the center of the agglomerate.

2.3 Fuel Cell Transport Simulation

Broadly speaking, fuel cell simulation can be broken down into two categories: macro-scale full cell simulation or larger, and component-specific simulations which can be used to analyze each component individually, in more detail than is possible in macro-scale simulations. Macro-scale models have been used to model both isothermal and non-isothermal models, single-phase models (where water is in the vapour form only) or multi-phase models in which water is present as both a vapour and liquid. They also have been used to model both steady-state operation and transient operation such as for cold-start and drive cycling. Macro-scale simulations such as in [85–87] often use volume-averaged or estimated effective transport properties as parameters in their models. Often these parameters are either experimentally derived or very rough estimates from empirical

correlations which usually are based off of only the porosity (such as using the Bruggeman correlation to estimate effective diffusivity). Instead of using such estimated parameters, it is possible to separately run simplified pore-scale simulations on pore-scale geometrically accurate reconstructions of the fuel component (CL, MPL, GDL, membrane) in order to obtain effective transport properties like the effective diffusivity and effective conductivities. These transport properties can then be used as inputs for the macro-scale simulation, theoretically improving their accuracy.

Focusing on the catalyst layer, within the category of pore-scale simulations, there are a number of methods relatively frequently employed in the literature. These include, but are not necessarily limited to: direct numerical simulation where the governing partial differential equations are solved on a computational mesh, typically with the finite volume method; lattice Boltzmann method (LBM) where the Boltzmann particle distribution is tracked in a lattice model of the domain; pore network modeling (PNM) which simplifies the pore geometry into pores and throats, massively reducing the computational cost at the expense of some accuracy; and the random walk methods (RW) where individual particles are tracked as they move randomly throughout the domain. A comprehensive review of modeling methods of PEMFC catalyst layers is given in [31].

2.3.1 Governing Equations

This section defines the generic governing equations which can be used to develop a single-phase model of transport in a PEMFC (neglecting the effects of liquid phase water).

The conservation equation for mass is given by:

$$\frac{\partial}{\partial t}(\varepsilon\rho) + \nabla \cdot (\varepsilon\rho\vec{u}) = S_m \tag{2.8}$$

where ε is the porosity, ρ is the gas density, $\vec{u}$ is the superficial gas velocity and S_m is a source term for mass. The first term represents transient accumulation of mass in the control volume and the second term represents changes in the mass flux [88].

Conservation of species is a form of the conservation of mass equation, but only applied to one species in the domain. Conservation of species for species i is given by

$$\frac{\partial}{\partial t}(\varepsilon C_i) + \nabla \cdot (-D_{\text{eff}} \nabla C_i) + \nabla \cdot (\vec{u} C_i) = \mathrm{S_{s,i}} \tag{2.9}$$

where C_i is the concentration of species i and D_{eff} is the effective diffusivity. The first term describes accumulation or removal of species i on a transient basis. The second term describes changes in species concentration caused by Fickian diffusion. The third term describes changes in concentration from convection. $\mathrm{S_{s,i}}$ is the source or sink term and accounts for changes in concentration due to chemical reaction where a species is either consumed or produced [88]. D_{eff} is typically either obtained from a separate, pore-scale simulation, determined from experiments or estimated from empirical correlations.

The conservation of momentum is described by the following equation [89]:

$$\frac{\partial}{\partial t}\left(\frac{\rho \vec{u}}{\varepsilon}\right) + \nabla \cdot \left(\frac{\rho \vec{u}\vec{u}}{\varepsilon^2}\right) = -\nabla P + \mu \nabla \cdot \left(\nabla\left(\frac{\vec{u}}{\varepsilon}\right) + \nabla\left(\frac{\vec{u}^T}{\varepsilon}\right)\right) - \frac{2}{3}\mu \nabla\left(\nabla \cdot \left(\frac{\vec{u}}{\varepsilon}\right)\right) + S_m \tag{2.10}$$

where the P is the fluid pressure and μ is the mixture average viscosity. In the equation, the first term on the left-hand side represents accumulation of momentum with time and the second term represents the net rate of momentum convected into the control volume. On the right-hand side, the first term represents momentum

changes due to a pressure difference, and the second term represents momentum changes due to viscosity effects. S_m is the momentum source term where for a catalyst layer, it is given by

$$S_m = -\frac{\mu}{K}\vec{u} \tag{2.11}$$

where μ is the gas viscosity, and K is the gas permeability [88].

In PEMFC simulation, the governing equation for conservation of charge is applied to both the transport of electrons in the solid phase and the transport of protons in the electrolyte phase. The conservation of charge for electrical current is given by:

$$\nabla \cdot \left(-\sigma_{e^-,\text{eff}} \nabla V_{e^-}\right) = S_{e^-} \tag{2.12}$$

where $\sigma_{s,\text{eff}}$ is the effective electrical conductivity, V_{e^-} is the solid phase potential and S_{e^-} is a source term for volumetric transfer current at solid-ionomer interfaces. Similarly, the conservation of charge for proton transfer in the ionomer is given by

$$\nabla \cdot \left(-\sigma_{\text{H}^+,\text{eff}} \nabla V_{\text{H}^+}\right) = S_{\text{H}^+} \tag{2.13}$$

where $\sigma_{\text{H}^+,\text{eff}}$ is the effective ionic conductivity in the ionomer phase, V_{H^+} is the ionomer phase potential and S_{H^+} is a source term for volumetric transfer current at solid-ionomer interfaces. In both cases, the first term accounts for the conduction/diffusion of charge and is determined from the effective electrical conductivity of the material [88].

The conservation equation for energy governs the transport and generation of heat in the model and is given by:

$$\frac{\partial}{\partial t}\left(\left(\rho c_p\right)_{\text{eff}} T\right) + \nabla \cdot \left(\left(\rho c_p\right)_{\text{eff}} \vec{u} T\right) = \nabla \cdot \left(k_{\text{eff}} \nabla T\right) + S_{\text{e}} \tag{2.14}$$

where c_p is the mixed average specific heat capacity, k_{eff} is the effective thermal conductivity and T is the temperature. The first term on the left-hand side describes changes in temperature on a transient basis. The second term accounts for heat transfer through convection of gaseous species. The third term accounts for the diffusion (conduction) of heat and the magnitude of this term is determined by the effective thermal conductivity, k_{eff}, which is itself determined by the CL structure, materials and phase fractions. The final term S_{e} is a source term for the heat gained or released from chemical reactions, Ohmic heating and/or phase change [88]. For example, in the cathode catalyst layer of a PEMFC, the source term S_{e} is given by [89]

$$S_{\text{e}} = -\frac{jT\Delta S}{2F} + j|\eta_{act}| + \|\nabla V_{e^-}\|^2 \sigma_{e^-,\text{eff}} + \|\nabla V_{\text{H}^+}\|^2 \sigma_{\text{H}^+,\text{eff}} \tag{2.15}$$

where j is the cathodic reaction rate, ΔS is the change in entropy associated with the reaction, F is Faraday's constant, V_{e^-} is the electric potential in the carbon phase, and V_{H^+} is the ionic potential in the ionomer phase. The first term accounts for the heat generated from the exothermic reaction. The second term accounts for additional heat released due to the activation overpotential and the third and fourth terms account for the Ohmic heating caused by electric and ionic conduction in the carbon and ionomer phases, respectively.

What is evident from all of these governing equations is that an effective transport property is required for each one. The following sections describe some of the literature in which these effective properties are computed from simulations on pore-scale reconstructions.

2.3.2 Ionic Conductivity

In the cathode catalyst layer of a PEMFC, protons are conducted from the membrane to the catalyst reaction sites via an ionically conducting polymer (ionomer). This transport can be modelled using Ohm's Law where

an effective ionic conductivity is a volumetrically averaged parameter that describes the conductivity of protons through the catalyst layer. The steady state diffusion equation for ionic conductivity is given by:

$$\nabla \cdot \left(\sigma_{\mathrm{H}^+,\mathrm{eff}} \nabla V_{\mathrm{H}^+}\right) = 0 \tag{2.16}$$

where σ_{H^+} is the ionic conductivity of the medium and V is the local voltage potential. In the catalyst layer, the ionic conductivity is usually expressed as an effective property where the intrinsic conductivity must be reduced due to the presence of non-conducting phases and due to the tortuosity of the conducting phase.

Mukherjee et al [90] conducted direct numerical simulations on a pore-scale model of the cathode catalyst layer, studying the effect of temperature and relative humidity on the proton transport in the ionomer. They found that higher temperatures and lower relative humidities degraded the ionic transport. This goes to show that management of the water in the catalyst layer is very important. There needs to be enough water to maintain humidification of the ionomer, but not so much that it floods the catalyst layer and blocks the oxygen diffusion.

Kang et al [27] used their sphere-based simulated annealing reconstructions of the CL to investigate the effect of varying the ionomer distribution on the effective ionic conductivity. They used an off-set method in which the film thickness on each carbon sphere is varied by overlaying a larger sphere of ionomer on top of the carbon sphere, but off-set from the center such that the film thickness around the carbon sphere varies, which resulted in some agreement with their experimental data except at lower relative humidities where the model overpredicted the ionic conductivity. Barbosa et al [91] used a multi-scale reconstruction method to estimate ionic conductivity efficiency. Their work confirms that the ionic conductivity is increased with an increasing ionomer load, but at the expense of the electrical conductivity, as an increasing ionomer fraction results in a lower carbon volume fraction.

In their stochastic sphere-packing reconstructions, Lange et al [8,19,20] estimated the effective ionic conductivity for different carbon particle sizes and different ionomer volume fractions. In these studies, the effective ionic conductivity varied from approximately 3% of the bulk property for ionomer volume fractions as low as 15% and smaller carbon particles with a uniform radius of 12nm, up to 18% of the bulk property for larger carbon particles of 20nm radius and ionomer volume fractions up to 35%. This shows that ionic conductivity is increased both when there is more ionomer present and when the carbon particle size is larger. The latter is because larger particle sizes reduce the overall surface area of the model, so the ionomer film thickness is increased for the same ionomer volume fraction.

2.3.3 Electrical Conductivity

Electrons are transported in the PEMFC CL primarily through the carbon phase since air and the ionomer are both electrically insulating. While oxygen diffusion is generally the rate-limiting step in the mass transfer process, it is still important to understand the effect of the pore structure on the electrical conductivity of the CL. Electron transfer in the CL occurs at a much faster timescale than gas diffusion, and so can be considered to be quasi-steady. Therefore, the steady state diffusion equation for electrical conduction can be described by [88]:

$$\nabla \cdot \left(\sigma_{e^-,\text{eff}} \nabla V_{e^-}\right) = 0 \tag{2.17}$$

where σ_{e^-} is the electrical conductivity of the medium and V is the local electric potential. In a catalyst layer, the electrical conductivity is typically expressed as an effective conductivity and is reduced by the presence of non-conducting phases and the tortuosity of the carbon conducting phase.

To this effect, numerous studies have also attempted to estimate the effective electrical conductivity of the CL, both through numerical modeling and experimental measurements.

Inoue et al [39] used a random walk method to study the electrical conductivity of the PEMFC microporous layer and compared this to experimental measurements. In the random walk method, the conductivity is estimated by comparing the mean square displacement of the random walk to what its expected value would be in the bulk ionomer. They found that the electrical conductivity was strongly dependent on the contact resistance and on the PTFE binder.

Barbosa et al [91] used a multi-scale reconstruction method to estimate the electrical conductivity. Their work confirms that the electrical conductivity is decreased with an increasing ionomer load or increased porosity as this reduces the overall volume fraction of carbon in which the electrons can be transported.

Wu et al [23] estimated the effective electrical conductivity of a CL reconstructed using a sphere-based annealing method. They obtained tortuosity values for the carbon phase through an LBM simulation and found that their model predicted tortuosity values approximately 5x bigger than those expected by the Bruggeman correlation. This corresponds to effective electronic conductivities that are 5x smaller than the Bruggeman prediction.

In the sphere-packing work of Lange et al [8,19,20], the effective electrical conductivity was estimated using a finite volume method on a Cartesian grid mesh. Their findings show that the effective electrical conductivity increases from 0.025 S/cm for a carbon volume fraction of 0.25, up to 0.25 S/cm for a carbon volume fraction of 0.45. However, they also found that the effective electrical conductivity was fairly insensitive to the carbon particle size. Singh et al [16] extended this work to use a sphere-based simulated annealing instead of only the sphere-packing, and determined the effective electrical conductivity to be 0.586 ± 0.045 S/cm, but they note that this could be overestimating the actual value by a factor of ~20, citing Saab et al [92] and their experimentally measured result of $(2.6 \pm 0.2) \times 10^{-2}$ S/cm. They attribute this to interfacial contact between carbon particles in a real CL being mediated by the ionomer, as opposed to the perfect contact in the model.

However, it may also be noted that the bulk electrical conductivity of carbon itself can vary by several orders of magnitude depending on the atomic arrangement (for example, carbon fibers, carbon black, graphite, and carbon nanotubes all have different bulk phase electrical conductivities despite all being carbon), so a direct comparison between experimental and numerical results must take this into account.

2.3.4 Thermal Conductivity

The reaction of oxygen and hydrogen to generate water is an exothermic reaction. This heat is generated in the comparatively small volume of the catalyst layer and must dissipate out from the catalyst layer through the rest of the fuel cell and into the surroundings. Moreover, fuel cells are most efficient when operated isothermally [93], so it is important to consider the rate at which heat can be removed from the catalyst layer. This is a critical transport process that any catalyst layer design must consider.

While ionic and electrical conduction are primarily single-phase phenomena in the catalyst layer, thermal conduction occurs through the carbon/platinum solid phase, the ionomer phase, the gaseous phase and through any water that has formed as a byproduct of the reaction. Determination of the effective thermal conductivity can be accomplished by simplifying the conservation of energy equation to Fourier's Law:

$$\nabla \cdot (k_{\text{eff}} \nabla T) = 0 \tag{2.18}$$

The heat flux is given by :

$$q = -k_{\text{eff}} \nabla T \tag{2.19}$$

The thermal conductivity in the catalyst layer is theoretically affected by the intrinsic conductivity of each of its components, in addition to the porosity and liquid water saturation and can be expressed as follows:

$$k_{\text{eff}} = f(k_{\text{C}}, k_{\text{Pt}}, k_{\text{ion}}, k_{\text{w}}, k_{\text{g}}, \varepsilon, \phi_{Ion}, m_{\text{Pt}}, s) \tag{2.20}$$

where k_{eff} is the effective, volumetrically averaged, thermal conductivity of the catalyst layer as a whole, k_{C}, k_{Pt}, k_{ion}, k_{w}, and k_{g} are the thermal conductivities of the carbon black, platinum, ionomer, water, and gas species, respectively. ε is the CL porosity, ϕ_{ion} is the ionomer volume fraction, m_{Pt} is the Platinum loading and s is the liquid water saturation.

A number of studies, both experimental and numerical, have been conducted to estimate the effective thermal conductivity of the PEMFC catalyst layers.

Burheim et al conducted an experimental study of the thermal conductivity of the PEMFC CL [94] under varying compression and water saturation. They found that dry CLs had an effective thermal conductivity of 0.07- 0.11 W/mK.

Barragan et al [95] performed a combined experimental/numerical study estimating the through-plane thermal conductivity of PEMFC ion-exchange membranes. They found that the intrinsic thermal conductivity of the membrane ranged from 0.04 W/mK to 0.42 W/mK.

Lange et al [19] conducted a numerical study to investigate the effects of water vapor and temperature on the effective transport properties. For this purpose, they used the finite volume method (FVM) on their reconstructions generated through a sphere-packing technique. The FVM was solved directly on the Cartesian mesh. They found that the effective thermal conductivity was reduced with increasing porosities and that it was increased with increasing ionomer volume fraction.

Ahadi et al [96] developed a closed-form analytical expression for the thermal conductivity in the CL using a multi-scale reconstruction approach. They found that the thermal conductivity was insensitive to ionomer coverage and the radius of the aggregates, but that it was affected by the carbon particle radius, porosity, platinum loading and ionomer loading.

2.3.5 Mass Diffusion in Porous Media

In PEM fuel cells, oxygen is transported from the flow channels in the bipolar plate through the gas diffusion layer and into the catalyst layer where it reacts with protons and electrons to form water and heat according to Eq. (1.3). To obtain the effective diffusivity, the conservation of species equation can be simplified by assuming steady state diffusion without convection or reaction, the conservation of species equation can be reduced to

$$\nabla \cdot (D_{\text{eff}} \nabla C_{\text{i}}) = 0 \tag{2.21}$$

The diffusive flux is described by Fick's law:

$$N_A = -D_{\text{eff}} \nabla C_{\text{i}} \tag{2.22}$$

In a porous medium, the diffusivity is reduced by a) porosity and b) longer and more tortuous path to navigate. The effective diffusivity is then defined as

$$D_{\text{eff}} = {}^{\varepsilon}/_{\tau} * D_{\text{bulk}} \tag{2.23}$$

where ε is the porosity, τ is the tortuosity and D_{bulk} is the bulk phase diffusivity of the gas.

The effective diffusivity is a value which can then be plugged into a larger scale volume-averaged model. It can be obtained from pore-scale numerical simulation, from empirical correlations or from experiments.

Various empirical correlations have been developed to approximate the effective diffusion coefficient such as the Bruggeman [32] correlation for a packing of spheres:

$$\frac{D_{\text{eff}}}{D_{\text{bulk}}} = \varepsilon^{1.5} \tag{2.24}$$

More recently, Das et al developed an empirical correlation specifically for PEMFC CL, based on a multi-layered coated sphere. Their model extends to predicting diffusion in a 4-phase catalyst layer (gas phase, carbon phase, ionomer phase, and water). However, their relation simplifies to the following for a two-phase (air and solid) structure where the secondary phase is impermeable to gas.

$$D_{\text{eff}} = D_{\text{bulk}} + \frac{3(1-\varepsilon)D_{\text{bulk}}}{\varepsilon - 3} \tag{2.25}$$

Zhao et al [49] measured the effective diffusivity coefficient (EDC) using a modified Loschmidt cell and found it to be 4.6 $e10^{-7}$ m^2/s at 75 °C, which is an order of magnitude smaller than that which is predicted by either the Bruggeman correlation or the correlation proposed by Das et al [33]. Their work showed that as the Pt loading is increased, the porosity and mean pore size decrease, which decreases the EDC.

Yu et al [73,97] measured the EDC of catalyst layers with either amorphous carbon or graphitized carbon with I/C ratios of 0.5, 1.0, 1.5. They found that increased I/C ratios drastically reduced the EDC, which they attributed to the ionomer creating much more tortuous pathways. Salari et al [98] conducted a similar study using a modified Loschmidt cell to measure the relative gas diffusivity. They also found that increasing the I/C ratio decreases the diffusivity by approximately 80%. Their sample with I/C of 0.5 had a relative (dimensionless) diffusivity of 0.055, while the sample with I/C of 1.5 had a relative diffusivity of 0.012.

Inoue et al [39,99] conducted a theoretical examination of the EDC coefficients of the CL, MPL and GDL. They used a random walk method to simulate Knudsen diffusion, calculating the tortuosity as a ratio of the mean square displacement in the porous medium and that in free space. They compared the results of their simulation to their experimental results and found that the CL had an effective diffusivity an order of magnitude smaller than the MPL for the same porosity, despite them both being comprised of carbon black packings. This is attributed to the presence of the ionomer and the extra blockage it creates. At a porosity of

0.4, their experimental results give a relative EDC of 0.008. In [99] they extended the same work and found that the two most dominant factors affecting gas diffusion performance were the morphology of the ionomer adhesion to the carbon black support structure and the presence of large, isolated pores which do not contribute to diffusion.

Jung et al [72] used FIB-SEM images to reconstruct a dual-layered catalyst layer in three dimensions. They proceeded to investigate the diffusion using the random walk formulation as described by Becker et al [100]. Sabharwal et al [101] also reconstructed a catalyst layer from FIB-SEM images and performed continuum-scale diffusion simulations, ignoring Knudsen effects. They conducted extensive post-reconstruction structural analysis, including the two-point correlation, the chord length distribution, pore size distribution, percolating volume fraction by the two-point cluster functions and a geometric image tortuosity. They found that diffusion was anisotropic in the reconstruction which is attributed to the low resolution in the z-direction of the FIB-SEM images.

Zheng and Kim [102] investigate how the microstructure and the water saturation affect the gas diffusion in the catalyst layer. Contrary to the conclusions of Inoue et al [99], they found that variation in the ionomer structures only had minor effects on the effective diffusivity. They also found that the Knudsen effect on diffusion was more pronounced at higher levels of water saturation.

In their sphere-packing reconstruction studies, Lange et al [8,19,20] used the finite volume method to estimate the oxygen and water vapor effective diffusivities. In their simulations they use a local Knudsen number at each voxel, based on the pore size at that site. They found that larger carbon spheres tend to allow larger pores to form as well, which results in a higher effective diffusivity.

In order to determine the importance of the catalyst layer microstructure in determining the effective oxygen diffusivity, Gao and Zhang [103] performed pore-scale simulations on six different reconstructions of the

catalyst layer, including packing of cylinders, packing of isolated spheres, packing of touching spheres, two real CLs from experimental images, and one stochastically reconstructed catalyst layer. Comparing the oxygen diffusivities indicated that the structure of the catalyst layer only played a significant role at higher current densities where concentration overpotential due to mass transfer limitations is the dominant form of polarization.

2.3.5.1 Knudsen Diffusion in Catalyst Layer

Species diffusion in the PEMFC catalyst layer occurs at a nanometer length scale, where the continuum description of transport breaks down and Knudsen effects become significant. The diffusion regime can be characterized by the Knudsen number:

$$\mathrm{Kn} = \lambda / l_c \tag{2.26}$$

where λ is the mean free path of the diffusing species and l_c is the characteristic size of the medium, typically taken as the mean chord length or mean pore diameter. Tomadakis and Sotirchos [48] define three different diffusion regimes: 1) the continuum regime where $\mathrm{Kn} < 0.02$; 2) the Knudsen regime where $\mathrm{Kn} > 50$; 3) the transition regime where $0.02 \leq \mathrm{Kn} \leq 50$. In the continuum regime, diffusion occurs by Brownian motion, being characterized predominantly by molecule-molecule collisions and it can be modeled by conventional methods such as the finite difference, finite volume and finite element techniques. In the Knudsen regime, the scale of the diffusing medium is much smaller than the mean free path of the gas such that diffusion is characterized solely by molecule-wall collisions and molecule-molecule collisions are comparatively rare. For this reason, this regime is also sometimes referred to as the 'rarefied' regime. Knudsen diffusion can be estimated according to the analytical formula [104]:

$$D_{\mathrm{K}} = \frac{1}{3}\bar{v}l_c \tag{2.27}$$

where $\bar{v}$ is the mean thermal velocity of the gas. In between the Knudsen and continuum regimes is the transition regime, in which diffusion is characterized by a combination of Knudsen and continuum effects. This is typically the regime under which diffusion in the PEMFC CL occurs. Analytically, transition regime diffusion is often approximately by the Bosanquet relation [105]:

$$\frac{1}{D_{\mathrm{eff}}} = \frac{1}{D_{\mathrm{K}}} + \frac{1}{D_{\mathrm{continuum}}} \tag{2.28}$$

This relation holds for diffusion from the Knudsen to the continuum regime since in the extremes of either regime, the term for the other regime drops out.

The Navier-Stokes description of diffusion does not adequately describe these non-continuum Knudsen effects. Various alternative strategies have been proposed for estimating and/or simulating the Knudsen and transition regime diffusion. Finite volume methods have been employed using the analytical estimate given by the Bosanquet relation (Eq. (2.28)) to obtain local diffusion coefficients based on the local pore size at each voxel or node [39,99,101]. Similarly, Lange et al [8] and Singh et al [16] simulated oxygen and water vapor diffusion by solving the Stefan-Maxwell equations using a finite volume method, but they account for Knudsen diffusion by defining the local diffusion coefficient at each voxel as a function of the pore diameter, which is calculated as the average of the 13 lengths in different directions.

Alternatively, Knudsen effects have also been accounted for by using a higher-order lattice Boltzmann method with a D3Q64 lattice model (3D lattice model with 64 potential directions of travel for each node) [7], a multiple relaxation time LBM on simplified catalyst layer geometry [106] and an LBM on a D3Q19 lattice

model, accounting for local Knudsen effects using a local diffusion coefficient at each node, based on the local pore size, in a similar fashion to how Lange et al [8] and Singh et al [16] accounted for Knudsen effects in the FVM model.

A third class of methods for simulating Knudsen diffusion are known as Monte Carlo methods, based on stochastic sampling of the pore space. These generally entail some sort of random walk (RW) which tracks particle trajectories in the pore space.

It is well known that an unbiased random walk is able to simulate regular Brownian motion (ie continuum scale diffusion). Tomadakis and Sotirchos [48] pioneered the adaptation of the random walk model to simulate Knudsen and transition regime diffusion in fibrous media. In their method, molecules are simulated as walkers in the pore space, and travel in a step-by-step procedure, changing direction when they either collide with a wall or when they reach the end of their mean free path λ, chosen from an exponential distribution centered on the mean free path calculated at the simulation temperature and pressure. After traveling the distance λ, it is assumed that the molecule collides with another molecule (though no actual collision takes place as the molecules perform their walks independently), and a new random direction is chosen. This same approach was adopted by Berson et al[74]. Becker et al [100] found that this method can diverge $D \to \infty$ for $\mathrm{Kn} \to \infty$, where it should converge to $D_{\mathrm{K}} = {}^{1}/_{3}\, \bar{v} l_c$. They modified the random walk method to calculate the Knudsen regime diffusion without imposing a mean free path restriction: Molecules only change direction after hitting a wall, which effectively isolates the Knudsen regime diffusion. They separately calculate the bulk regime diffusion through a FVM solving Fick's Law. They then calculate the effective diffusion coefficient by merging the Knudsen and bulk regime diffusivities with the Bosanquet relation (Eq. (2.28)). This second approach was adopted as is by Jung et al [72] and Hutzenlaub et al[66].

More recently in 2022, Ferguson et al [107] implemented the same algorithm as Tomadakis and Sotirichos [48], with massively increased computational power compared to the 1993 original work, using OpenMP parallelization and a high-performance distributed computer network. They also convert the voxelized grid into a triangulated iso-surface using Lewiner's marching cubes algorithm [60,61] used for detecting molecule-wall collisions which increases the accuracy compared to using the voxelized grid. They also redefine the characteristic length such that the tortuosity of the medium is one value based on the geometry and independent of the diffusion regime or species being simulated. They validate the method using a direct simulation monte Carlo method that is well-established but computationally intensive.

2.3.6 Other Modeling Frameworks

2.3.6.1 Pore network modeling

Pore network modeling (PNM) is a simulation method which can be described as extremely coarse meshing. In this method, a catalyst layer is subdivided into a network of pores connected to each other via constrictions ('throats'). A system of linear equations is solved where transport occurs through the throats. This method is able to greatly reduce the computational cost of simulation and is highly customizable but does sacrifice pore-scale accuracy. One of the main advantages of PNM over other pore scale models is the ability to simulate transport over a much larger domain than is conventionally realistic. It is also a convenient framework for modeling multi-phase transport [108] .

El Hannach et al [109] employed PNM to analyze water management and electrical conductivity in a PEMFC CL. They used a cubic lattice as their network and assigned throat diameters randomly using an experimental pore size distribution. The study highlights the importance of using a structure with an accurate pore size distribution to obtain an accurate estimate of the effective oxygen diffusivity.

Aghighi et al [110,111] also used PNM to simulate both sides of an MEA, including the GDL, CL and membrane and investigate the effects of water blockage on the polarization curve. They found that water blockage affects not only the concentration polarization, but also Ohmic polarization due to the longer travel distance.

Jung et al [112] compared the use of PNM and direct numerical simulation (DNS) on a catalyst layer that was reconstructed from an XCT image. They confirmed that the DNS simulation is more accurate, while PNM is much faster, and can provide reliable results given a couple of fitting parameters.

In summary, PNM is a simulation method which employs great structural simplifications to provide very efficient computing algorithms but may not be as accurate as pore-scale simulation using DNS. It is not used in this thesis because it does not make much sense to go through the effort to precisely reconstruct the CL microstructure, just to lose most of the detail through the process of generating the network.

2.3.6.2 Lattice Boltzmann

The lattice Boltzmann method is an alternative method for performing computational fluid dynamics (CFD) simulations on pore-scale structures. Unlike more conventional CFD methods in which the continuum Navier-Stokes equations are discretized and solved on a structured or unstructured mesh, the lattice Boltzmann method has more in common with particle-based methods.

The lattice Boltzmann method simulates transport through tracking of particle distributions. For this purpose, it uses the Boltzmann particle distribution. In short, the method consists of two steps: a streaming step in which particles move from one node to another, and a collision step in which particles are given a new direction. A thorough explanation of the lattice Boltzmann method is given in [113].

Due to the nature of the streaming process, the LBM is already optimized to be used on a regular Cartesian grid, whereas solvers using the FVM are often more stable when implemented on a more optimized mesh. For this reason, the LBM is often employed for the complex geometries found in porous media which can be very challenging to mesh successfully. However, the LBM is also extremely memory-intensive, which can be a significant drawback to the method.

2.4 Summary of Chapter

This chapter began by introducing and defining various metrics and statistical correlations by which porous media such as the CL can be characterized. Different methods of reconstruction of PEMFC catalyst layers are then discussed, including a) direct reconstruction from experimental images, b) stochastic fabrication-based methods such as sphere-packing and variations of the agglomerate growth algorithms, c) simulated annealing (both the conventional pixel-swapping method and the sphere-based simulated annealing technique). A brief acknowledgement is made of newer methods of reconstruction involving deep learning and generative adversarial networks. Methods of segmentation of the reconstruction into multiple phases are then discussed.

Following the discussion on porous media characterization and reconstruction, the modeling of transport and the concept of effective transport properties in fuel cells is introduced. The governing equations of transport in fuel cells are defined and discussed. A literature review is put forward on the estimation of the effective transport properties, including the effective ionic, electrical, and thermal conductivities as well as the effective oxygen diffusivity. The latter is considered in more detail, as additional methods for simulation of species diffusion are required in order to account for Knudsen-based effects.

Chapter 3

Methodology

3.1 Methodology: Stochastic Reconstruction and Characterization

3.1.1 Numerical Representation of Structure

The effectiveness of the reconstruction method depends on identifying and using statistical correlations which are computationally efficient, and which can accurately describe a heterogenous random microstructure. In order to describe the microstructure numerically, Torquato [41] defines an indicator function, which when applied to the pore and solid structure of the CL is defined by:

$$I(\boldsymbol{x}) = \begin{cases} 1, & \boldsymbol{x} \in \text{solid} \\ 0, & \boldsymbol{x} \in \text{pore} \end{cases} \tag{3.1}$$

where $\boldsymbol{x}$ is the position within the computational domain.

For implementation of sphere-based simulated annealing (SSA) another digitized field $J(\boldsymbol{x})$ is introduced of the same dimensions of $I(\boldsymbol{x})$ which tracks the number of spheres n which overlap each voxel.

$$J(\boldsymbol{x}) = \begin{cases} 1, 2, \dots n, & \boldsymbol{x} \in \text{solid} \\ 0, & \boldsymbol{x} \in \text{pore} \end{cases} \tag{3.2}$$

The volume fraction of phase i is then just the average of the indicator function:

$$\phi^i = \langle I^i(\boldsymbol{x}) \rangle \tag{3.3}$$

3.1.2 Correlations

Images of porous media can be characterized using n-point correlation functions. This section defines and describes the computations of the two-point correlation, lineal path distributions, chord length distributions and pore size distributions. Note that the 1-point correlation is equal to the phase volume fraction as defined in Eq. (3.3).

3.1.2.1 Two-Point Correlation

Mathematically, the two-point correlation S_2 of a given phase i is defined for a statistically homogenous and isotropic material [41] as

$$S_2{}^i(r) = P\{I(\boldsymbol{x}) = 1,\ I(\boldsymbol{x}+\boldsymbol{r}) = 1\} \tag{3.4}$$

where $\boldsymbol{r}$ is a vector of magnitude $|\boldsymbol{r}|$ away from position $\boldsymbol{x}$**.** Conventionally, the two-point correlation of an image is computed linearly in each Cartesian direction, then averaged. As a brute force method, this is relatively simple to implement but can be computationally expensive for large images. Another much more efficient way of computing the two-point correlation of an image is through its autocorrelation. Using the convolution theorem and the discrete fast Fourier Transform algorithm, the two-point correlation can be computed by taking the inverse Fourier transform of the image's Fourier power spectrum [114–116], which is much faster than the brute force counting method of two-point correlation computation. Patelli and Schuëller [115,116] show that the two-point correlation can be computed from

$$S_2(r) = IFFT\{FFT[I(r)] * \overline{FFT[(I(r)]}\} \tag{3.5}$$

where FFT is the forward fast Fourier transform and $IFFT$ is the inverse fast Fourier transform. The bar represents the complex conjugate. They claim that this method requires $O(N\ln N)$ operations, compared to the

$O(N^2)$ required for the conventional approach, where N is the domain length. They also show that it is a very efficient operation to implement on GPU.

In this thesis, the two-point correlation is computed using this fast Fourier transform method, employing an algorithm modified from *porespy* [55] so that it operates on GPU using *cupy* [117]. Some basic speed testing showed an approximately 50x speed improvement from 7 seconds for the brute force method to 0.15 seconds (computer specifications given in section 3.1.4) for the Fourier Transform method for a 250^3 voxel image (which for a 200,000 iteration simulated annealing is the difference between 360 hours and 8 hours). The computation proceeds as follows:

1. Take the forward 3-dimensional Fourier transform of the image.
2. Obtain the Fourier power spectrum by taking the absolute value of the square of the transform.
3. Obtain auto-correlation by taking the inverse N-dimensional Fourier transform of the power spectrum.
4. For each distance data point r, a radial sum is computed from the auto-correlation image in an annulus (in 2D) or a spherical shell (in 3D) between r and $r + 1$ from the center of the image and normalized by the number of elements in that annulus/spherical shell.
5. The values obtained are scaled by the porosity of the image to obtain the two-point correlation as defined by Eq. (3.4).

It has been shown the two-point correlations for the solid and void phases in a binary structure are linearly dependent [41]. Thus, there is no value in including both the solid and pore phase two-point correlations in simulated annealing reconstruction methods [41]. This thesis uses the pore phase two-point correlation, referred to hereafter as simply S_2. It would be equally valid to use the solid phase.

Additionally, it is important to note that the 2D S_2 for an isotropic structure is the same as for the 3D structure [41]. This is also demonstrated in Figure 3.1, where the blue line indicates the two-point correlation extracted directly from a 3D packing of spheres of uniform radius, while the orange line shows the two-point correlation extracted from every 2D slice of the same domain, taken in 3 directions and averaged together. It is evident that the 2D and 3D S_2 correlations are equivalent, which means that it is valid to optimize the 3D two-point correlation extracted from the reconstruction with the 2D S_2 correlation from the reference image.

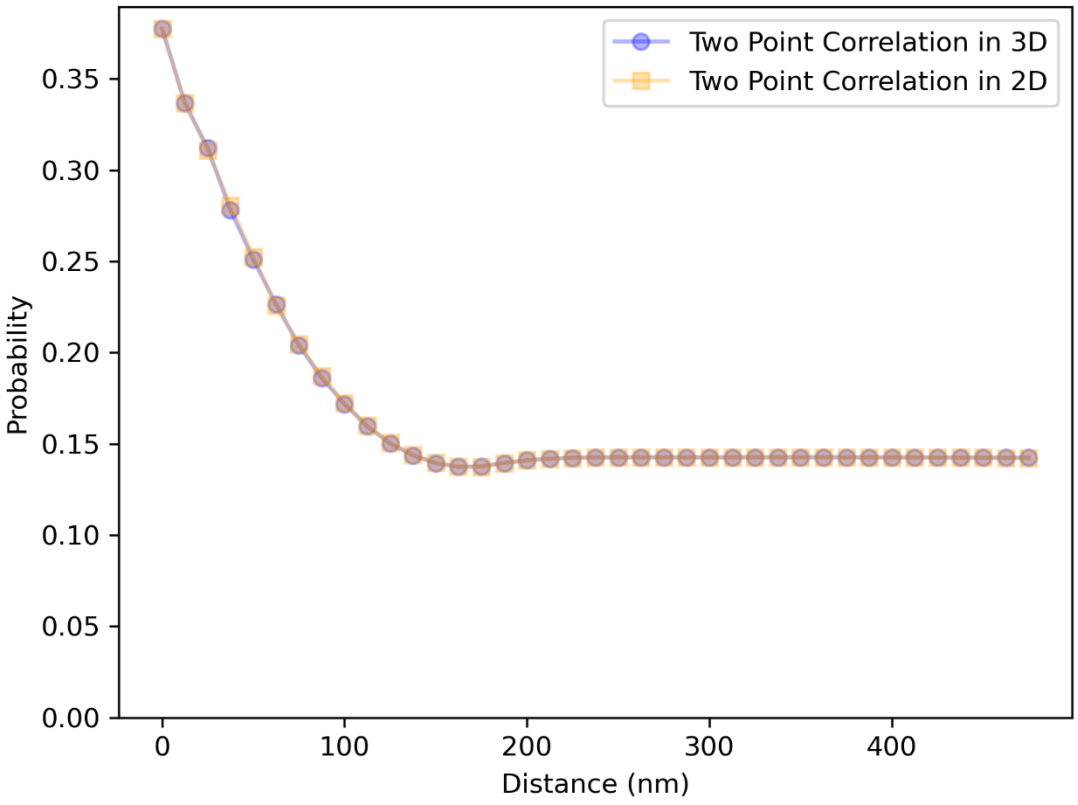


Figure 3.1: Two-point correlation of a random packing of uniform radius spheres. The blue circles indicate the two-point correlation extracted directly from the 3D structure, and the orange circles

indicate the two-point correlation extracted from all of the 2D slices of the domain and average together.

3.1.2.1.1 Degeneracy of the Two Point Correlation

One of the reasons that the two-point correlation may be considered insufficient to be the sole optimization parameter in simulated annealing is that the information contained in the distribution is not adequate to specify a unique structure. In other words, the same distribution may be obtained from multiple different structures with distinct differences in structure. This is referred to as 'structural degeneracy' and has been demonstrated by Gommes et al [118].

For the purposes of a quick demonstration, 3 distinct structures with the same porosity but significantly different geometries are shown in Figure 3.2. Image A has a random web-like structure labelled 'blobs'; image B is a sphere pack of uniform radius overlapping spheres, labeled 'spheres', and image C contains uniform radius overlapping cylinders, simulating a fibrous-type material, labelled "fibers". Figure 3.3 shows the two-point correlations for each of these structures. While they are not identical, they are very similar to each other. To emphasize this point, the images were generated randomly using *porespy* [55]over only 100 iterations, calculating the two-point correlation on each one and keeping the image which had the smallest error compared to the other two. No other optimization was performed showing the ease with which it is possible to create images with similar two-point correlation.

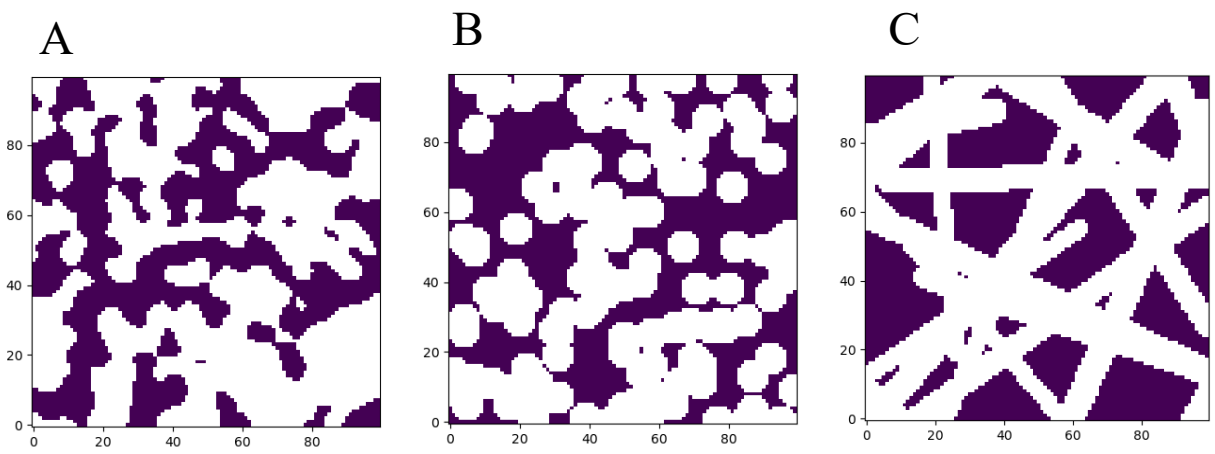


Figure 3.2: Images of distinct structures with the same porosity that have a similar two-point correlation.

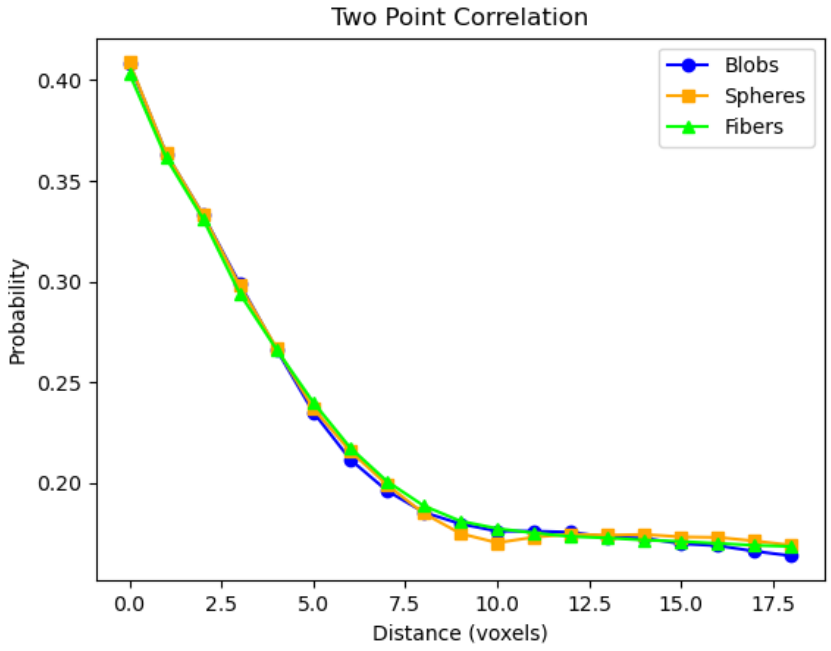


Figure 3.3: Two-point correlations for each of the images in Figure 3.2.

3.1.2.2 Lineal Path Distribution

The lineal path distribution can be computed numerically by applying a linear distance transform in each orthogonal direction and counting the occurrence of each length. This also can be used to check for anisotropy in the image. The lineal path distribution calculation was implemented by slightly modifying an algorithm from *porespy* [55] which computes a linear distance transform in each orthogonal direction then generates a probability density distribution of the lengths. The linear distance transform is depicted in Figure 3.4, while a close-up view close-up view with the counting shown is depicted in Figure 3.5. The linear distance transform computation (using the solid phase as an example) is detailed as follows:

1. For each orthogonal direction, begin at the edge of the image. If the edge voxel is solid, assign a value of 1; if it is pore, assign a value of 0.
2. Continue to move linearly away from the edge. Contiguous solid voxels get incrementing values, while pore voxels will always be assigned a value of 0 (as depicted in Figure 3.5). This continues until a solid-void interface is encountered, at which point the counter terminates and void voxels.
3. When a void-solid interface is encountered, the counting restarts from 1 again.
4. Repeat steps 2 and 3 until reaching the opposite edge of the image.

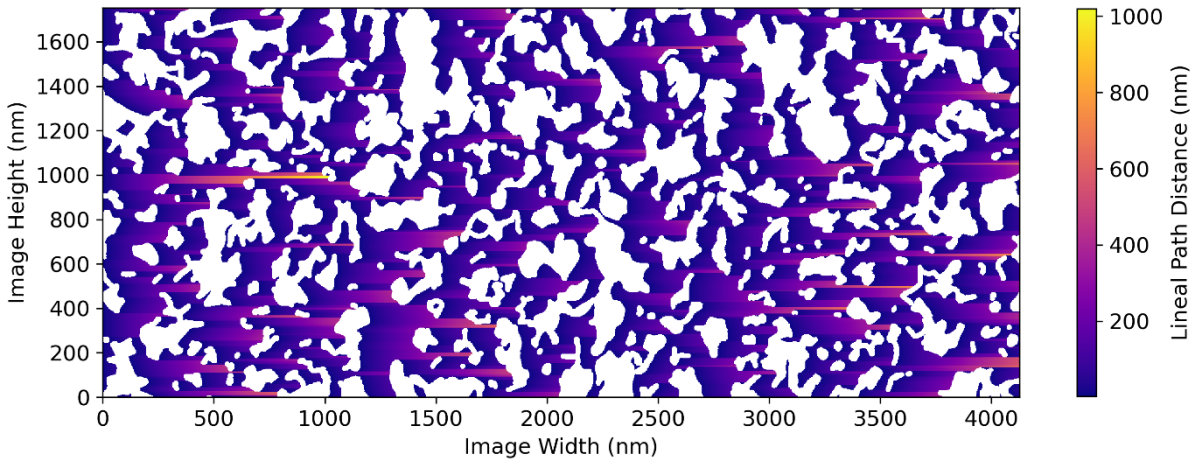


Figure 3.4: An example image of the lineal distance transform applied to the solid phase in the horizontal direction on an SEM image of the catalyst layer. Brighter colors indicate a longer uninterrupted lineal path distance. Original SEM image obtained from [17], **image shown is modified to show lineal path distance transform.**

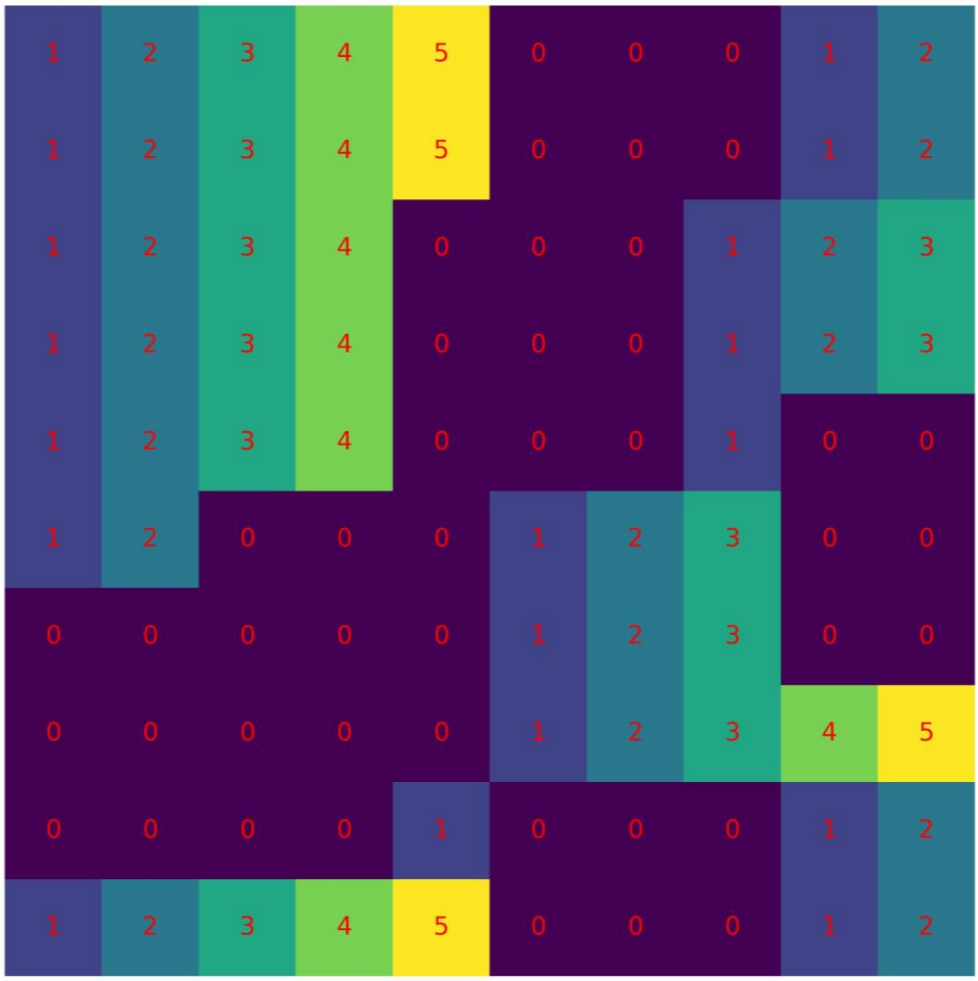


Figure 3.5: Depiction of lineal path distribution.

3.1.2.3 Chord Length Distribution

Chords are defined as the linear distance from one phase interface to the next one. The chord length distribution is thus the relative abundance of the different chord lengths in a porous medium. The mean chord length is the first moment of the chord length distribution and is a useful metric for determining the characteristic length of each phase. When applied to the pore phase, the mean chord length has also been employed as a characteristic

mean pore size. The chord length distribution is also closely related to the lineal path distribution [41], so it is typically considered redundant to optimize both the chord length distribution and the lineal path distribution in simulated annealing processes.

The mean chord length of a porous medium can be estimated from the slope of its two-point correlation, according to Eq. (2.5).

Alternatively, the chord length distribution and mean chord length of a sample image can be obtained directly by measuring the number of pixels from one phase interface to the next in orthogonal directions. Figure 3.6 depicts the chord lengths in the vertical direction applied to an example of a porous media image. Note also that chords connected to the boundary of an image are omitted, as they are artificially shortened.

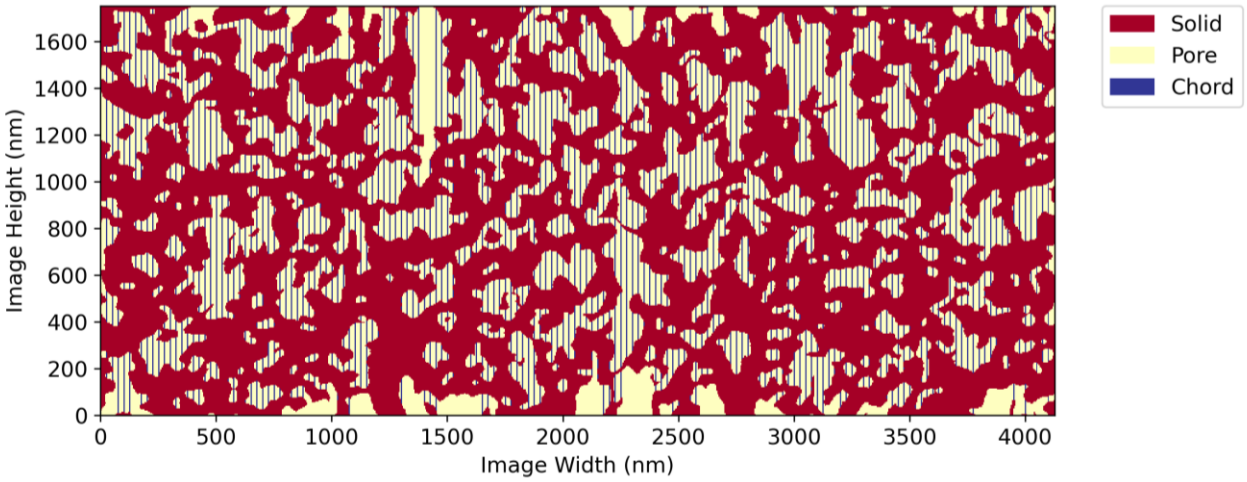


Figure 3.6: An example image of the chords filter applied to a CL SEM image. Red and yellow represent the solid and void phases, respectively. The blue lines represent chords drawn in the pore phase (separated from each other for visualization purposes). Chords that would touch the image boundary are artificially shortened and so are removed in order to avoid skewing the data. Original SEM image obtained from [17], **image shown is modified to show chords.**

3.1.2.4 Pore and Particle Size Distribution

Determination of the pore size distribution (PSD) and particle size distribution (PaSD), referred to collectively as the radial size distributions (RSD), of the reconstruction is useful in at least two ways: 1) It can be used as a model validation by comparing the RSDs of the reconstruction to that of the reference image, and 2) the PSD is a means to obtaining the average pore size, which is useful when determining effective transport properties such as the effective diffusivity. However, the RSDs are inherently three-dimensional properties, so that those of the 3D reconstruction cannot be directly compared to those extracted from the 2D reference image. For this reason, a 2D analog of the RSDs is computed from the 3D reconstruction by computing it independently on each slice of the image in each axial direction and then averaging the values. Unless otherwise stated, in this thesis, PSD, PaSD and RSD refer to the 2D analogues of the actual pore, particle and radial (combined pore and particle) size distributions. While the 2D analogue is not a true representation of the experimental 3D pore distributions, it is an effective way to compare the RSDs of the reconstructions to that of the 2D reference image. It is assumed that if the 2D RSDs match, the 3D RSDs will be reasonably similar as well.

In the first part of this thesis, the RSDs were computed by using a local thickness function from *porespy* [55], in which each voxel in the phase of interest is assigned a value equal to the radius of the largest sphere that will fits entirely within that phase and which encompasses the voxel, noting that the voxel does not have to be in the center of the sphere, as is the case for the Euclidean distance transform.

The local thickness function works by sequentially inscribing spheres of decreasing size into the pore space, and adding up the volume of largest spheres that will fit in the pore space. Spheres are inserted in order from the largest to smallest, without overwriting previous values. Each individual voxel contains the value of the largest sphere which can be inscribed in the pore space containing that voxel. An example of the output of local thickness function is pictured in Figure 3.7. Conveniently, the local thickness function is phase-agnostic,

and can be applied to characterize the solid phase particle size distribution as well and so can be used to obtain both the PSD and PaSD from images of porous media.

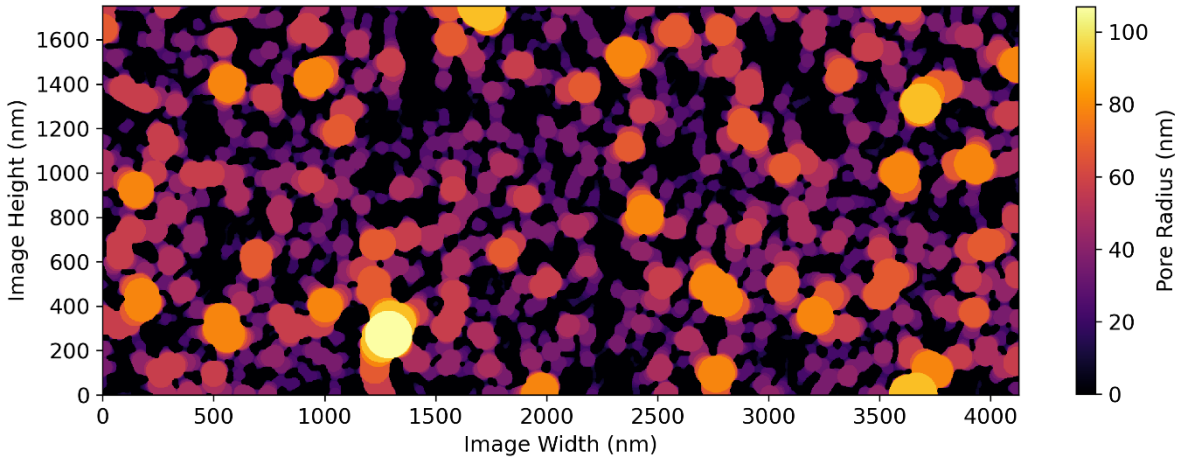


Figure 3.7: An example illustration of the local thickness filter applied to the pore phase of a CL SEM image. Brighter colours indicate larger pore sizes. Original SEM image obtained from [17], **image shown is modified to show local thickness filter.**

The pore size distribution is inherently a 3D characteristic of a porous structure [41], and to the best of this author's knowledge, it is not possible to obtain the true pore size distribution from a 2D image. This presents an issue in reconstructing a 3D structure using only information extracted from a 2D reference image. However, it is possible to obtain a 2D local thickness distribution or "circle size distribution" from a 2D image, which is an analogue of the 3D pore size distribution. If this distribution is extracted from a reference 2D image of a catalyst layer, it is then possible to use this to optimize a reconstruction structure using simulated annealing. Unfortunately, the implementation of the local thickness in *porespy* [55] was not fast enough to be

practical for simulated annealing purposes. For this reason, an alternative implementation was developed using a medial axis basis (described in the next section), which is much faster but also sacrifices some accuracy.

3.1.2.4.1 Medial Axis Implementation of Local Thickness

In order to optimize the RSDs in the simulated annealing method, it was necessary to develop a faster implementation of the local thickness function. For this purpose, this thesis presents a new implementation of the local thickness function using a medial axis-based method, similar to that which is implemented in ImageJ [119]. The implementation in this thesis was developed for the purpose of directly optimizing the pore size distribution in simulated annealing. The medial axis of an image is an intermediate step of this local thickness computation and is taken to mean a skeleton of the pore or solid space with each pixel on the skeleton containing the value of the distance to the nearest phase interface.

The medial axis local thickness is computed as follows:

1. Start from a binary image where the phase of interest is indicated by voxels tagged with unit value, and other voxels are tagged as zeros (Figure 3.8A).
2. Take the Euclidean distance transform of the image, shown in Figure 3.8B.
3. Conduct an image thinning skeletonization process such as fast skeletonization in 2D only using the Zhang and Suen method [120] or a slightly slower implementation developed by Lee et al [121] that works for both 2D and 3D images. Both of these skeletonization algorithms are implemented in *scikit-image* [122] and were used as is. This yields a skeleton of the original image, where the skeleton is one voxel thick along the centre line of the pore space and is represented by voxels of unit value, whereas everywhere else is zero. The skeleton is shown in Figure 3.8C.

4. Obtain the medial axis, or ‘distance ridge’ by taking the value of the Euclidean distance transform everywhere along the skeleton. This is accomplished by elementwise multiplication of the distance transform from step 1 with the skeleton, yielding the medial axis skeleton, as shown in Figure 3.8D.
5. Overlay the distance ridge onto the original domain image, as shown in Figure 3.8E.
6. Fill pore space to generate local thickness image (Figure 3.8F):
    a. Insert a sphere in 3D or a circle in 2D, centred on each point along the distance ridge, of radius equal to the value of the centre voxel on the distance ridge. Each voxel within the radius of the sphere/circle is also assigned the value of the radius. In the case where a voxel is contained in multiple spheres centred on different points of the distance ridge, only the larger value is kept. This process is accomplished by looping through every pixel on the skeleton, however the loop is sped up by using just-in-time (jit) compilation implemented by the python package *numba* [123].
    b. Step 6a is able to nearly fill all the pore spaces, but there can remain some pore spaces that were not reached by the inserted spheres due to being in sharp corners or smaller enclosed pore spaces which the skeletonization process discarded. These remaining voxels are assigned a value equal to the Euclidean distance transform at that position.
7. The last step is to take all the non-zero values in the flattened local thickness representation of the image and compute the histogram, yielding the solid or void phase radial size distribution.

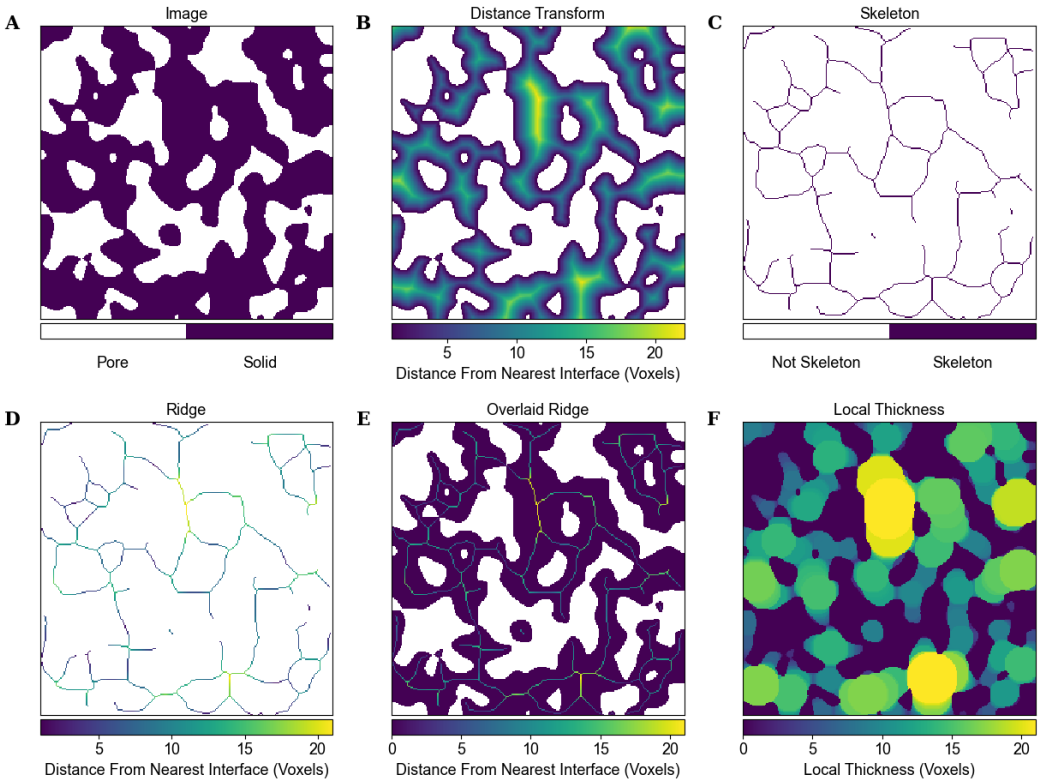


Figure 3.8: Medial Axis Implementation and Progression Depiction. A) Binary Image of Porous Medium (artificially generated). B) Euclidean distance transform on image from (A); C) a skeleton of the pore phase; D) Medial Axis Ridge (skeleton with values from distance transform); E) Medial axis ridge overlaid on image of original pore space; F) Local Thickness image generated by adding spheres to each location on ridge in (E).

To determine the 2-dimensional RSD of a 3D image, the 2D local thickness filter is applied to each 2D slice in a reconstruction image in each of the x, y and z directions. For each direction, a 3D array is saved consisting of all the 2D slices stacked on top of each other. When a sphere is moved, it is removed from one location and inserted at a second location. Since the pore size distribution is being calculated on a per-slice (2D) basis, the majority of the slices do not change for each individual sphere move. A significant computation time saving measure is therefore to only recalculate the pore size distribution for the slices which intersect with either the

old or new sphere locations. For example, a sphere with a radius of 7 intersects with 15 slices in each orthogonal direction. If the sphere is moved to a far location with no shared slices, a maximum of 90 slices will be affected. Recalculating the local thickness filter for only those 90 slices is much more efficient than recalculating every slice in the reconstruction (which would be 750 slices for a reconstruction of size 250^3 voxels). In order to enable this time-saving feature, a copy of the medial axis filtered model must be kept for each orthogonal direction (3 copies for a 3D model).

In addition to the above, it is possible to compute the local thickness for a subset of the slices in the reconstruction. In this study, the pore size distribution for calculated for every third slice in the reconstruction. Due to the nature of using spheres and assuming a radius>=1, each configuration change is going to change a number of adjacent voxels. By sampling every third slice, the algorithm is able to further reduce computation time without significantly affecting the reliability of the algorithm. Additionally, by sampling in each Cartesian direction, the structure should remain isotropic since no contiguous set of voxels is left unsampled.

3.1.2.5 Surface Area

Numerically, the interfacial surface area for images of porous media can be obtained by two primary methods. It can be estimated from Eq. (2.7) using the slope of the two-point correlation, or alternatively, it can be computed by generating a triangular mesh from the reconstructed geometry using a marching cubes algorithm [60], and then measuring the mesh surface area. The latter method is used in this thesis and was implemented using functions provided in the *scikit-image* image processing library in python [122]. Additionally, it should be noted that the marching cubes algorithm requires a 3D image to produce reliable results, which requires entering the entire stack of 2D FIB-SEM images together, while specifying the spatial dimensional anisotropy. While this provides a reasonable result, there remains some uncertainty due to the resolution of pores between each slice in the z-direction.

Ultimately, these numerical surface area estimates should be validated against experimentally determined values, such as those determined using the method of standards of porosimetry [57] or Brunauer–Emmett–Teller (BET) methods [59].

3.1.2.6 Validation of the Correlation Functions

In order to validate the correlation functions used in the simulated annealing process (the two-point correlation, lineal path distribution and pore size distribution), an artificial image is generated of uniform radius overlapping spheres, for which an analytical expression is known for the two-point correlation and lineal path distribution. To the author's knowledge there is no analytical distribution for the pore size distribution of overlapping spheres, so to validate the new medial-axis based implementation of the local thickness function, the distribution produced by this algorithm is compared to that of the local thickness algorithm already implemented in porespy.

For the purpose of the following validations, an artificial image of overlapping spheres is generated on a 250 x 250 x 250 voxel domain with a uniform sphere radius of 5 voxels and a porosity of 0.50. A slice from this image is shown in Figure 3.9.

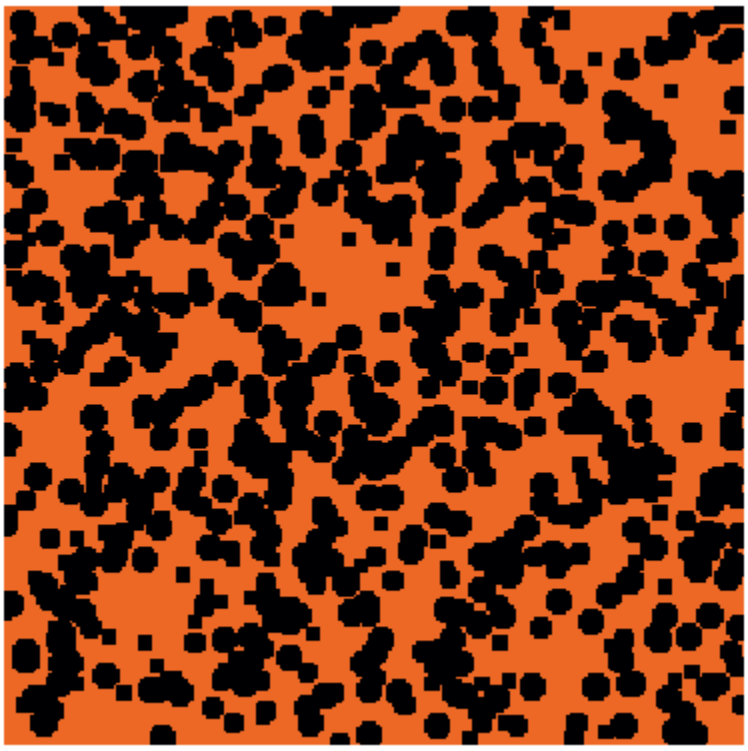


Figure 3.9: Artificial image generated of a sphere pack of overlapping spheres with a uniform radius of 5 voxels and a porosity of 0.50 in a domain of size 250 x 250 x 250 voxels. A 2D slice is shown of this 3D domain.

The two-point correlation for a packing of overlapping spheres of uniform radius R is given by [124]

$$S_2 = \begin{cases} \exp(-\rho V(r)) & \frac{r}{R} < 2 \\ \frac{8\pi R^3}{3} & \frac{r}{R} \geq 2 \end{cases} \tag{3.6}$$

where ρ is the sphere packing density given by

$$\rho = -3\ln(\varepsilon)/4\pi R^3 \tag{3.7}$$

and $V(r)$ is an expression for the union volume of two spheres of radius R and a distance r away from each other (centre to centre) given by:

$$V(r) = \frac{4\pi R^3}{3} * (1 + \frac{3r}{4R} - \frac{r^3}{16R^3}) \tag{3.8}$$

Figure 3.10 shows the two-point correlation for the pack of overlapping spheres. Close, but not identical agreement, is found between the analytical expression and the two-point correlation of the image. The small discrepancy where the empirical correlation is overestimating the analytical expression is likely due to the binning of the distances in the numerical computation, whereas the distances in the analytical expression are continuous. Even with this small difference, it is expected that the numerical two-point correlation will be sufficiently accurate for the reconstruction process.

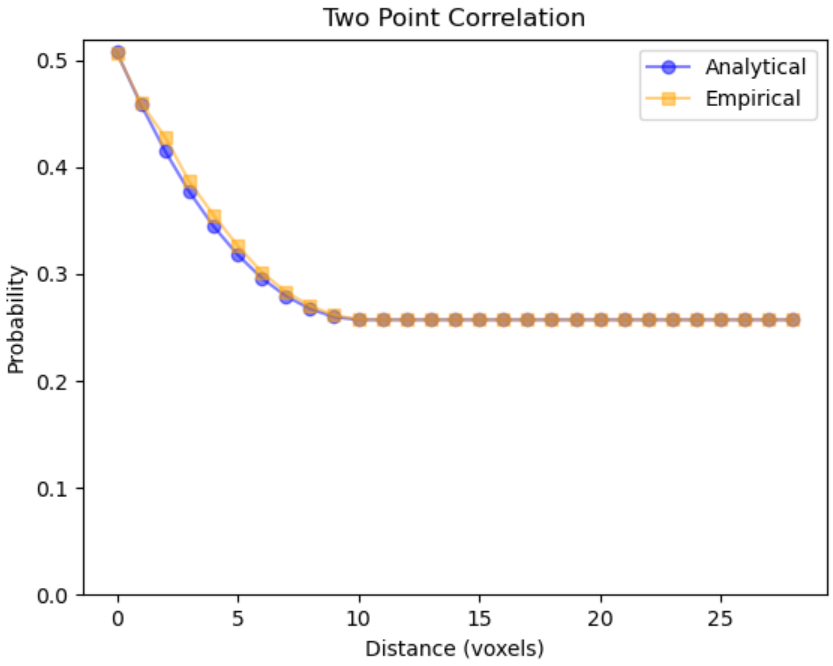


Figure 3.10: Two-point correlation validation of empirical function using the analytical expression for a pack of overlapping spheres.

The lineal path distribution was also validated via another analytical expression for the packing of overlapping spheres. This expression is derived in [124] and given by

$$L^i = \phi^{i\left(1+\frac{3r}{4R}\right)} \tag{3.9}$$

The lineal path distribution extracted from the sphere pack is plotted against the analytical expression in Figure 3.11. In this case, there can be no doubt that the image-based correlation is an accurate description of the analytical distribution since the plots are nearly perfectly overlapping. The binning issue that arises with the two-point correlation may not appear in this case because the lineal path distribution is calculated in axial directions on the Cartesian grid and so the distances that are evaluated are all discrete integers.

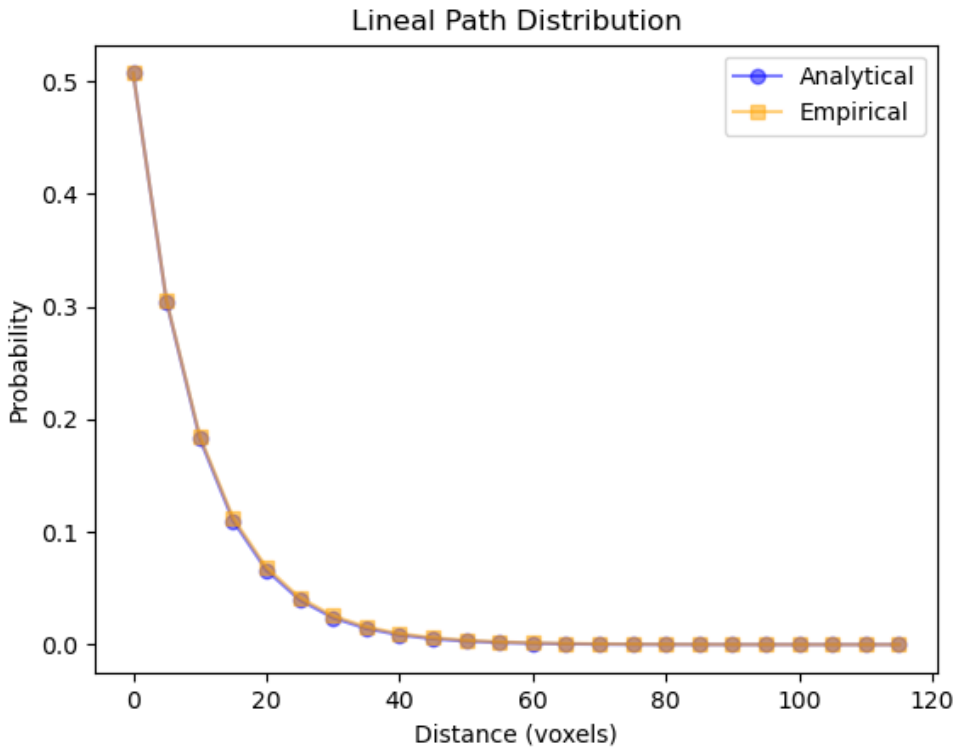


Figure 3.11: Lineal path distribution validation of numerical implementation using the analytical expression for a pack of overlapping spheres.

Unlike the lineal path and two-point correlations, to the best of the author's knowledge there is no analytical expression for the pore size distribution of a pack of overlapping spheres. For this reason, the newly implemented medial axis-based computation of the local thickness RSDs is compared to the version of the local thickness already implemented in *porespy*. The distributions for both versions are shown in Figure 3.12, which demonstrates that both functions are computing very similar distributions. The version in *porespy* may be marginally more accurate because it works by inserting spheres at every point in the distance transform, rather than only using the distance ridge. However, the speed improvement of using the medial axis computation is what makes use of the pore size distribution in simulated annealing computationally practical so the minor reduction in accuracy is considered worthwhile for this purpose.

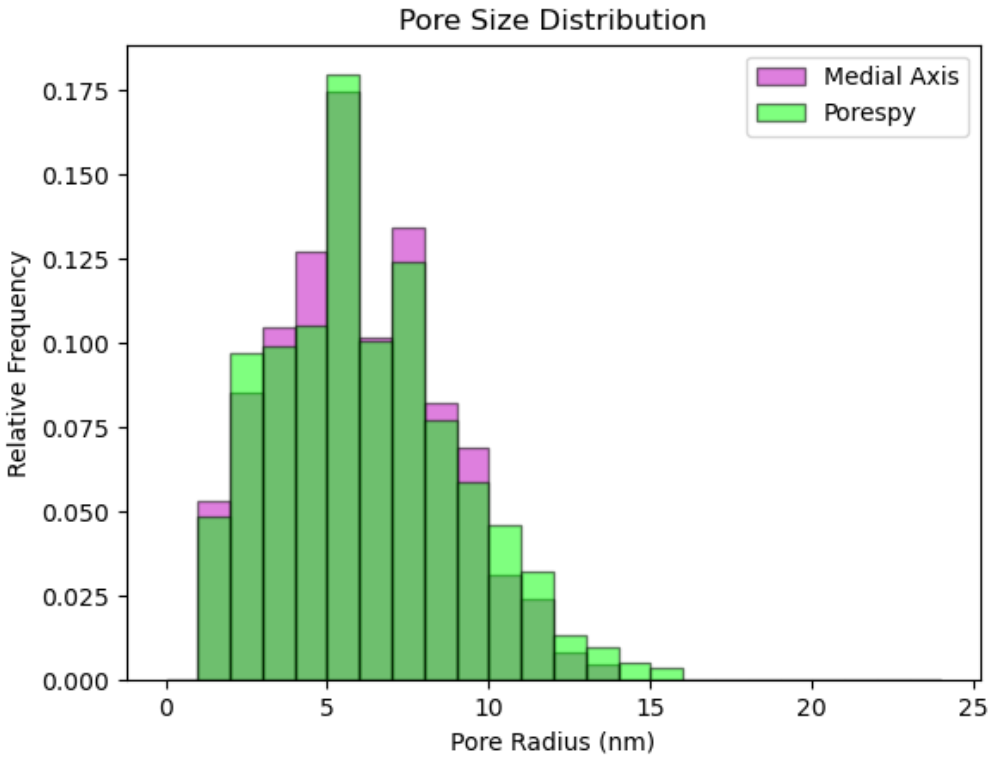


Figure 3.12: 2D analogue of the pore size distribution calculated by the medial axis method the distance transform method in *porespy*.

3.1.3 Stochastic Reconstruction

3.1.3.1 Image Processing

The reconstruction process begins with processing of the experimental reference images. In this case, the binarized images have been taken from the work done by Thiele, Ziegler et al. [17,18], in which the full set of segmented and binarized FIB-SEM images of a PEMFC CL were made available through the supplementary materials of their publication. The image set is a stack of 113 images of dimensions 1651 x 701 pixels that have a voxel resolution of 2.5nm in the x-y plane (or the so-called in-plane), but each section is 30nm apart, effectively limiting the resolution in the z-direction (or the through-plane direction) to 30nm. Due to the artificial anisotropy caused by the limited resolution in the z-direction, the reference two point and lineal path correlations were obtained independently for each 2D image in the stack, and then averaged across 113 images for each correlation. This effectively simulates using a single 2D image, if the assumption is made that the image is representative of the CL as a whole. The averaged two-point correlation and lineal path distributions (hereafter referred to as the reference correlations) were used as the target configuration in the simulated annealing reconstruction.

Additionally, due to limitations in computational power, it was not possible to reconstruct a 3D model of the entire field of view of the reference images. Instead, the reference image was scaled down in the in-plane direction from 2.5nm to the desired reconstruction resolution. The correlations for the reference were calculated (Figure 3.13) and it was determined that all three correlations had decayed at least 99% or more of their theoretical total decay from $r = 0$ to $r = \infty$ by approximately 500 nm. The two-point correlation requires that the image has a minimum dimension of at least 2x the max radial distance measured. Additionally, the lineal path distribution is artificially shortened at the edges of the image, so a doubling of the required dimension reduces the skew caused by the edge effect. Therefore, in order to include enough space in the

domain to properly represent the correlations and avoid significant edge effects, an additional 250nm padding was added to the required minimum dimension in the image, resulting in a reconstruction of size 1250 x 1250 x 1250 nm with a voxel resolution of 5nm when reconstructed in a 250^3 voxel domain This represents a 2x coarsening of the reference resolution from 2.5nm to 5nm. In order to easily compare the reference distributions with the reconstruction, the reference image was resized to have a 5nm resolution as well.

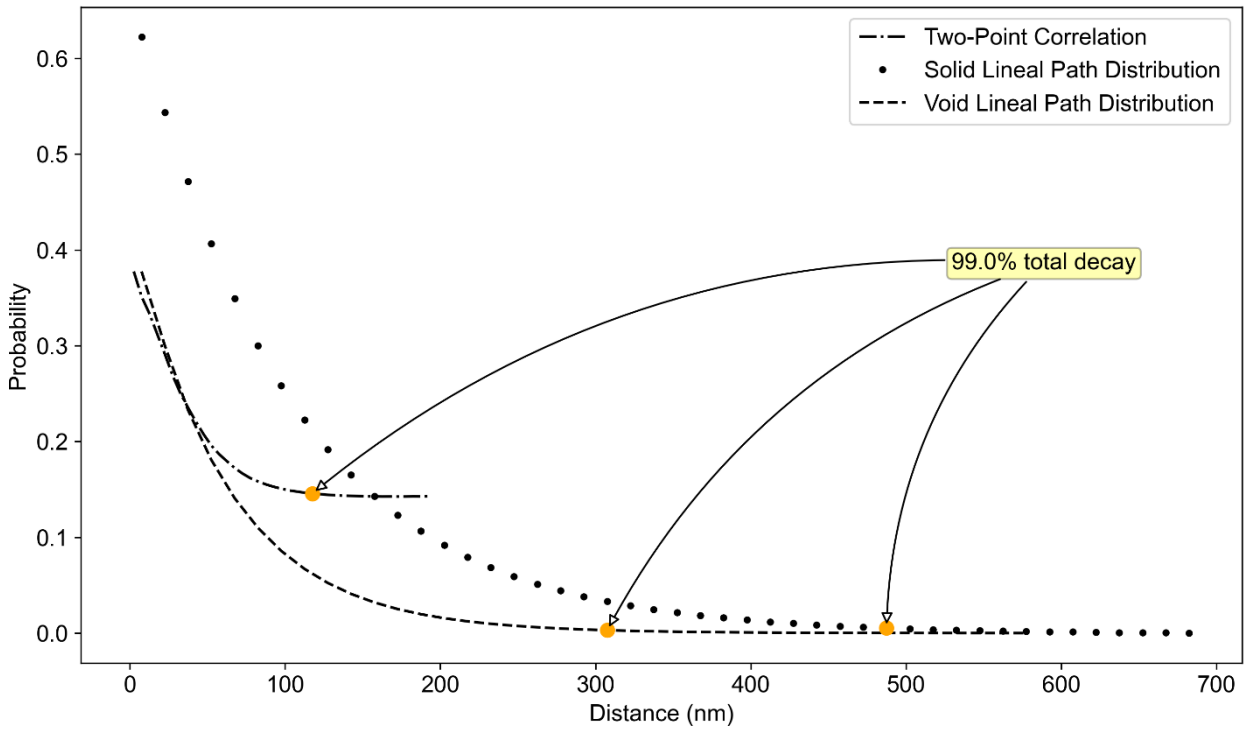


Figure 3.13: Reference image two-point correlation, and solid and pore phase lineal path distributions. Arrows point to where all three correlations have reached at least 99% of their theoretical total decay.

3.1.3.2 Sphere Packing

As discussed previously, the hybrid reconstruction method consists of two distinct phases: 1) the fabrication-based sphere-packing phase, also referred to as the initialization phase, in which factors such as CL porosity

and carbon particle sphere size distributions are considered; and 2) the simulated annealing phase in which the structure is fine-tuned using stochastic structural descriptors such as the porosity, two-point correlation lineal path distributions and radial size distributions.

One of the advantages of including the sphere-packing reconstruction in hybrid reconstruction is the speed at which it is possible to generate a structure which is coarsely representative of the reference structure. By tuning the parameters used, it is possible to drastically reduce the number of iterations required of the simulated annealing stage, which is orders of magnitude more computationally intensive than sphere-packing (the sphere-packing approach is usually completed in under a minute, whereas the simulated annealing takes at minimum several hours on a computer with a 12^{th} gen Intel i7 processor, 16 GB RAM and an 8GB Nvidia RTX 3050 GPU).

Solid spheres of variable size are iteratively inserted into the initially empty reconstruction domain $I(\boldsymbol{x})$ of size $N \times N \times N$ voxels until the reference image porosity is reached. Sphere sizes are defined by their radius and the radius measurement is the magnitude of the distance from the center to center of each voxel and is inclusive, such that a sphere with a zero-voxel radius is a single voxel, a sphere with unit radius is 3 voxels in diameter, a sphere with radius of 2 voxels is 5 voxels in diameter and so on. 2D depictions of the sphere representations are shown in Figure 3.14.

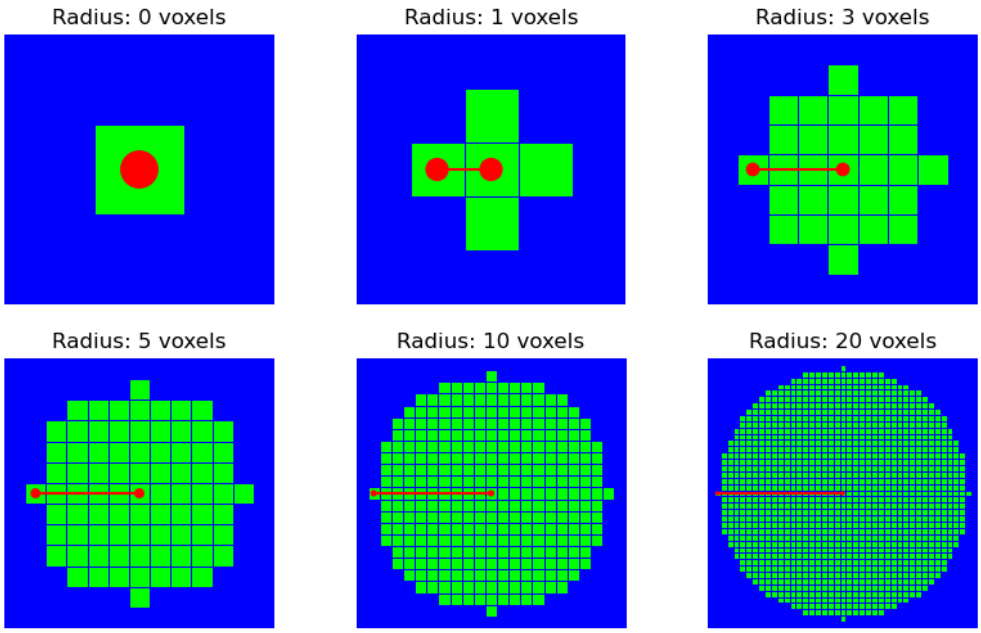


Figure 3.14: 2D depictions of sphere representation.

The computational domain is periodic, such that spheres wrap around the edge of the domain. Spheres are allowed to overlap with each other using the probability given by

$$P(\nu) = \exp\left(-\frac{\nu}{\epsilon_\nu}\right) \tag{3.10}$$

where ν is the ratio of overlap volume to sphere volume and ϵ_ν is a clustering parameter controlling the sphere overlap.

3.1.3.2.1 Variable Radius

A first approximation of the initial sphere size distribution was obtained by computing the local thickness of the solid phase of the reference image and fitting this distribution to a variety of probability distributions in

the *scipy.stats* python library, including the Normal distribution, the double Weibull, Maxwell, Rayleigh, alpha, and an exponentially modified Normal. The exponentially modified Normal distribution given by:

$$f(r,\mu,\sigma,\lambda) = \frac{\lambda}{2}\exp\left(\frac{\lambda}{2} * (2\mu + \lambda\sigma^2 - 2r\right)\operatorname{erfc}\left(\frac{\mu + \lambda\sigma^2 - r}{\sqrt{2}\sigma}\right) \tag{3.11}$$

was found to have the best fit for the data where r is the sphere radius, μ is the distribution mean, σ is the distribution standard deviation and λ is the exponential shape parameter. Note that in the *scipy.stats* implementation, a shape parameter K is used where $K = {}^{1}/_{\sigma\lambda}$. Further adjustment of the mean, standard deviation and shape parameter was accomplished by performing an initial simulated annealing reconstruction while allowing the spheres to increment or decrement their radius on each iteration. This slightly reduced the mean and standard deviation of the distribution, though the shape parameter remained relatively constant. The final distribution parameters are tabulated in Table 3.1.

A summary of the sphere-packing phase of hybrid reconstruction is as follows:

1) Initialize empty list of sphere coordinates and radii, and computational domains $I(\boldsymbol{x})$ and $J(\boldsymbol{x})$ as periodic boxes of $N \times N \times N$ voxels.
2) Determine sphere radius based on the probability distribution given by Equation (3.11) with the parameter values given in Table 3.1. Randomly select a set of unoccupied coordinates within the domain for the sphere center.
3) Insert sphere into $I(\boldsymbol{x})$ domain.
4) Calculate the volume of the overlap region and accept or reject based on probability given by Equation (3.10).
    a. If accepted, add sphere coordinates and radius to the list of spheres.

b. If rejected, revert $I(\boldsymbol{x})$ to the previous configuration (it is useful to keep a *safe* copy and a *trial* copy of the domain to make reversion to an acceptable configuration easy).

5) Repeat steps 2-5 until the reference image porosity has been achieved.

Table 3.1: Parameters input to an exponentially modified Gaussian distribution used to generate the sphere radii inserted into reconstruction.

Parameter	Value
Mean (μ)	14.1 nm
Standard Deviation (σ)	10.2 nm
Shape Parameter (λ)	0.073

3.1.3.2.2 Constant Radius

For the constant radius sphere-packing, the process is the same as described in section 3.1.3.2.1, except that at step 2, the sphere radius is a pre-specified parameter and is not chosen randomly. The sphere radius used was equal to 7 voxels or 35 nm at a resolution of 5nm/voxel. This provided a reasonable structure using between 7000 and 8500 spheres which is in the same range as the variable radius sphere-packing.

3.1.3.3 Sphere-based Simulated Annealing

The sphere-packing phase is used to obtain a coarse reconstruction of the reference CL, while the simulated annealing phase is meant to fine-tune the structure so that it is statistically equivalent to the reference image structure.

A significant advantage of simulated annealing is the ability to optimize multiple parameters simultaneously. This thesis has employed simulated annealing to optimize the CL structure according to a variety of combinations of the porosity, the two-point correlation, pore and solid phase lineal path distributions and the radial size distributions of the reference image. This ensures that at minimum, the reconstruction should have similar interfacial surface area, characteristic length, mean pore and particle size and phase volume fractions as the reference image. The simulated annealing algorithm implemented in this study was adapted from that used in [7,27].

An 'energy cost' function E (hereafter referred to simply as the energy) is first defined as a measure of the statistical difference between the reconstruction and the reference configurations, computed as the sum of squared differences of the distributions relative to their reference states. The individual energy costs associated with the porosity, S_2, L^v and L^s , PSD and PaSD (denoted E_ϕ, E_{S_2}, E_{L^v}, E_{L^s} E_{PSD}, and E_{PaSD}) are expressed in the equations given below:

$$E_\phi = \left\| \phi^v - \phi^v{}_{\mathrm{ref}} \right\| \tag{3.12}$$

$$E_{S_2} = \sum_{r=0}^{r_{\max,S2}} [S_2(r) - S_2(r)_{\mathrm{ref}}]^2 \tag{3.13}$$

$$E_{L^v} = \sum_{r=0}^{r_{\max,L^v}} [L^v(r) - L^v(r)_{\mathrm{ref}}]^2 \tag{3.14}$$

$$E_{L^s} = \sum_{r=0}^{r_{\max,L^s}} [L^s(r) - L^s(r)_{\mathrm{ref}}]^2 \tag{3.15}$$

$$E_{\mathrm{PSD}} = \sum_{r=0}^{r_{\max,PSD}} [PSD(r) - PSD(r)_{\mathrm{ref}}]^2 \tag{3.16}$$

$$E_{\mathrm{PaSD}} = \sum_{r=0}^{r_{\max,PaSD}} [PaSD(r) - PaSD(r)_{\mathrm{ref}}]^2 \tag{3.17}$$

where r_{max} is set for each correlation such that it includes the entirety of the distance r over which that correlation varies, as seen in Figure 3.13. The reference distribution is denoted with a subscript 'ref' while the reconstruction distribution has no subscript. The area under the curve for each correlation S_2, L^v, L^s, PSD and PaSD is normalized to unity so that the energy cost of each correlation is comparable to that of the others.

The nature of allowing some sphere overlap in sphere-based simulated annealing means the porosity of the reconstruction can vary from iteration to iteration due to sphere clumping or spreading (unlike conventional simulated annealing where precisely one void pixel and one solid pixel are swapped at each iteration, so the porosity will never change). Therefore, it is useful to include the porosity energy cost E_ϕ directly in the simulated annealing energy cost function to prevent the porosity of the reconstruction to shift away from the reference porosity. The porosity of a model is a single value instead of a distribution, which is why its energy cost is expressed as the magnitude of the absolute error, rather than a sum of squared errors.

This thesis presents 6 reconstruction methods:

Method 0: A reconstruction using only the sphere-packing of the model domain and no simulated annealing. The structure generated using this method serves as the starting point of Methods 1-5.

Method 1: The conventional sphere-based simulated annealing which optimizes the two-point correlation and the porosity. Simulated annealing begins from the final structure generated by Method 0.. The simulated annealing energy cost is given by:

$$E_{\text{Method 1}} = E_{\phi} + E_{S_2} \tag{3.18}$$

Method 2: The newly implemented sphere-based simulated annealing which optimizes the porosity, two-point correlation, and lineal path distributions for solid and void phase. Simulated annealing also begins from the final structure generated by Method 0, and the energy cost function is given by:

$$E_{\text{Method 2}} = E_{\phi} + E_{S_2} + E_{L^v} + E_{L^s} \tag{3.19}$$

Method 3: Simulated annealing also begins from the final structure generated by Method 0, and the energy cost function is given by:

$$E_{\text{Method3}} = E_{\phi} + E_{PSD} + E_{\text{PaSD}} \tag{3.20}$$

Method 4: The newly implemented sphere-based simulated annealing which optimizes the porosity, two-point correlation, and pore and particle size distributions. Simulated annealing also begins from the final structure generated by Method 0, and the energy cost function is given by:

$$E_{\text{Method 4}} = E_{\phi} + E_{S2} + E_{PSD} + E_{PaSD} \tag{3.21}$$

Method 5: Simulated annealing also begins from the final structure generated by Method 0, and the energy cost function is given by:

$$E_{\text{Method 5}} = E_{\phi} + E_{S2} + E_{L^v} + E_{L^s} + E_{PSD} + E_{PaSD} \tag{3.22}$$

A flow diagram of the simulated annealing algorithm is presented in Figure 3.15. At each iteration, a sphere is chosen from the domain and moved to a new location. Additionally, the sphere radius is given a 1/3 chance each of increasing by 1 voxel, staying the same, or decreasing by 1 voxel. In this way, the sphere size distribution can dynamically adapt to the target structure and the algorithm is given some additional freedom to increase or decrease the sphere radii in order to maintain the correct porosity.

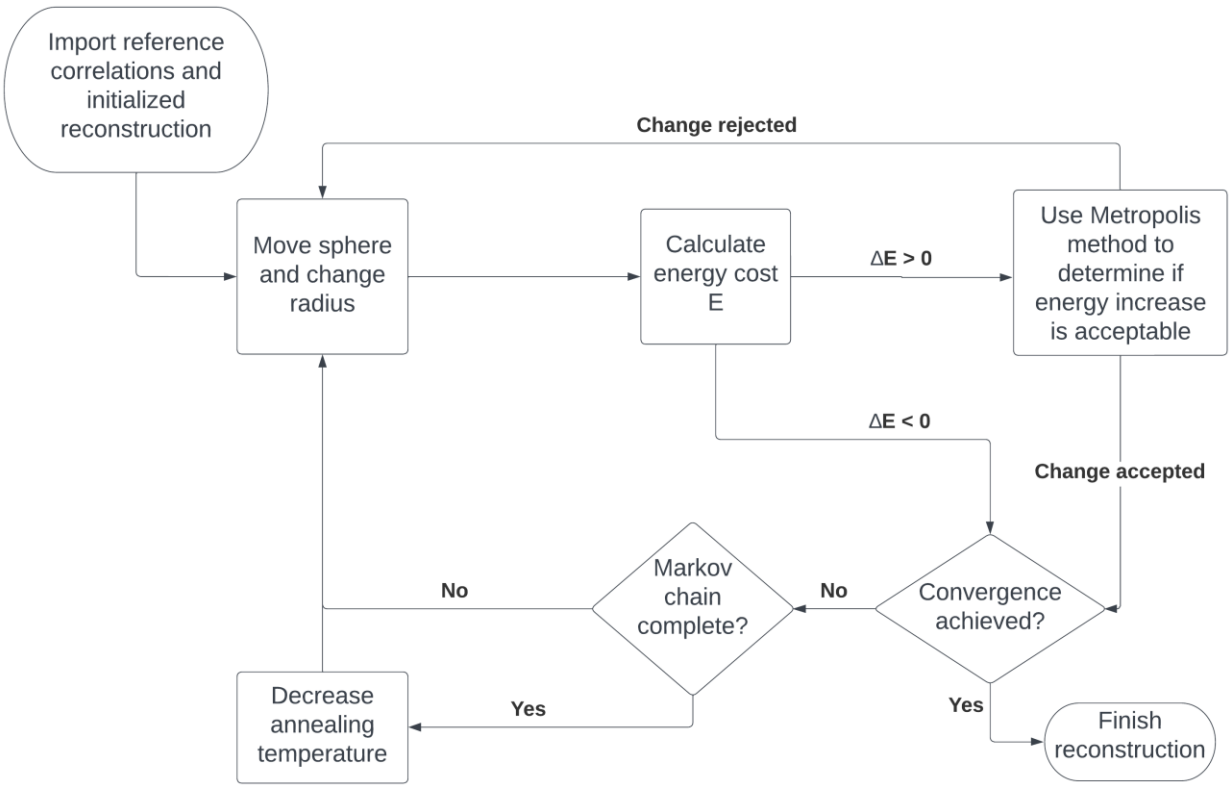


Figure 3.15: Sphere-based simulated annealing algorithm flow control.

After moving a sphere, the correlations and associated energy cost are recomputed and the change in energy is calculated as

$$\Delta E = E_n - E_{n-1} \tag{3.23}$$

where E_n is the energy in the current trial configuration, and E_{n-1} is the energy of the previous successful configuration. If the sphere movement results in an energy decrease, the new configuration is accepted unconditionally. If the sphere movement results in an energy increase, it is only accepted using a Metropolis probability scheme:

$$P(\Delta E) = \begin{cases} \exp\left(\frac{-\Delta E}{T}\right), & \Delta E > 0 \\ 1, & \Delta E < 0 \end{cases} \tag{3.24}$$

where T is a fictitious annealing temperature which is used to control the probability of accepting energy increase due to sphere movements. This is done in order to avoid being trapped in local minima within the energy landscape, as could occur if no energy increasing movements were permitted. The temperature is reduced based on a cooling schedule such that as the structure 'solidifies', it becomes increasingly difficult to make energy increasing movements. The same starting configuration can have different absolute energy values depending on which correlations are employed. Since the probability of accepting an energy increase is dependent on the absolute value of the energy increase, it is necessary to define the initial temperature in a way that a similar acceptance rate will be achieved regardless of which correlations are being optimized. In order to achieve this, the energy of the starting configuration is computed, and then the initial temperature is calculated so that it gives a 25% chance of accepting a 0.1% of initial energy increase in the energy by rearranging Eq (3.21):

$$T_i = -0.001 * \frac{E_i}{\ln(0.25)} \tag{3.25}$$

Thereafter, the temperature is decreased on a geometric scale given by

$$T^{(m+1)} = \lambda T^{(m)}, \qquad 0 < \lambda < 1 \tag{3.26}$$

where λ is a pre-specified scaling factor and m is the Markov chain number. The rate of temperature decrease is controlled by a Markov chain mechanism such that the temperature decrease is only induced after N_{MC} accepted configuration changes or N_T total trial configurations at each temperature where $N_{MC} < N_T$,

theoretically allowing the configuration to reach an equilibrium at each temperature stage. The program terminates after a pre-set maximum number of trial iterations. In this thesis, each reconstruction was run to 200,000 total iterations.

3.1.3.4 Segmentation and Post Reconstruction Processing

The simulated annealing reconstruction algorithm described up until now, in sections 3.1.3.1 to 3.1.3.3 generates a binary structure composed of pore space and a solid phase. However, the solid phase in a conventional PEMFC CL is actually made up of a carbon black backbone, an ionically conducting polymer (ionomer), and platinum particles acting as a catalyst. As the gas phase is already segmented, the structure can be used as is to obtain the gas phase diffusivity (neglecting gas diffusivity through the ionomer). However, in order to perform transport simulations in the solid phase, it is necessary to segment the solid into its component materials of carbon, ionomer and platinum. To achieve this end, the solid phase is first segmented into solid/ionomer, and then the remaining solid phase is segmented into carbon and platinum.

Before segmentation begins, a couple of pre-processing steps are necessary in order to ensure the resulting structure is realistic and will yield accurate results. One of the issues with the original reconstruction is that there is no mechanism that ensures that all solid spheres are interconnected, nor is there a mechanism that prevents the formation of blind pores. The presence of non-connected solid phase (essentially floating solids) is both physically unrealistic and will cause singularities when solving the linear algebra matrix. On the other hand, blind pores (pores not connected to the rest of the pore space) are not necessarily unrealistic, as these do exist, but they do not contribute to the transport in the gas phase and similarly to the disconnected solid phase, they will cause issues when solving the transport equations. Generally, the reconstructions generated have a very low percentage of blind pores or disconnected solid phase, so removal of these does not have a large effect on the overall structure or porosity but does fix the issues with formation of singularities.

A number of segmentation methods were investigated by Lange et al. [8], which include 1) specifying an ionomer thickness and all cells greater than the carbon radius, but less than the sum of the carbon radius and ionomer thickness are tagged as ionomer cells; 2) randomly assigning carbon-adjacent cells to be ionomer until reaching a pre-specified ionomer phase fraction 3) random ionomer growth with a uniform distribution; 4) random ionomer growth with a normal distribution and 5) random ionomer growth with a gradient. For methods 3-5, they subdivide their domain into smaller sections in order to implement the distribution.

For the sake of simplicity and in the absence of concrete data on the actual distribution of ionomer in the CL, the method chosen for segmentation in this thesis is most similar to method 2 of Lange et al, where random cells adjacent to the carbon phase are chosen and tagged as ionomer or platinum until a prespecified phase volume fraction is met.

The following describes the step-by-step process of segmentation implemented in this thesis.

1. Start from the reconstruction output from simulated annealing, where pore voxels are tagged with a 0 and solid voxels are tagged as a 1.
    a. Blind pores are removed using the *porespy* [55] function *fill_blind_pores*, which converts these cells to solid phase, tagging them with a 1.
    b. Floating solids are removed using the *porespy* [55] function *trim_floating_solids*, which converts these cells to pore phase, tagging them with a 0.
2. All of the indices of the solid phase voxels on the solid/pore interfaces are determined. This is accomplished by performing a 1-voxel morphological erosion of the solid phase.
3. These edge voxels are then randomly and sequentially tagged as ionomer phase by assigning a value of 2 until the pre-specified phase fraction of ionomer is achieved. If there are insufficient edge voxels, a secondary erosion of the remaining untagged solid phase is performed, and the process is repeated

until the desired ionomer phase fraction is achieved. This ensures that the ionomer will form a randomly distributed layer covering the carbon agglomerates, while also not changing the porosity of the image.

4. Steps 2 and 3 are repeated for the platinum phase (with a tagging identifier value of 3), using a pre-specified platinum phase fraction and performing the erosion on the remaining untagged solid phase (ignoring both the pore and ionomer phases).
5. Any remaining solid phase after tagging all the ionomer and platinum is assumed to represent carbon (with a tagged value of 1).
6. Finally, a second check of the connectivity is conducted to ensure that the entire carbon phase remains interconnected, and that no ionomer is disconnected. Any carbon phase that is disconnected is converted to ionomer, and any ionomer that is disconnected is converted to carbon.

A 2D slice of the final segmentation of the reconstruction is depicted in Figure 3.16 where the platinum is represented by red dots, the ionomer is a thin green layer which covers the grey carbon phase. The pore phase is represented by the white colour.

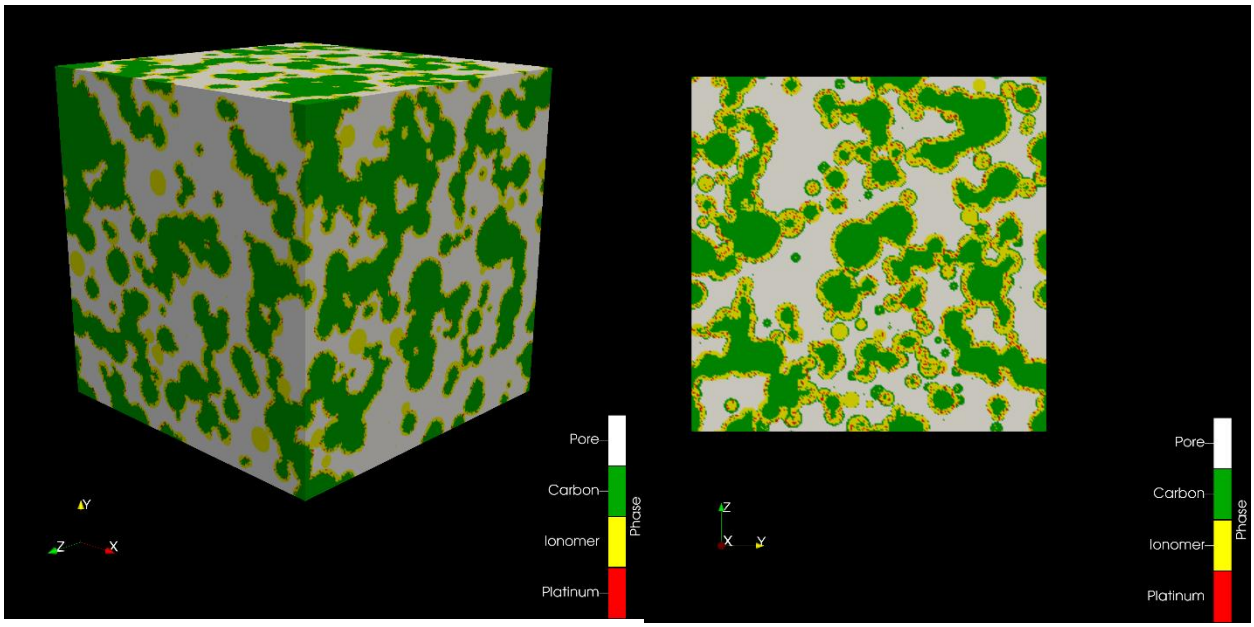


Figure 3.16: An example of the segmented reconstruction. Note that due to the 3D nature of the reconstruction, solids that appear to be floating and pores that appear blind in the 2D depiction are not actually disconnected from the rest of their respective phases, as they are connected in the 3rd dimension.

3.1.4 Numerical Implementation

The codes used in this thesis are entirely written in the Python programming language, making ample use of the numpy [125], scipy [126] and matplotlib [127] libraries. The main simulated annealing algorithm was developed in-house, while many of the correlation and characterization functions have been taken and/or adapted from porespy [55], an open-source library that specializes in "quantitative image analysis of porous microstructures".

The programs are all run on a Windows PC with a 20-core 12th gen Intel i7 processor, 16 GB RAM and an 8GB Nvidia RTX 3050 GPU.

3.2 Methodology: Transport Simulation and Effective Properties

In this section, each of the reconstruction methods that were used to generate catalyst layer structures are evaluated by computing the effective transport properties of each structure. The following sections describe the implementation for computation of the effective thermal conductivity, effective electrical conductivity, effective ionic conductivity and effective oxygen diffusivity.

3.2.1 Implementation of Solvers for Thermal, Ionic and Electrical Conductivity

Thermal conductivity, electrical conductivity and ionic conductivity are expected to exhibit continuum scale behavior and as such, the effective conductivities for these transport mechanisms can be found by solving the Navier-Stokes equations using a finite volume method.

These effective properties were computed by running a finite volume method simulation for each property. This code was run in Python using functions from the Porous Microstructure Analysis (PuMA) library [128]. The solvers in this library are designed to be run directly on the Cartesian mesh output of the reconstruction and segmentation processes, and so no additional meshing was required. The solvers are converged using a biconjugate gradient stabilized method [129].

Each conductivity was solved individually in each of the 3 cartesian directions. Dirichlet boundary conditions were assigned to the ends of the domain in the direction of interest, while periodic boundary conditions were used on the sides of the domain for the reconstructions and symmetric boundary conditions were used to simulate the transport on the reference image as it is not periodic.

The Cartesian mesh is represented by a 3D array with voxels whose values indicate which phase they belong to. Voxels with a value of ‘0’ indicate the gas/air phase, a value of ‘1’ indicates the carbon/platinum solid

phase and a value of '2' indicates the ionomer phase. Each phase is assigned a bulk conductivity, which are listed in Table 3.2.

Table 3.2: Bulk Phase Conductivities

	Ionic Conductivity (S/cm)	Electrical Conductivity (S/cm)	Thermal Conductivity (W/cm-K)	Oxygen Diffusivity (m^2/s)
Carbon	0	0.564 [39]	0.3 [130]	0
Ionomer	0.03 [131]	0	0.25 [95]	0
Air	0	0	0.0255 [132]	2.19E-5

One thing to note is that the platinum contribution to all transport mechanisms is neglected in these simulations. That is to say, the transport simulations are conducted on reconstructions that are segmented into pore, carbon and ionomer only, while the platinum volume fraction is small enough that it is assumed to have negligible effect on the simulation results. This simplification was made because inclusion of the platinum in the simulations caused numerical instability which then had difficulty in converging.

3.2.2 Random Walk Model for Knudsen Diffusion

3.2.2.1 Theory

While finite volume solvers can be successfully implemented for continuum-based solutions, species diffusion in the PEMFC catalyst layer is significantly affected by Knudsen effects because the average pore size is

approximately on the same order as the mean free path of oxygen, and so the continuum description of transport breaks down and the Navier-Stokes equations cannot be used to accurately solve diffusion in this case. In a PEMFC CL, the average pore size is approximately 60-70 nm, while the mean free path of oxygen at 298K and 1 atm is approximately 75 nm. This yields a Knudsen number $\mathrm{Kn} = \frac{\lambda}{l_c} \cong 1$, which is considered to be in the transition regime, indicating that both Knudsen effects and continuum effects are significant and that each has an approximately equal effect on the effective diffusivity. This means that a diffusion simulation of the PEMFC CL must take both into account in order to obtain an accurate result.

To solve for the transition regime effective diffusivity in the reconstructed catalyst layers, a modified random walk (RW) method was developed and implemented which is able to predict the effective diffusion rate, while accounting for the non-continuum Knudsen effects. The algorithm is based off of the work of Tomadakis and Sotirchos [48], but is not identical. It also borrows ideas from Becker et al [100] and more recently that of Ferguson et al [107].

It is well established that regular Brownian diffusion (continuum diffusion) can be modeled by an unbiased random walk with a step size $d_{step} \ll l_c$, where l_c is the characteristic length of the domain [133,134]. In this continuum random walk scenario, the walker changes direction after every step. In an unbiased random walk, the mean displacement of a walker has an expected value of zero, which is not very useful. For this reason, the mean squared displacement $\langle \boldsymbol{r}^2 \rangle$ is used to compute the diffusion, where the angular brackets are used to denote the mean. The mean squared displacement at a time t is defined as:

$$\langle \boldsymbol{r}^2 \rangle_t = \sum_{i}^{N_{\text{walkers}}} \|\boldsymbol{x}_{ti} - \boldsymbol{x}_{0i}\|^2 \Big/ N_{\text{walkers}} \tag{3.27}$$

The expected value of the $\langle \boldsymbol{r}^2 \rangle$ in a free space simulation is the product of the number of steps taken and the square of the step size[135]:

$$E(\langle \boldsymbol{r}^2 \rangle) = N_{\text{steps}} * {d_{\text{step}}}^2 \tag{3.28}$$

Note that the number of steps taken is actually the number of direction changes, which is an important consideration for the modification of this concept to Knudsen diffusion, as will be described shortly.

This scenario can be adapted to account for Knudsen diffusion by adjusting the step size: instead of having the walker change direction after every step, it continues in the same direction until it either collides with a wall or until it has travelled a distance equal to its mean free path λ, at which point it is assumed it has collided with another molecule. The mean free path of a gas is defined according to kinetic theory as:

$$\lambda = \frac{\mu}{P}\sqrt{\frac{\pi RT}{2 * MW}} \tag{3.29}$$

where μ is the gas dynamic viscosity, P is the pressure, R is the ideal gas constant, T is the temperature and MW is the molecular weight of the gas. Molecules are assumed to travel at the mean thermal velocity, $\bar{v}$, also derived from kinetic theory and defined as:

$$\bar{v} = \sqrt{\frac{8kT}{\pi * MW}} \tag{3.30}$$

where k is the Boltzmann constant. After running the random walk with sufficient walkers and for a sufficient number of steps such that a steady state is reached, the effective diffusivity can be calculated according to the Einstein equation [133,134,136,137] as:

$$D_{\text{eff}} = \varepsilon \frac{\langle \boldsymbol{r}^2 \rangle}{2\Delta t_{\text{total}}} \tag{3.31}$$

where ε is the domain porosity and Δt_{total} is the time elapsed for the number of steps taken. It is defined as the total distance travelled divided by the mean thermal velocity:

$$\Delta t_{\text{total}} = d_{\text{total}} / \bar{v} \tag{3.32}$$

An alternative method of obtaining the effective diffusivity is through a tortuosity argument, where the tortuosity τ is defined as the ratio of the expected free space mean square displacement to the simulated mean square displacement:

$$\tau = \frac{E(\langle \boldsymbol{r}^2 \rangle_{\text{free space}})}{\langle \boldsymbol{r}^2 \rangle_{\text{porous}}} = \frac{N_{\text{steps}} * {d_{\text{step}}}^2}{\langle \boldsymbol{r}^2 \rangle_{\text{porous}}} \tag{3.33}$$

Note that in the Knudsen modified random walk, the walkers continue in a straight line until reaching their mean free path distance. Each mean free path is divided into smaller steps, which means that the walkers do not change direction at every step. This means that the above equation needs to be modified as follows:

$$\tau = \frac{E(\langle \boldsymbol{r}^2 \rangle_{\text{free space}})}{\langle \boldsymbol{r}^2 \rangle_{\text{porous}}} = \frac{N_{\lambda} * \lambda^2}{\langle \boldsymbol{r}^2 \rangle_{\text{porous}}} \tag{3.34}$$

where N_{λ} is the number of mean free paths the walkers would be expected to undergo and is defined as:

$$N_{\lambda} = N_{\text{steps}} * \frac{d_{\text{step}}}{\lambda} \tag{3.35}$$

The effective diffusivity using the tortuosity method is defined as:

$$D_{\text{eff}} = \frac{\varepsilon}{\tau} D_{\text{bulk}} \tag{3.36}$$

3.2.2.2 Random walk Implementation

This section lays out the practical details of implementation of the random walk procedure.

3.2.2.2.1 Image Pre-processing

One important restriction in order to run this random walk simulation is that the entirety of the pore space must be percolating. This requires identification of the primary pore region, and subsequent removal of all pores not connected to this region (referred to as blind pores). This includes any pore space on the edge of the image as this can create a blind pore when the image is tiled symmetrically. If these blind pores are not removed, any walkers populated into the blind pores will never be able to escape no matter how long the simulation is run for, and this would erroneously reduce the calculated diffusion rate. This would be evident as the diffusion plot would never reach a true steady state. Removal of the blind pores does not affect the integrity of the simulation because while there may be blind pores in an actual catalyst layer, they cannot contribute to the transport of the gaseous reactants. Removal of blind pores is accomplished using functions from the *porespy* library and is discussed in further detail in the segmentation section of this thesis (section 3.1.3.4).

The issue of image percolation is even more significant when conducting the random walk on a 2D image. In 2D images of a 2-phase structure, it is not possible for both phases to be completely percolating and the likelihood of having blind pores is much higher. For this reason, it is recommended to use this algorithm only on 3D structures or images, as the results from a 2D simulation would be misleading.

3.2.2.2.2 Determination of Step Size

While a continuum random walk may use discrete steps, for the Knudsen random walk, each mean free path λ of the gas is split into many smaller steps, so that the algorithm may account for the effect of molecule-wall collisions and the resulting effective shortening of the measured mean free path due to molecule-wall

collisions. For this work, the mean free path of oxygen was split into 50 individual steps. More steps per mean free path will reduce the noise but also will also mean the simulation will take longer to reach a steady state value.

One important note about step size is that it should not be more than 1 voxel in whatever resolution the image being studied has, as this could allow the walkers to pass through thin walls as the collision with the wall would not be registered.

3.2.2.2.3 Determination of Direction

When a walker encounters a wall or reaches the end of its mean free path, a new direction must be chosen. Diffuse reflection is assumed to occur, since specular reflection is unrealistic at this scale due to surface roughness and the voxelated nature of the reconstruction [48]. One method of implementing diffuse reflection is to use only the axial directions, commonly done in continuum diffusion regime random walk simulations. However, for Knudsen diffusion it is preferable to select a direction vector from a continuous angle on a unit sphere.

A uniformly distributed unit direction vector is determined [138] by first generating two random numbers u, v from uniformly distributed interval $[0,1]$. Two intermediate variables q and f are computed:

$$q = \cos^{-1}(2u - 1) - \frac{\pi}{2} \tag{3.37}$$

$$f = 2\pi v \tag{3.38}$$

The direction unit vector is then:

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} \cos f * \cos q \\ \cos q * \sin f \\ \sin q \end{bmatrix} \tag{3.39}$$

where the magnitude of the vector is equal to unity so that the distance can easily be scaled to the step size of the walker.

3.2.2.2.4 Boundary Treatment

When a walker reaches the edge of an image, the boundary effect must be accounted for. A periodic boundary condition can be implemented if the image is already periodic, as is the case for the reconstructions generated for this work. This is implemented by tiling the image as necessary when the first walker exits the boundaries of the current image.

However, if the image is not periodic as is the case for the reference image in this work and for experimental images in general, then implementation of periodic boundary conditions will result in erroneous results as the edges will not match. In this case, it is preferable to use symmetric boundary conditions. Symmetric boundary conditions are implemented by tiling the *reversed* image so that the edges will always match. Care must be taken when implementing this to ensure that no blind pores are generated when flipping the image (pores on the surface that are not connected to the main pore space must be removed).

3.2.2.2.5 Step by Step Algorithm

The algorithm for simulating Knudsen diffusion proceeds as follows:

1. Calculate the mean free path of the gas given a particular temperature and pressure using Eq. (3.29).
2. Compute the mean thermal velocity of each molecule relative to the other molecules is given by Eq. (3.30).

3. Given the velocity and the mean free path, calculate the time taken by each molecule to complete one mean free path:

$$\Delta t_{\text{mfp}} = \lambda / \bar{v} \tag{3.40}$$

4. Subdivide each mean free path into *steps_per_path* number of smaller segments, denoted N_{spp}. For the purposes of this thesis, each mean free path was split into 50 separate steps. Thus, the distance travelled by each walker at each time step is $d_{\text{step}} = \frac{\lambda}{N_{spp}}$ and $\Delta t_{\text{step}} = \frac{\Delta t_{\text{mfp}}}{N_{spp}}$
5. Initialize a 3D array of dimensions N_{walkers} x N_{steps} x 3 for recording each walker's Cartesian coordinates at each time step. This is referred to as the *path array.* Locations are saved in continuous space with coordinates of type *float.*
6. Initialize the domain with the N_{walkers} inserted into randomized locations in the pore space of the image. Record walker positions in the first slice of the *path array.*
7. Each walker is assigned an initial direction vector, which is saved in a separate 2D array of dimensions equal to (N_{walkers} x 3), referred to as the *direction array*.
8. A third array is initialized to keep track of the position of the last direction change for each walker, referred to as the *start_path array*.
9. Increment the time step by moving each walker by d_{step} in its vector direction.
10. Check if any walkers have exceeded their mean free path.
    a. Any walkers that have reached the end of their mean free path are assigned a new direction vector and their recorded location in the *start_path* array is updated to their previous legal position.
    b. Re-execute the step for these walkers.

11. Check if any walkers have collided with the wall.
    a. The locations of walkers in the *path array* are rounded to the nearest integer and the voxel at this coordinate location is checked in the domain. If it is a solid voxel, the walker has collided with a wall.
    b. They are assigned a new direction vector and their location in the *start_path* array is updated to their previous legal position. If the new direction vector results in another collision with the wall, a new vector is randomly assigned until all walkers have an acceptable trajectory.
    c. Re-execute the step for these walkers.
12. Once all collisions are resolved, record the new walker positions in the corresponding location of the *path array.*
13. Repeat steps 9-12 until reaching the end of the prescribed number of steps in the simulation. This should be set so that the calculated diffusion reaches a steady state.

3.2.2.3 Validation of Random Walk Method

In order to validate the random walk method described, two alternative calculations were used: 1) method proposed by Becker et al [100] which combines a Knudsen regime random walk with a finite volume method and 2) Combination of an analytical estimate of the Knudsen diffusivity with the finite volume method continuum diffusivity.

Method 1 has two parts: 1) a high Knudsen random walk where Knudsen effects dominate the continuum effects, and 2) a continuum diffusion simulation using a finite volume solver. The Knudsen-only and continuum-only diffusivities can be combined via the Bosanquet relation [105]:

$$D_{\text{eff}} = (D_{\text{Kn}}^{-1} + D_{\text{bulk}}^{-1})^{-1} \tag{3.41}$$

The Knudsen-only diffusivity can be obtained by using the random walk algorithm implemented in this thesis, but setting the mean free path to infinity, or some arbitrarily large number such that walkers only ever change direction due to collisions with the wall of the pore space; otherwise, they continue in a straight line forever. This simulates the extreme case of Kn $\gg$ 1 and is the value that diffusion will approach as $\lambda \to \infty$. The Knudsen effective diffusivity D_{Kn} is then calculated according to Eq. (3.31), where the time elapsed Δt_{total} is calculated as the total distance traveled d_{total} divided by the mean thermal velocity $\bar{v}$.

The continuum regime diffusivity is obtained by a finite volume method implemented in PuMA [128] which computes the pore phase tortuosity. The effective continuum diffusivity is obtained through by combining Eqs (3.33) and (3.36), using the bulk phase diffusivity of oxygen ($2.19 * 10^{-5} \frac{m^2}{s}$).

A second, analytical method of estimating the effective diffusivity is also employed. In this method, the Bosanquet relation is still used, and the continuum regime diffusivity is still obtained by the finite volume method, as in the Becker et al [100] validation method. The difference is the Knudsen regime diffusivity is calculated according to the analytical formula

$$D_{\text{Kn}} = \frac{1}{3} l_c \bar{v} \tag{3.42}$$

where l_c is the characteristic length, computed as the mean chord length of the domain using the chord length filter as described in section 3.1.2.3.

The two validation methods and the method used in this paper were used to estimate the transition regime effective diffusivity of artificially generated sphere packs with sphere radii of 25nm over a range of porosities.

Each sphere pack is of dimensions 100 x 100 x 100 voxels with a nominal resolution of 5nm. The results of these simulations are shown in Figure 3.17.

The pure random walk method (as implemented in this thesis) is shown in black circles and is in reasonable agreement with the Becker method (validation method 1, shown green pluses) at low porosities but the two begin to diverge at around 60% porosity. This can be explained by the use of the Knudsen only random walk in the Becker method: As the porosity increases, there is a smaller chance of molecule-wall collisions, so D_{Kn} quickly increases, enabling the effective diffusivity to quickly approach the bulk diffusivity. In comparison to validation method 2 (effective diffusivity predicted by analytical formula Eq (3.42), shown as blue triangles) at higher porosities, the random walk method predicts a higher diffusivity. While the results are not exactly same between each of the methods, they are sufficiently close that the author felt confident in the ability to use the implemented random walk method to obtain estimates of the structural effects on the effective diffusivity of the reconstructed catalyst layers.

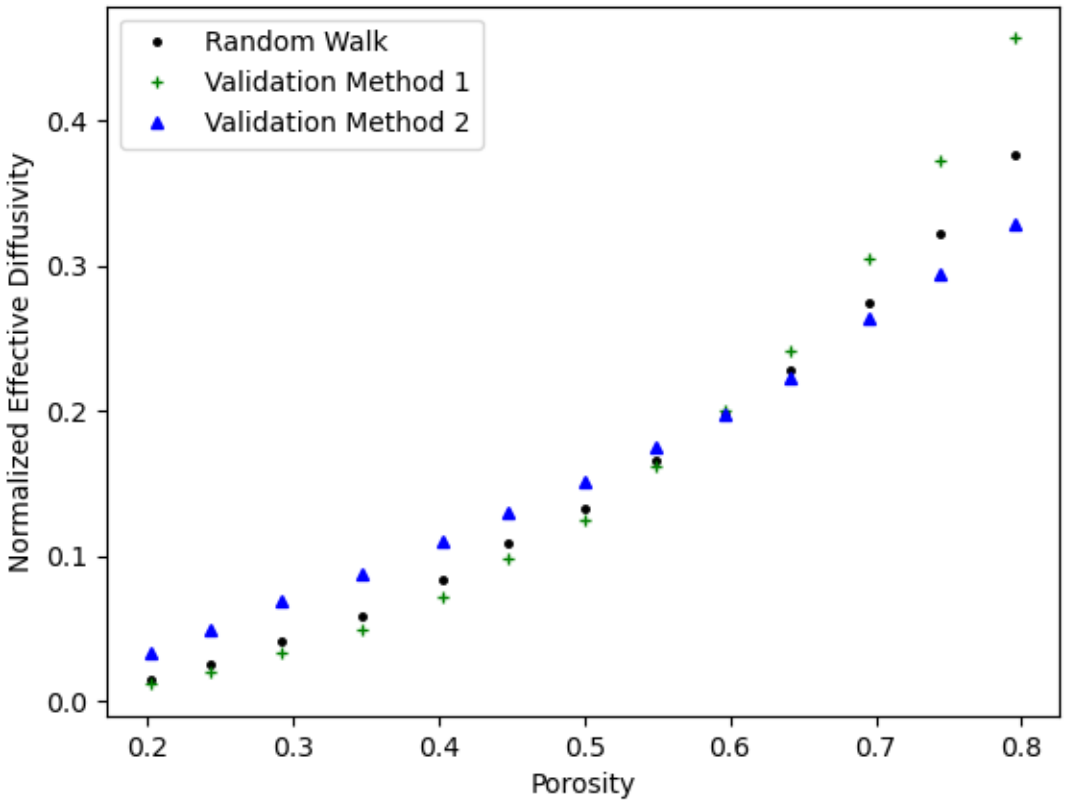


Figure 3.17: Computation of transition regime effective diffusivity for a domain of overlapping spheres of 25 nm radius at varying porosity, and three different computation methods.

3.3 Chapter Summary

This chapter has described the methodology used for the development and implementation of the reconstruction algorithms using sphere-based simulated annealing and various combinations of the n-point correlations. In the second part of the chapter, the methodology for pore-scale simulation of the reconstructions is presented. A finite volume method is used to solve for the effective thermal, ionic and electrical conductivities. A modified random walk method is presented which is used to simulate diffusion, while accounting for Knudsen effects.

Chapter 4

Results and Discussion

4.1 Reconstruction

Part 1 of the reconstruction section of this thesis is based on the work that was published in a paper in a special issue of Energy [139]. This section focuses on using the lineal path distribution in simulated annealing to optimize the structure of the PEMFC catalyst layer. Based on the results of this study, further investigation was carried out, and a new correlation function (the 2D analogue of the pore size distribution) was introduced to the simulated annealing process in order to attempt to further improve the accuracy. This extension of the original study is presented in Part 2 of the reconstruction section.

4.1.1 Part 1: Lineal-Path Distribution Reconstructions

Using the hybrid stochastic method outlined in Section 3.1.3, two 3D reconstructions were generated for each of Methods 0, 1 and 2. Each reconstruction method was used to 1) generate a reconstruction initialized with a variable radius (VR) distribution sphere pack as given by Eq. (3.11) and 2) generate a reconstruction from a sphere pack initialized with a constant radius (CR). Starting from a constant radius sphere pack is less ideal but is presented to demonstrate the effectiveness of the simulated annealing procedure more clearly. Note that for both the VR and CR cases, once simulated annealing begins, the sphere size is allowed to increment or decrement as described in Figure 3.15.

For convenience, a naming convention is used in Part 1 which begins with either *VR* or *CR*, for reconstructions which start with the variable radius sphere pack or the constant radius sphere pack, respectively. This is followed by the simulated annealing *Method* number used to generate the final structure

(0,1, or 2), as defined in section 3.1.3.3. For example, VR2 is the reconstruction initialized with the variable radius sphere pack, then optimized with simulated annealing *Method 2.* CR1 is the reconstruction initialized with a constant radius sphere pack and optimized with simulated annealing *Method 1.* Table 4.1 summarizes the notation used for the reconstructions involved in Part 1 of this thesis.

Table 4.1: Notation used for Part 1 reconstructions.

Reconstruction label	Correlations Optimized	
CR0	-	Reconstruction initialized with a constant radius sphere pack using **Method 0** (sphere-packing only, no simulated annealing)
CR1	S_2	Reconstruction initialized with a constant radius sphere pack using **Method 1** (S_2 optimized, starts from CR0)
CR2	S_2, L^v, L^s	Reconstruction initialized with a constant radius sphere pack using **Method 2** (S_2, L^v, L^s optimized, starts from CR0)
VR0	-	Reconstruction initialized with a variable radius sphere pack using **Method 0** (sphere-packing only, no simulated annealing)
VR1	S_2	Reconstruction initialized with a variable radius sphere pack using **Method 1** (S_2 optimized, starts from VR0)
VR2	S_2, L^v, L^s	Reconstruction initialized with a variable radius sphere pack using **Method 2** (S_2, L^v, L^s optimized, starts from VR0)

Figure 4.1 shows the sphere size distributions for each of the six generated reconstructions. In the VR cases, the sphere size distribution changes only slightly over the course of simulated annealing, most likely because the starting distribution is skewed and allowing an equal chance of increasing or decreasing the radius on each iteration would eventually shift it towards a Normal distribution. Indeed, when observing the CR case where

the starting configuration uses a uniform initial sphere radius of 7 voxels or 35nm, the final sphere size distribution resembles a slightly left-skewed Normal distribution.

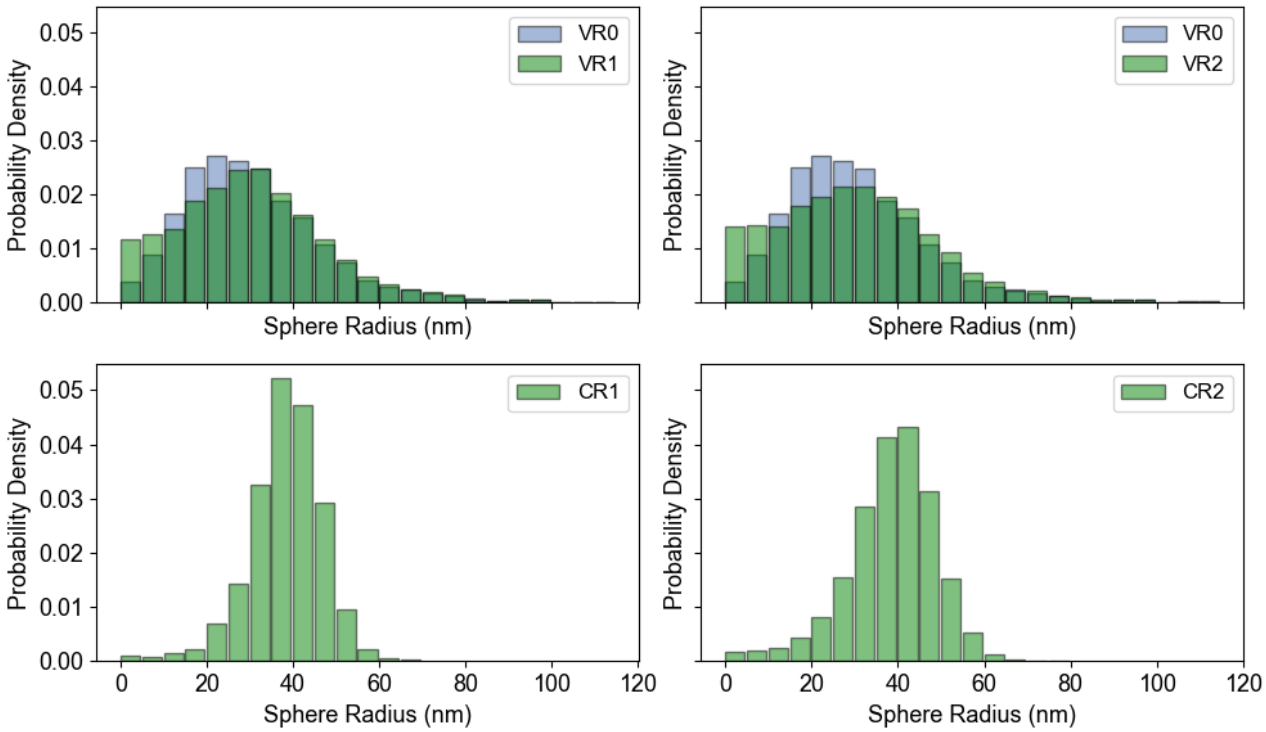


Figure 4.1: Distribution of spheres inserted into the reconstructions for each of the reconstructions. CR0 and VR0 are shown in blue and in each plot because they are the starting structure for the reconstructions using Methods 1 and 2.

The simulated annealing reconstruction was conducted with the parameters listed in Table 4.2. One of the 2D images out of the 113-image reference image set and 2D *slices* taken from each of the 3D reconstructions are displayed in Figure 4.2. 2D images are displayed to better observe the morphological effects of each of the reconstruction methods as these effects can be difficult to discern from a 3D visualization, an example of which is displayed in Figure 4.3. An observable increase in the amount of clumping has occurred in the Method 2 reconstructions (Figure 4.2 D,G), which is attributed to the inclusion of the lineal path distribution. The

clumping effect is most noticeable in the change from CR0 (Figure 4.2E) to CR2 (Figure 4.2G). CR0 begins with uniform radius spheres randomly distributed throughout the model domain. CR2 shows that many of the spheres have clumped into larger agglomerates which also results in larger pore spaces being formed.

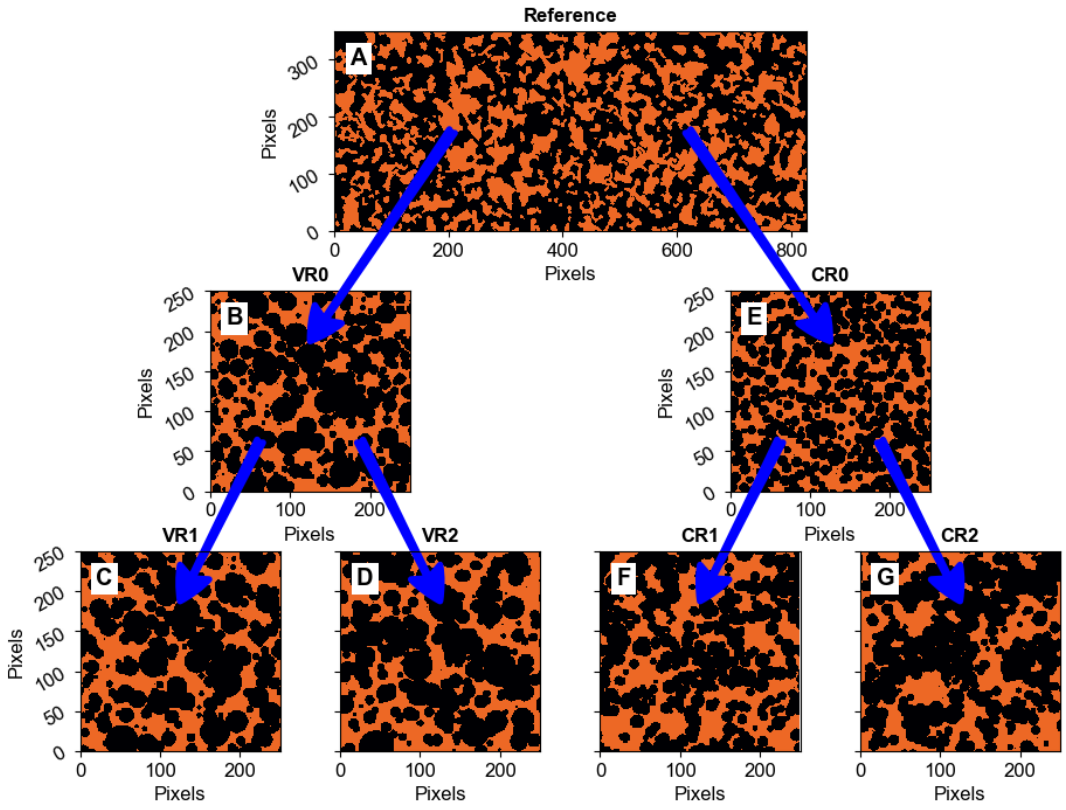


Figure 4.2: Sectioned (30th) slice of each of the reconstructions. A) full reference FIB-SEM image [18]**; B, C, D) Reconstructions using a variable radius initial sphere pack; E, F, G) Reconstructions starting from a constant radius sphere pack. Black and orange colors represent the solid and void phases, respectively.**

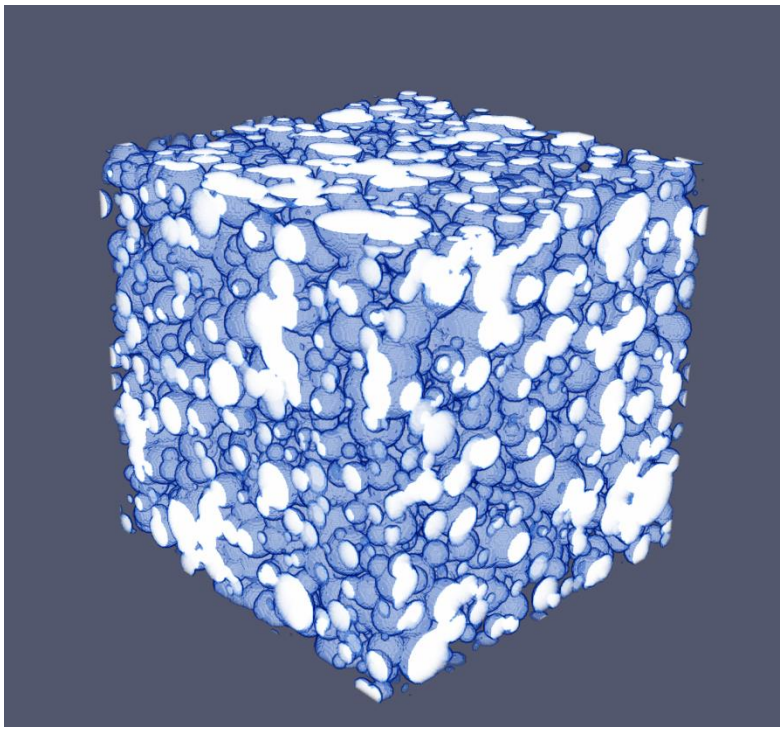


Figure 4.3: Reconstruction 3D visualization of the solid phase for VR2.

Table 4.2: Parameters used in the reconstruction of the PEMFC CL

Parameter	Symbol	Value
Reconstruction domain size (voxels)	L	250
Reconstruction voxel resolution (nm)	res	5
Maximum total sphere trial movement per Markov chain	N_T	2000
Maximum successful sphere trial movement per Markov chain	N_{MC}	1000
Max number of Markov chains before program termination	X_{MC}	100
Energy convergence criterium	δ	0.00001
Temperature scaling factor	λ	0.9
Sphere overlap tolerance	ϵ_v	0.5
Total iterations	N	200000

4.1.1.1 Two Point Correlation and Lineal Path distribution

The correlations S_2, L^v and L^s are shown in Figure 4.4 and Figure 4.5 for the VR and CR reconstructions, respectively. For each plot, the correlation is shown for the reference structure, the initial sphere pack structure (the structure obtained by Method 0), and the final structure after simulated annealing. The relative error is plotted as the black dotted line and calculated as the absolute difference between each point in the final reconstruction and the corresponding point in the reference distribution, relative to the maximum value of the reference distribution.

Similar trends can be seen across the different reconstructions. In the two-point correlation (Figure 4.4 and Figure 4.5, plots A and B), minimal error is observed in the final reconstruction for both Methods 1 and 2. Additionally, even VR0, which did not undergo simulated annealing, shows good agreement with the reference two-point correlation. This is attributed to the prior optimization of the sphere radius distribution whereas the

initial structure CR0 severely underestimates the reference two-point correlation because it did not undergo optimization prior to simulated annealing. However, despite the non-ideal start, the simulated annealing algorithm is able to effectively optimize the two-point correlation for both CR1 and CR2. This emphasizes both the effectiveness of the simulated annealing procedure and that if no simulated annealing is to be done, the use of an accurate sphere size distribution is critical.

Achieving good convergence of the lineal path distributions (Figure 4.4 and Figure 4.5, plots C, D, E and F) is slightly more difficult than with the two-point correlation. A small spike in error is present at the first or second points of L^v and L^s for all reconstructions. As small values of distance correspond to the short-range characteristics of the structure, this effect can be attributed to using a sphere-based system to reconstruct an image which is not inherently spherical in nature. Further improvement in voxel resolution or the inclusion of more spheres with smaller diameters may resolve this issue, but this then runs into exponentially increasing computational costs (as the average sphere radius approaches 1 voxel, the model stops being a sphere-based reconstruction and becomes the conventional voxel-based simulated annealing). When comparing the convergence of the lineal path distributions for *Method 1* and *Method 2,* it is evident that the maximum value of the error spike is higher in reconstructions using *Method 1* than those using *Method 2.* This is expected since *Method 2* directly optimizes L^v and L^s, but it also indicates that *Method 2* produces a more accurate structure.

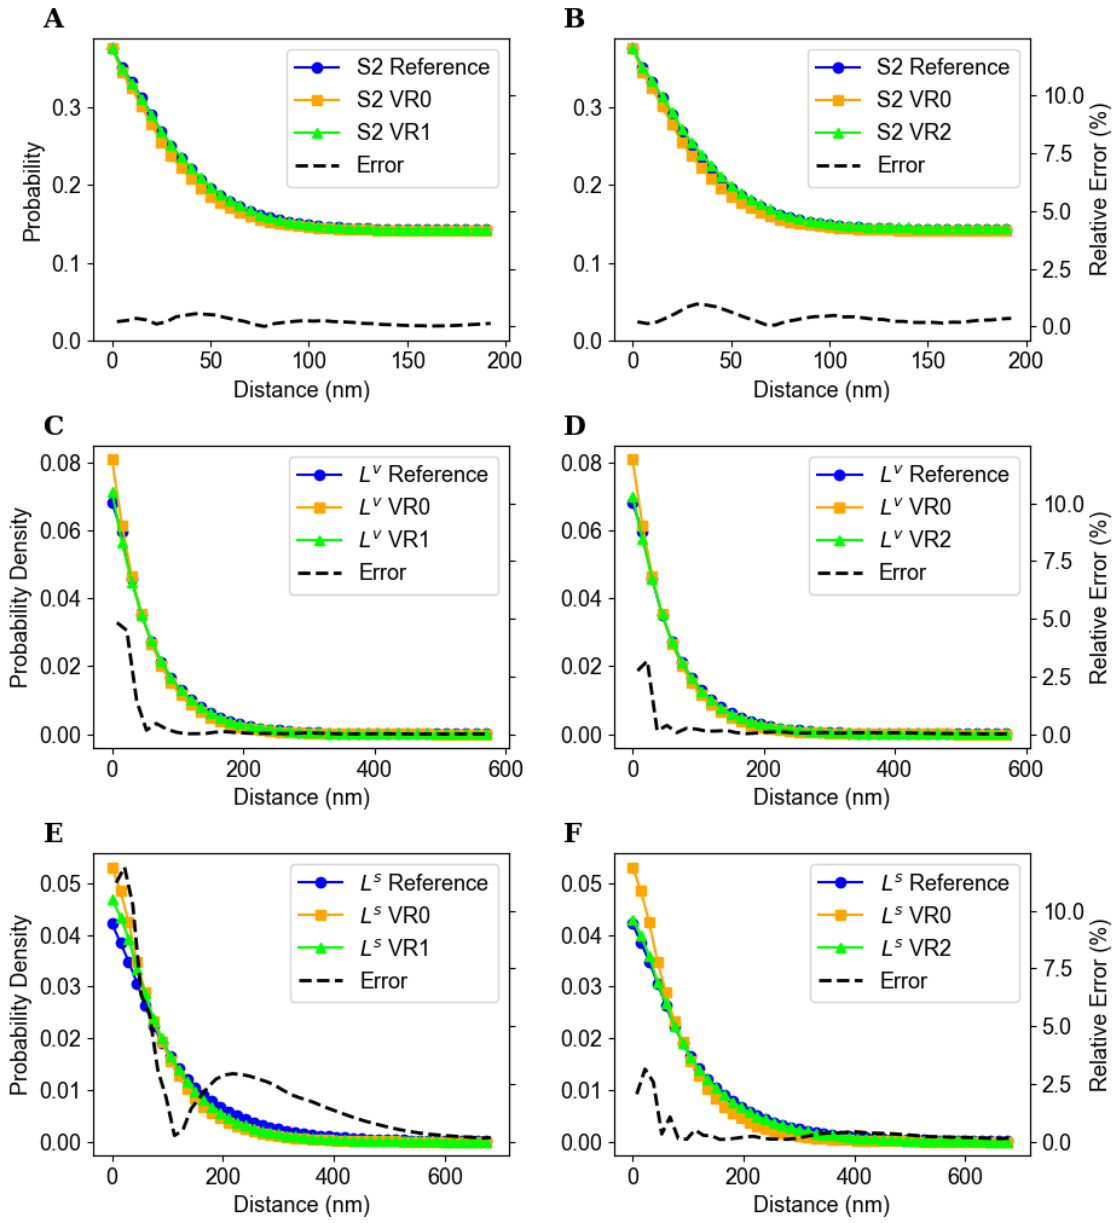


Figure 4.4: Two-point correlation, lineal path distribution void phase and lineal path distribution solid phase for variable radius reconstructions. VR1 is pictured on the left (A, C, E) and VR2 is pictured on the right (B, D, F).

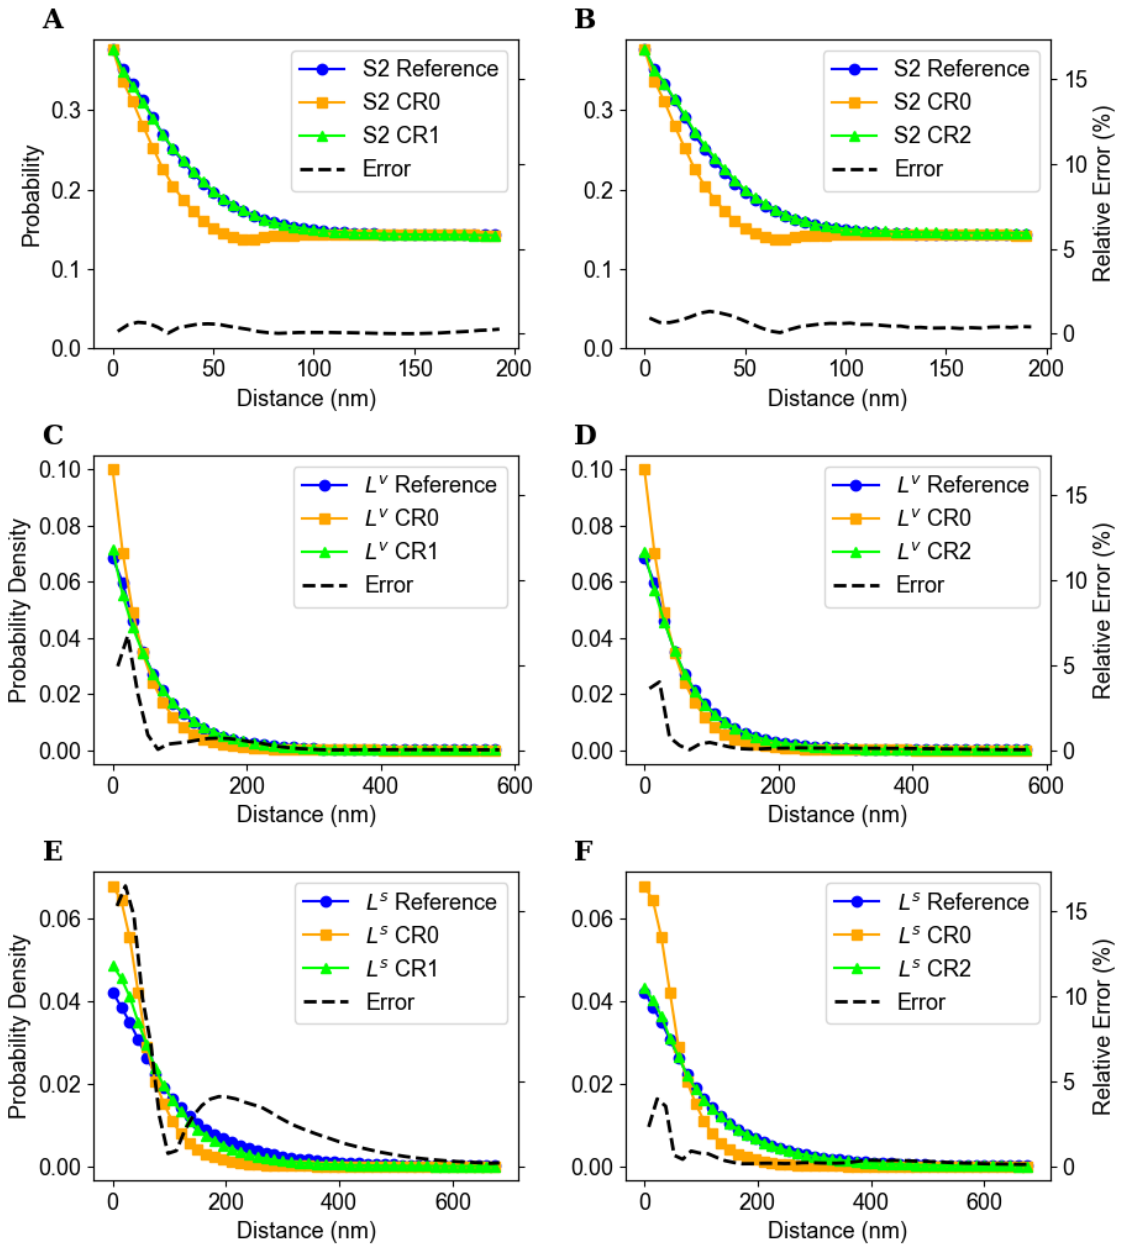


Figure 4.5: Two-point correlation, lineal path distribution void phase and lineal path distribution solid phase for constant initial radius reconstructions. CR1 is pictured on the left (A, C, E) and CR2 is pictured on the right (B, D, F).

4.1.1.2 Comparison of pore size distributions

The reconstruction results were validated by comparing the pore size distributions (PSD) of the reconstruction to that of the experimental reference image. The PSD was computed using the local thickness filter in the python module *porespy* [55] as discussed in section 3.1.2.4. Unlike the two-point correlation and the lineal path distribution, the PSD cannot be compared directly between 2D and 3D realizations since it is harder to fit a 3D sphere into a space where a flat circle may fit (for this reason, the 2D PSD is shifted to higher values than the 3D PSD). Due to the low resolution in the through-plane direction of the reference image, the 3D PSD of the reference was unsuitable for direct comparison. Therefore, the 2D PSDs were computed for each 2D image in the reference image set and for each slice in the z-direction of the reconstruction, then averaged across all slices.

Figure 4.6 shows the PSDs for each of the reconstructions, overlaid on the PSD from the reference. Qualitatively, the PSDs agree with each other quite well, though the range of pore sizes in the reconstructions is narrower than the reference indicating that it is more difficult to accurately reconstruct the larger pore sizes. This may be because the two-point correlation is considered a short-range correlation, and the lineal path distribution at best describes medium-range structural characteristics. If this is the case, accurate reconstruction of the full range of pore sizes would require a correlation describing long-range structural connectedness in the simulated annealing. In order to be effective, this correlation would also require an efficient computational implementation and should be dimensionally agnostic (the 2D correlation should be equivalent to the 3D correlation). Another option may be to introduce additional fabrication-based limitations in the movement of spheres during the reconstruction to artificially create and maintain larger empty pore spaces.

Interestingly, while both *Method 1* and *Method 2* simulated annealing improved the pore size distribution in both the VR and CR cases, the error in the reconstructions initialized with a constant sphere radius (CR1 and

CR2) is less than the error in reconstructions initialized with the optimized variable sphere size distribution in VR1 and VR2, despite having started from a less accurate structure. Recalling that the sphere size distribution used to initialize the VR reconstructions was obtained from a local thickness analysis of the *solid* phase in the reference, this underlines the difficulty of controlling the pore size distribution by adjusting means of the solid phase sizes. Additionally, it is noted that *Method 1* performed better than *Method 2* in obtaining an accurate pore size distribution, which may be counter intuitive, however it can be attributed to the clumping effect that is observed when optimizing for the lineal path distribution and is even visually evident by observing the actual structure in Figure 4.2D and Figure 4.2G.

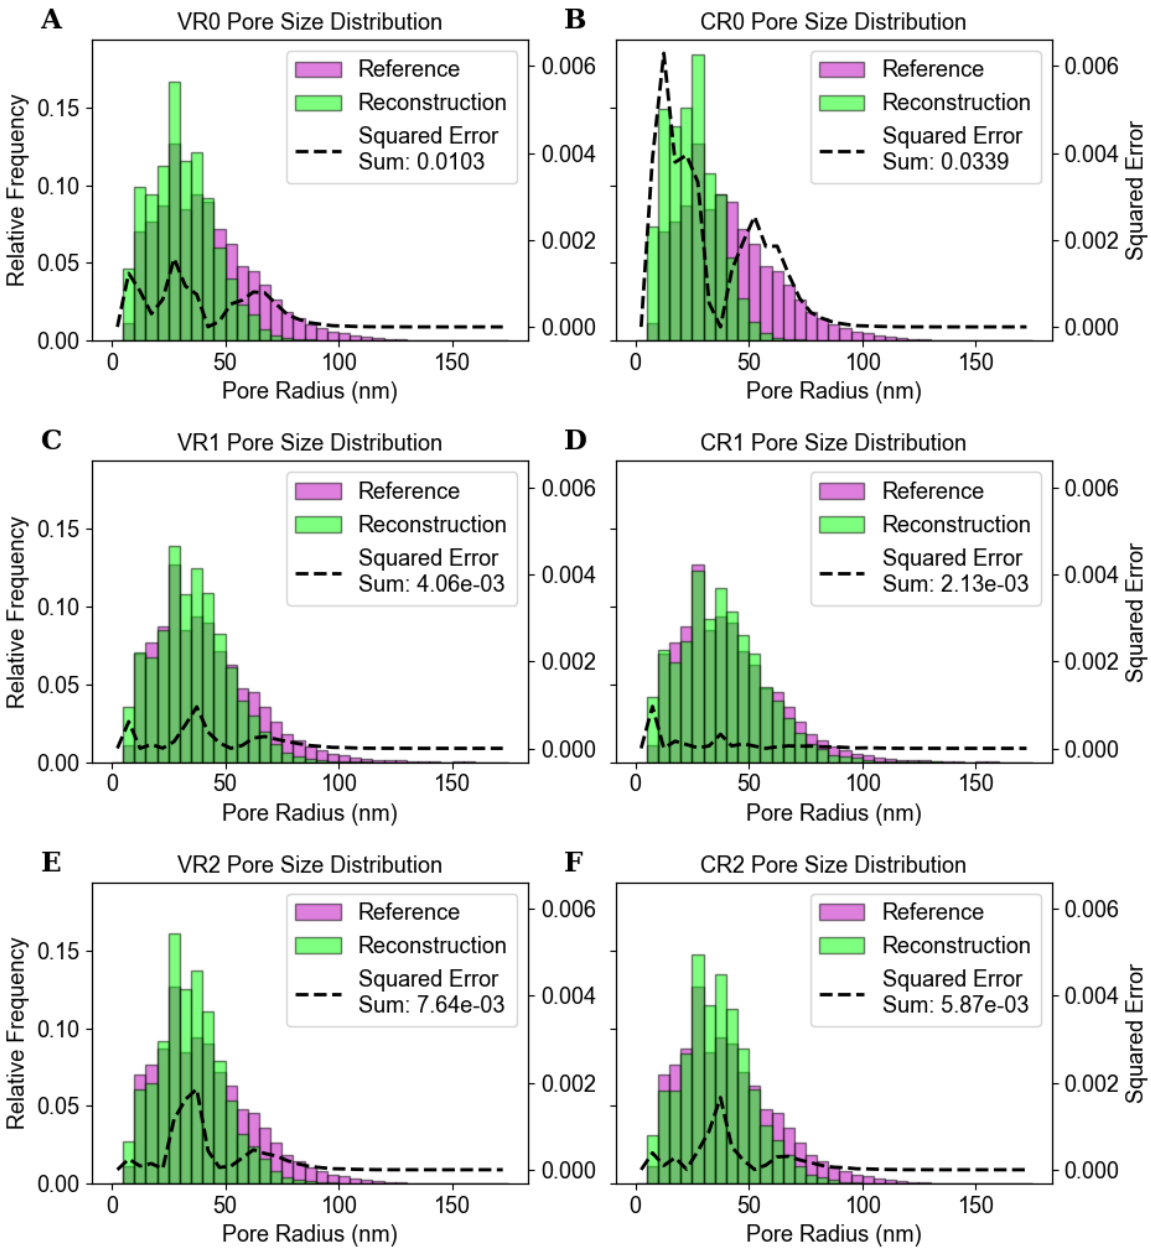


Figure 4.6: Pore Size Distributions for each reconstruction. Reconstructions are shown in green and the reference is shown in purple. The squared error is represented by the dotted black line and the sum of squared errors is shown in the legend.

4.1.1.3 Porosity

Figure 4.7 shows the porosity of the reference image and each of the reconstructions. While the porosity of the initialized sphere packs VR0 and CR0 is not exactly the same as the reference, they are already extremely close. After simulated annealing, and due to the inclusion of the porosity directly in the energy cost function, the porosity from the Method 1 and Method 2 reconstructions is virtually equal to that of the reference, down to a difference of a few voxels. The calculation of the porosity is very fast, so that there is very little reason not to include it in sphere-based simulated annealing, however to the author's knowledge this has not been done yet.

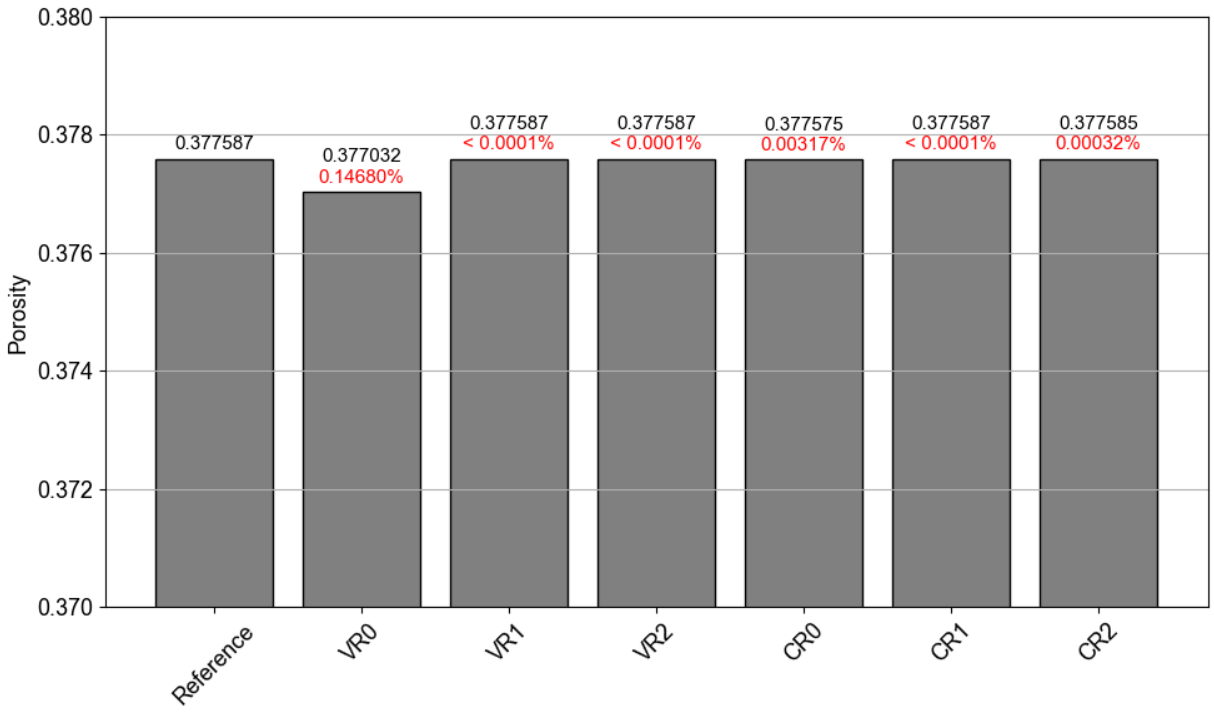


Figure 4.7: Porosities of the reference image and each reconstruction.

4.1.1.4 Specific surface area

Figure 4.8 shows the specific surface areas for the reference image and the reconstructions. Experimentally derived estimates for the specific surface area of PEMFC catalyst layer reported in the literature using the method of standard porosimetry range from 0.025-0.036 nm^2/nm^3 [49], which is close to the value found for the reference image, indicating that the marching cubes algorithm is a valid metric of the surface area. It also should be emphasized that the values from [49] are from a separate study than those from the reference image used in this study, which was taken from [18], and as such, variations in materials used, sample preparation and testing may be a reason for the difference between the reference image surface area (0.023 nm^{-1}) and those reported in the literature. More importantly to the work in this thesis however, is the ability of simulated annealing to match the reference image structure. The surface area is simply a metric by which the surface area of the reference can be compared to that of the reconstruction, and in this study is used as a means of model validation and to determine the relative effectiveness of each of the reconstruction techniques.

Good agreement was found for the surface area, particularly when comparing reconstructions using Method 1 and Method 2 to the reference image. This is expected since the specific surface area is directly related to the two-point correlation by Equation (2.7) and the two-point correlation is optimized in both of these methods. Despite having the same porosity as the reference, the surface area of both initial sphere packs VR0 and CR0 are significantly different from the reference, exhibiting errors of 19% and 52%, respectively. Meanwhile, the errors of the reconstructions using simulated annealing are much smaller and range from 2% to 6%. This is yet another indication that simulated annealing is required to generate an accurate structure rather than relying on a solely fabrication-based method.

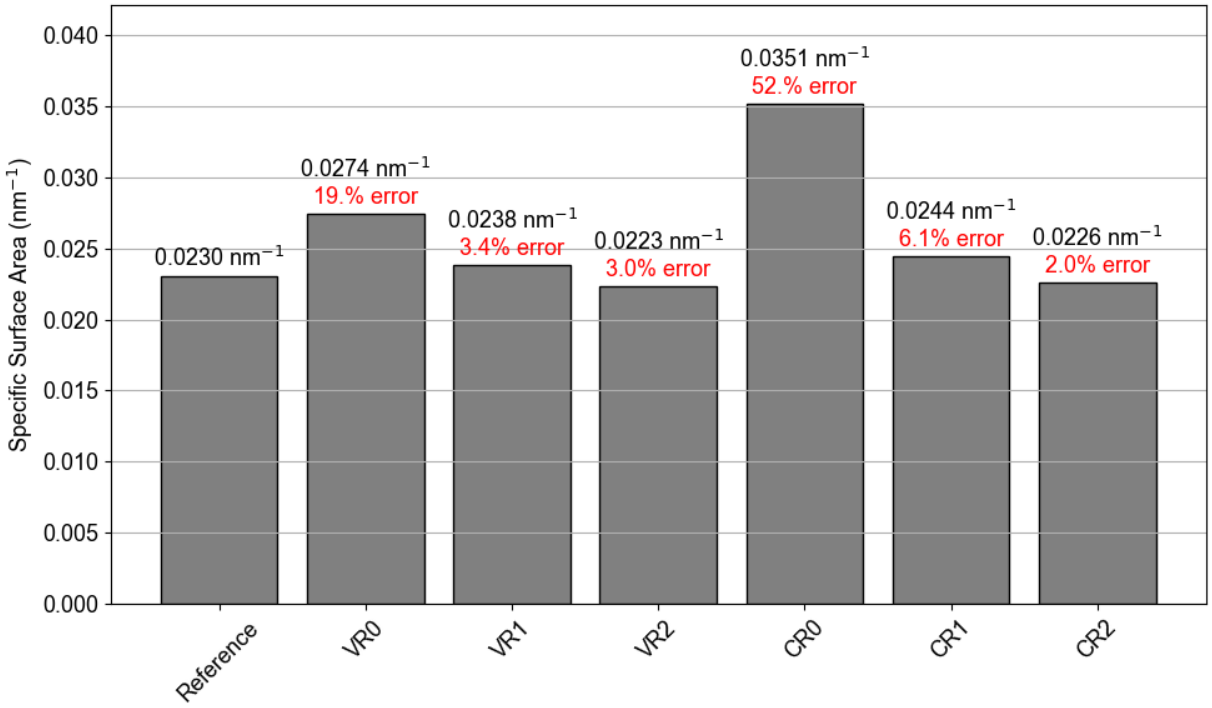


Figure 4.8: Specific surface area for the reference image and each of the reconstructions, computed using the marching cubes algorithm implemented in porespy [55]**.**

4.1.1.5 Mean Pore and Solid Phase Diameter

The mean pore and particle sizes, defined here as the spherical radii computed using the local thickness filter in *porespy* [55], are pictured in Figure 4.9 and Figure 4.10,, respectively. While an improvement in accuracy is seen after conducting simulated annealing (VR1, VR2, CR1 and CR2) over the non-optimized sphere pack (VR0 and CR0), the mean pore diameter consistently underestimates that of the reference. This is consistent with conclusions drawn from analysis of the pore size distributions in Section 4.1.1.2. On the other hand, the mean agglomerate particle diameter overestimates that of the reference. This can be attributed to the spherical

nature of the reconstruction and the method of computing the local thickness. Recalling that the local thickness filter used to obtain the pore size distributions works by inserting *inscribed* spheres into the domain, and the reconstruction solid phase is made up of spheres, so that the inscribed spheres will fit perfectly within the solid phase. Meanwhile, the pore phase will contain sharper intersections and corners where two solid spheres intersect, and the inscribed sphere in these locations will be smaller. On the other hand, neither the solid nor the pore phase of the reference image is spherical in nature, so the inscribed spheres from the local thickness filter will not be as biased in favour of the solid phase and a more balanced result is observed.

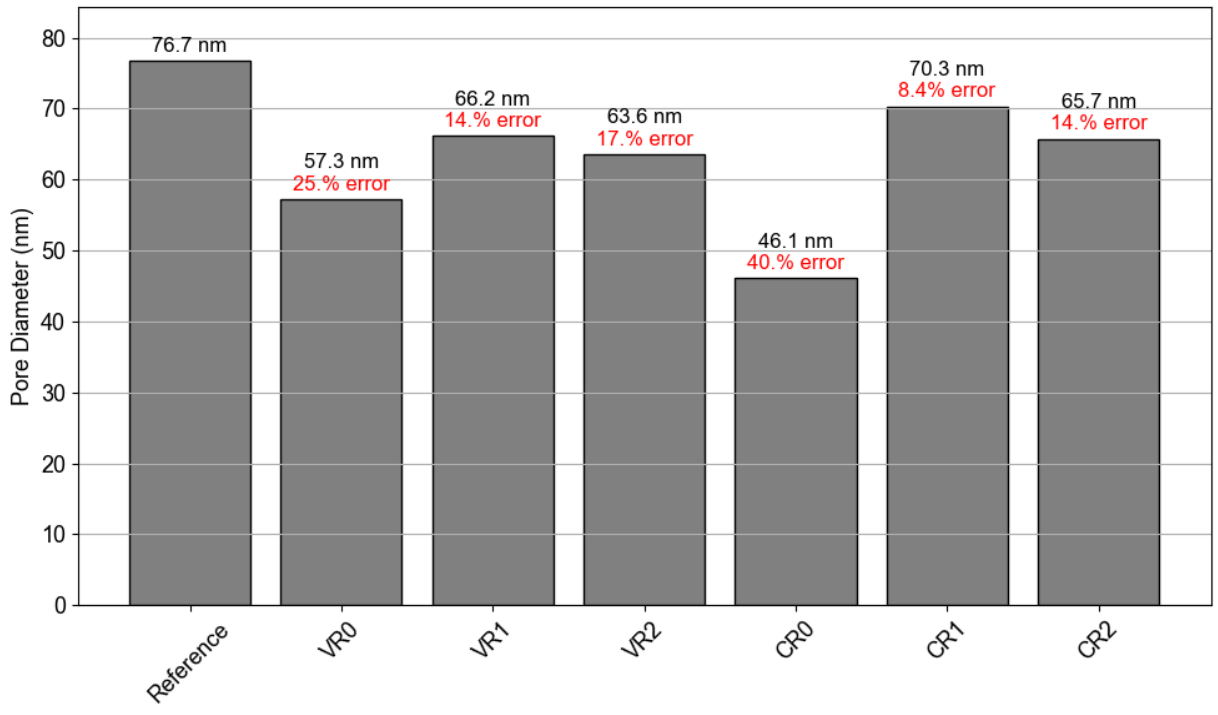


Figure 4.9: Mean pore diameter for the reference and reconstructions.

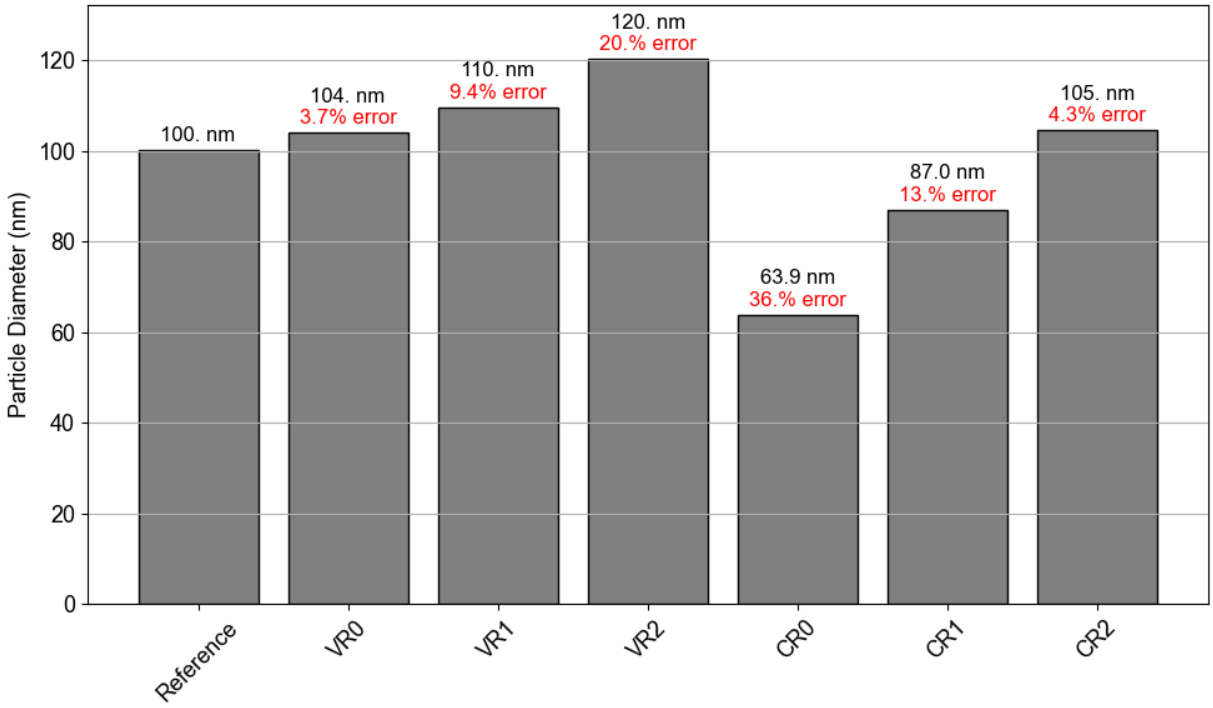


Figure 4.10: Mean particle diameter for the reference and reconstructions.

4.1.1.6 Mean Chord Lengths

The mean chord length is an alternative to the mean diameter as a measurement of the characteristic size of a porous structure. The mean chord lengths extracted using the chords image filter for the pore and solid phases are pictured in Figure 4.11 and Figure 4.12, respectively. The mean chord lengths in the *Method 1* and *Method 2* reconstructions demonstrate good agreement with those of the reference, where *Method 2* achieves slightly more accurate reconstructions than *Method 1*. This is expected because the mean chord length is related to the lineal path distribution, which was optimized in *Method 2* reconstruction. It is unsurprising therefore that VR2 has mean chord lengths closest to those of the reference.

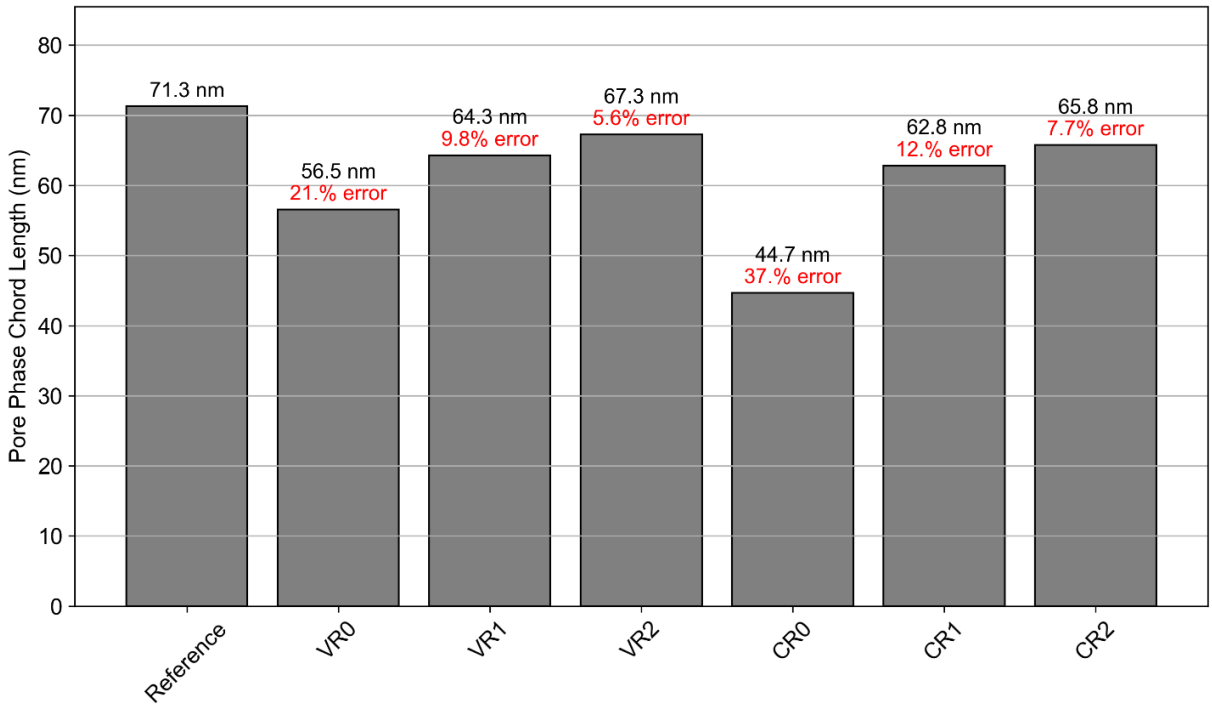


Figure 4.11: Mean chord length for the pore phase of the reference and reconstructions, computed using the chord image filter.

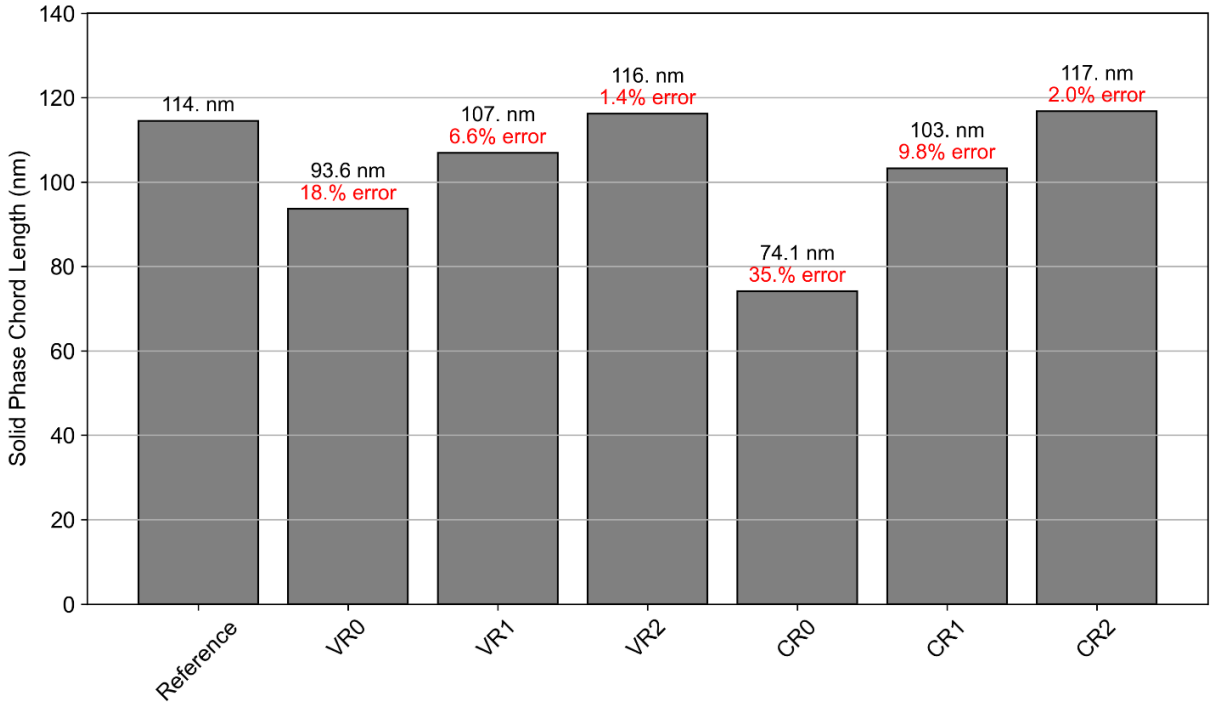


Figure 4.12: Mean chord length for the solid phase of the reference and reconstructions, computed using the chord image filter.

4.1.1.7 Summary

In this Part 1 of the reconstruction part of this thesis, an image-based hybrid stochastic reconstruction model for the catalyst layers in proton exchange membrane fuel cells was developed by incorporating the fabrication-based sphere-packing reconstruction method with a simulated annealing method that optimizes the two-point correlation and the lineal path distribution. The algorithm uses a two-step approach, first employing a sphere-packing followed by simulated annealing, benefiting from both the speed of the sphere-packing approach and the flexibility of simulated annealing. The reconstructions were initialized using either a variable sphere size distribution extracted from a focused-ion bean scanning electron microscope (FIB-SEM) image or the less realistic uniform sphere size, and these two methods are compared. The second step uses sphere-based

simulated annealing and optimizes the porosity, the two-point correlation and the lineal path distributions for solid and pore phases such that the final structure exhibits similar interconnectedness, surface area and porosity as the reference image. This algorithm *(Method 2*) is compared with the conventional sphere-based simulated annealing which only optimizes the two-point correlation (*Method 1*) as well as a purely fabrication-based reconstruction method (Method 0). Six 3D catalyst layer reconstructions were generated from a 2D experimental FIB-SEM image: For each of the three reconstruction methods, one reconstruction was generated using a variable radius initial sphere pack and one with a constant radius sphere pack. The resulting pore size distributions, porosities, specific surface areas, mean pore diameters and mean chord lengths were compared between the reconstructions in order to validate the model and determine the effectiveness of including the lineal path distribution in the simulated annealing algorithm. Both *Method 1* and *Method 2* optimized the porosity to within approximately 0.0001% of the reference porosity due to its inclusion in the energy cost function in both Methods. *Method 2* was able to achieve higher accuracy than *Method 1* or *Method 0* in the mean chord lengths and specific surface area. For the mean pore size, both *Method 1* and *Method 2* showed an improvement over Method 0, but neither was very accurate, though *Method 1* showed slightly less error than *Method 2*. This is attributed to the clumping effect of the lineal path distribution and future work should consider methods to directly optimize the pore size distribution or to limit this effect.

4.1.2 Part 2: Analysis of PSD-Based Reconstructions

This section of the reconstruction chapter presents an extension of the work in Part 1 to include simulated annealing optimization of the pore and particle size distributions. The study in Part 1 determined that use of the lineal path distributions to optimize the structure did a good job optimizing the chord lengths, surface area and porosity, but it struggled to reconstruct a structure with the correct radial pore and particle sizes. In Part 2, reconstructions are presented which directly optimize the pore and particle size distributions.

4.1.2.1 Notation

Results from using the lineal path distribution to optimize the structure of the PEMFC CL indicated that an excessive amount of clumping was occurring and while the results were improved over Method 1 where only the two-point correlation is optimized, the mean pore size and particle size were not very accurate.

For this reason, the fast, medial axis-based implementation of the pore size distribution (described in section 3.1.2.4.1) was developed and employed directly in the simulated annealing process. Three additional reconstruction methods incorporating combinations of the pore size distribution with the two-point correlation and lineal path distributions were used to generate reconstructions which were then compared to the variable radius reconstructions using Methods 1 and 2. The notation for these reconstructions is listed in Table 4.3. Note that for this section, the reconstructions with constant radius initialization have been dropped to reduce bloat.

In this section, each of the reconstructions for each method were generated three times with identical input parameters, but a different random starting configuration so that a standard deviation for the metrics can be obtained to determine the variability in the outcome of the reconstruction algorithm.

Unless otherwise noted, reported metrics are the average values across each of the three reconstructions generated with the same method and the reported measure of *sample* error is the standard deviation, while the *relative* error refers to the relative difference between the reconstruction metric and the corresponding metric from the reference image.

Table 4.3: Notations for reconstruction methods for PSD-based reconstruction analysis.

Label	Correlations Optimized	Description
M0	None	Reconstruction initialized with a variable radius sphere pack using **Method 0** (sphere-packing only, no simulated annealing)
M1	S_2	Reconstruction initialized with a variable radius sphere pack using **Method 1** (S_2 optimized, starts from VR0)
M2	S_2, L^v, L^s	Reconstruction initialized with a variable radius sphere pack using **Method 2** (S_2, L^v, L^s optimized, starts from VR0)
M3	PSD, PaSD	Reconstruction initialized with a variable radius sphere pack using **Method 3** (PSD, PaSD optimized, starts from VR0)
M4	S_2, PSD, PaSD	Reconstruction initialized with a variable radius sphere pack using **Method 4** (S_2, PSD, PaSD optimized, starts from VR0)
M5	S_2, PSD, PaSD, L^v, L^s	Reconstruction initialized with a variable radius sphere pack using **Method 5** (S_2, PSD, PaSD, L^v, L^ss optimized, starts from VR0)

4.1.2.2 Convergence

Each of the reconstructions were run until they reached 200000 total iterations (including failed and successful configuration changes). The configuration energy (normalized by the starting configuration energy for each method) is shown as a function of iteration number in Figure 4.13. It is evident that the reconstructions involving the PSD in the simulated annealing process (M3, M4 and M5) are slower to converge than M1 or M2. This is expected as these methods are optimizing correlations which are assumed to encode a higher quantity of structural information.

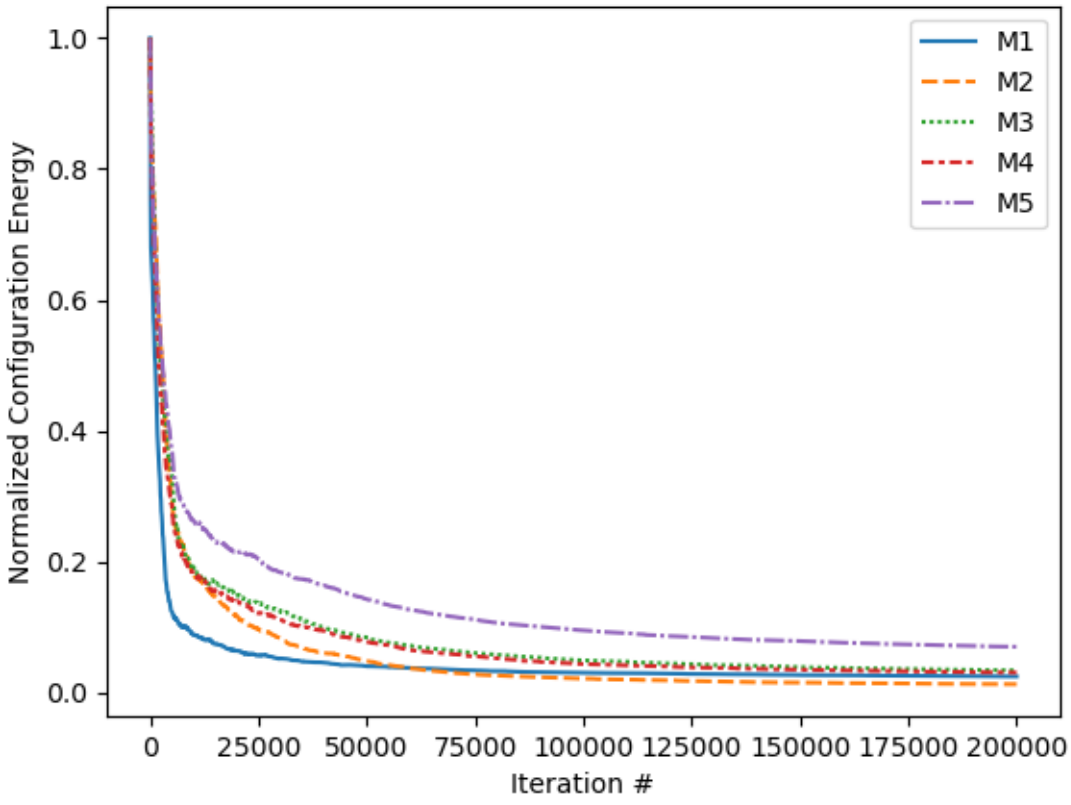


Figure 4.13: Normalized configuration energy as a function of reconstruction iteration number. The energy is normalized by the initial configuration starting energy for each method.

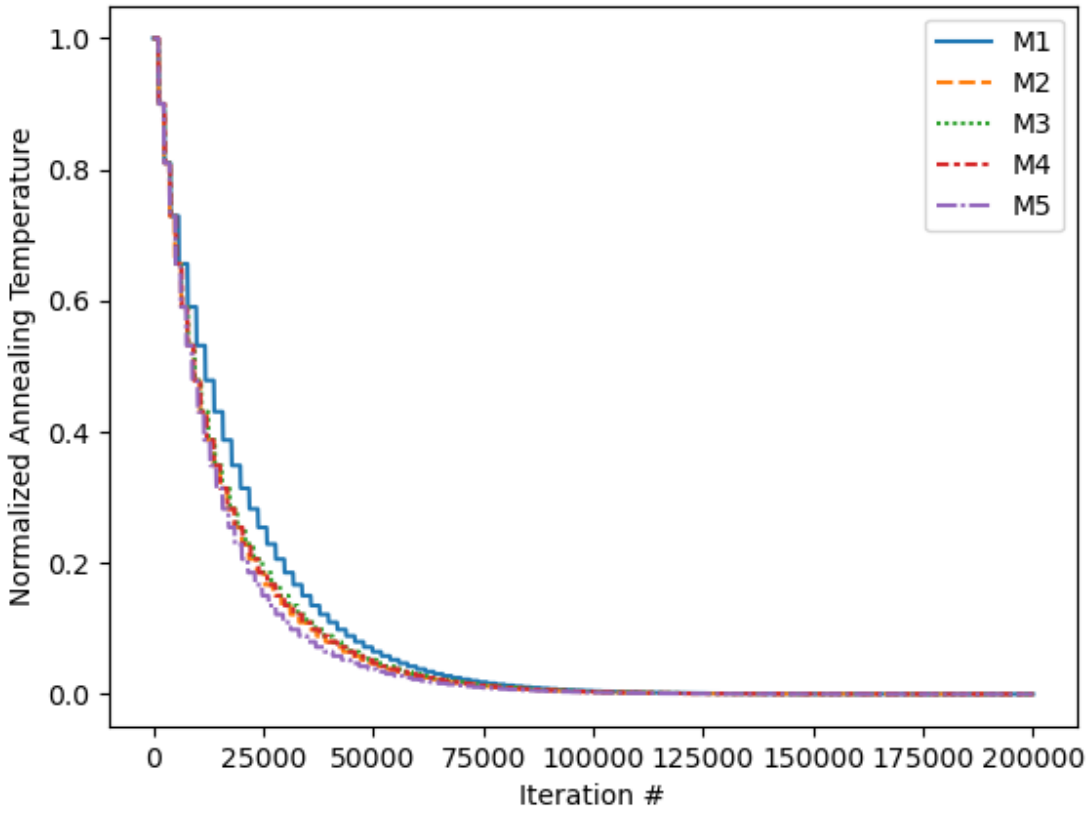


Figure 4.14: Annealing Temperature as a function of iteration number.

4.1.2.3 Timings

The simulation computation time per iteration and total simulation time for each of the reconstruction methods is listed in Table 4.4. Unsurprisingly, inclusion of the pore size distribution in simulated annealing significantly increases the computation time as compared to Methods 1 and 2. The computation of the two-point correlation using the fast Fourier transform method is very fast, in particular since it is implemented on the GPU. Likewise, the lineal path distribution in Method 2 is computed on the GPU and so is quite quick as well, adding only 0.05s per iteration to include it as compared to Method 1. Computation of the pore and particle size

distributions significantly increases the computation time by up to a factor of 4, but as is shown in the following sections, it also significantly increases the accuracy of the reconstructions.

Table 4.4: Simulation computation time for each reconstruction method.

Method	Time per iteration (s)	Total reconstruction time (s) (200000 SA iterations)
M1	0.2	40000
M2	0.25	50000
M3	0.37	74000
M4	0.75	150000
M5	0.85	170000

4.1.2.4 Comparison of structural correlations

This section will discuss how well each of the reconstruction methods is able to generate a structure with the same two-point correlation, solid/pore phase lineal path distributions and pore/particle size distributions. In each comparison, the sum of squared errors of the distributions is used as the measure of difference between the reconstruction and the reference.

Figure 4.15 shows the two-point correlations for each of the reconstructions. The squared error in the two-point correlation for each of the reconstructions is shown as a dotted black line, while the sum of the squared errors (SSE), which is the integral of the error line, is displayed in the legend. M0, which does not use simulated annealing at all, has the highest error of the reconstructions, but even so, the SSE is relatively low as compared to the other distributions (which will be discussed shortly). Given that the starting error is low, it is remarkable

that all reconstruction methods are able to reduce this error by a factor of 5 or higher. This includes M3, in which the two-point correlation is not included in the simulated annealing energy computation. The ease of optimization of the two-point correlation further emphasizes the need to include other correlations in the simulated annealing process.

Figure 4.16 shows the pore phase lineal path distributions for each of the reconstructions and Figure 4.17 shows the solid phase lineal path distributions for the same. In both cases, the least error is in the M2 reconstruction method, which is expected as this directly optimizes both lineal path distributions but does not consider the pore or particle radial sizes. Additionally, in every reconstruction for both the pore and solid phases, an error spike is observed at the small distances (<25 nm). In particular, in the solid phase, the reconstructions are overestimating the probability of finding these distances. This can be attributed to the spherical nature of the reconstructions. Use of spheres as the only structuring element means that the short distance characteristics of the reconstructions will be very uniform which is likely the cause of these error spikes.

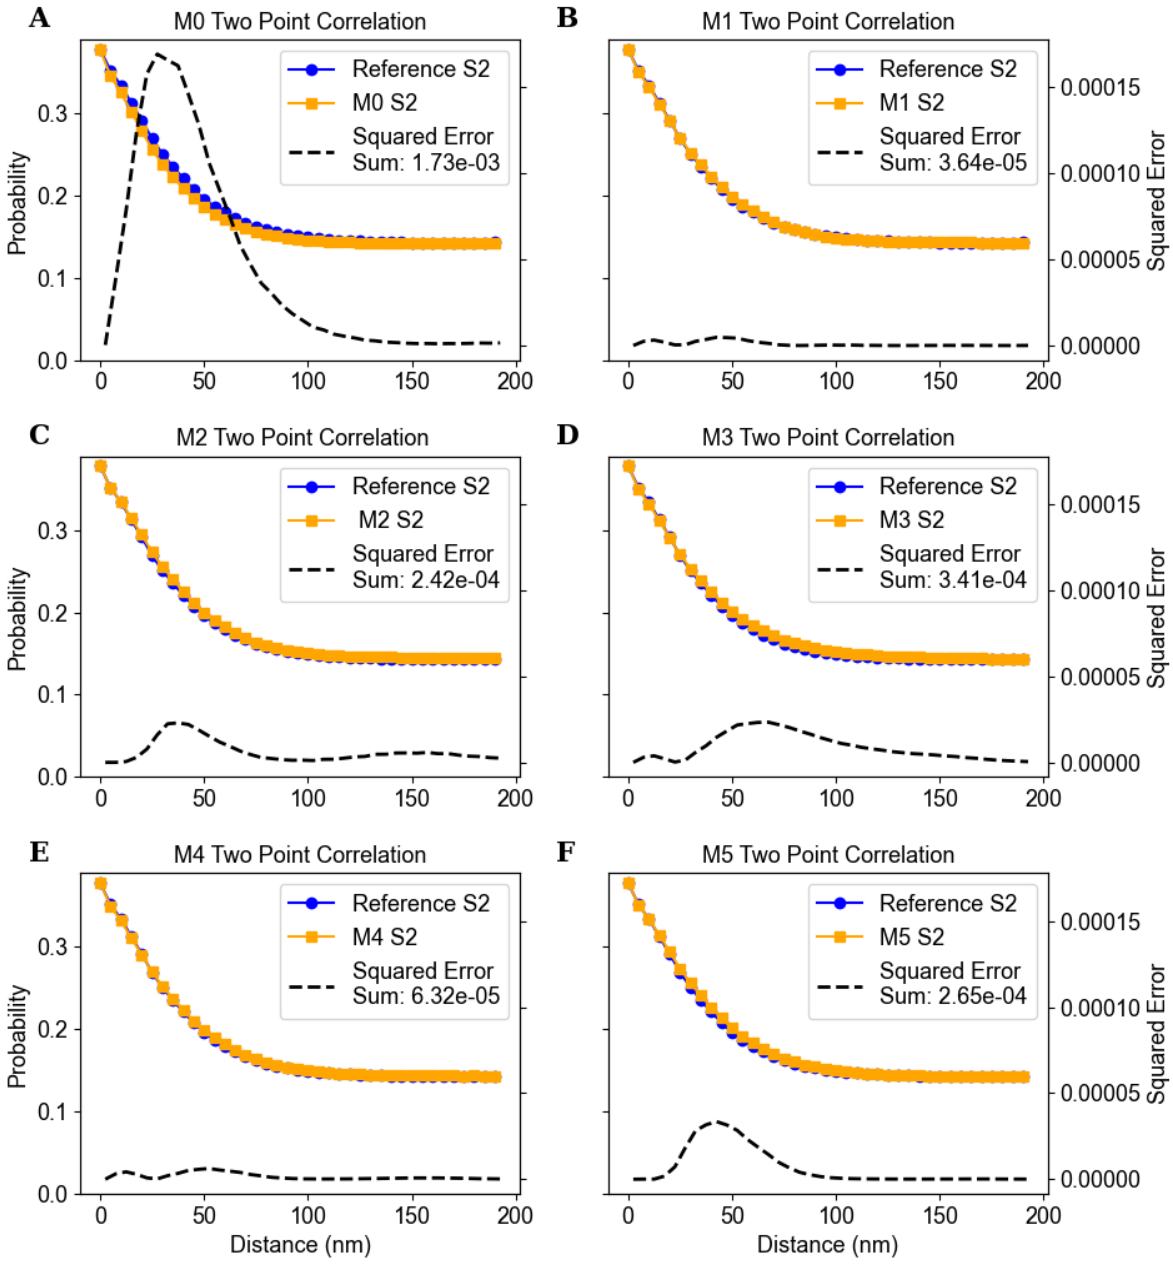


Figure 4.15: Two-point correlation for each of the reconstructions. The squared error is represented as a black dashed line, and the integral of the squared error is stated in the legend.

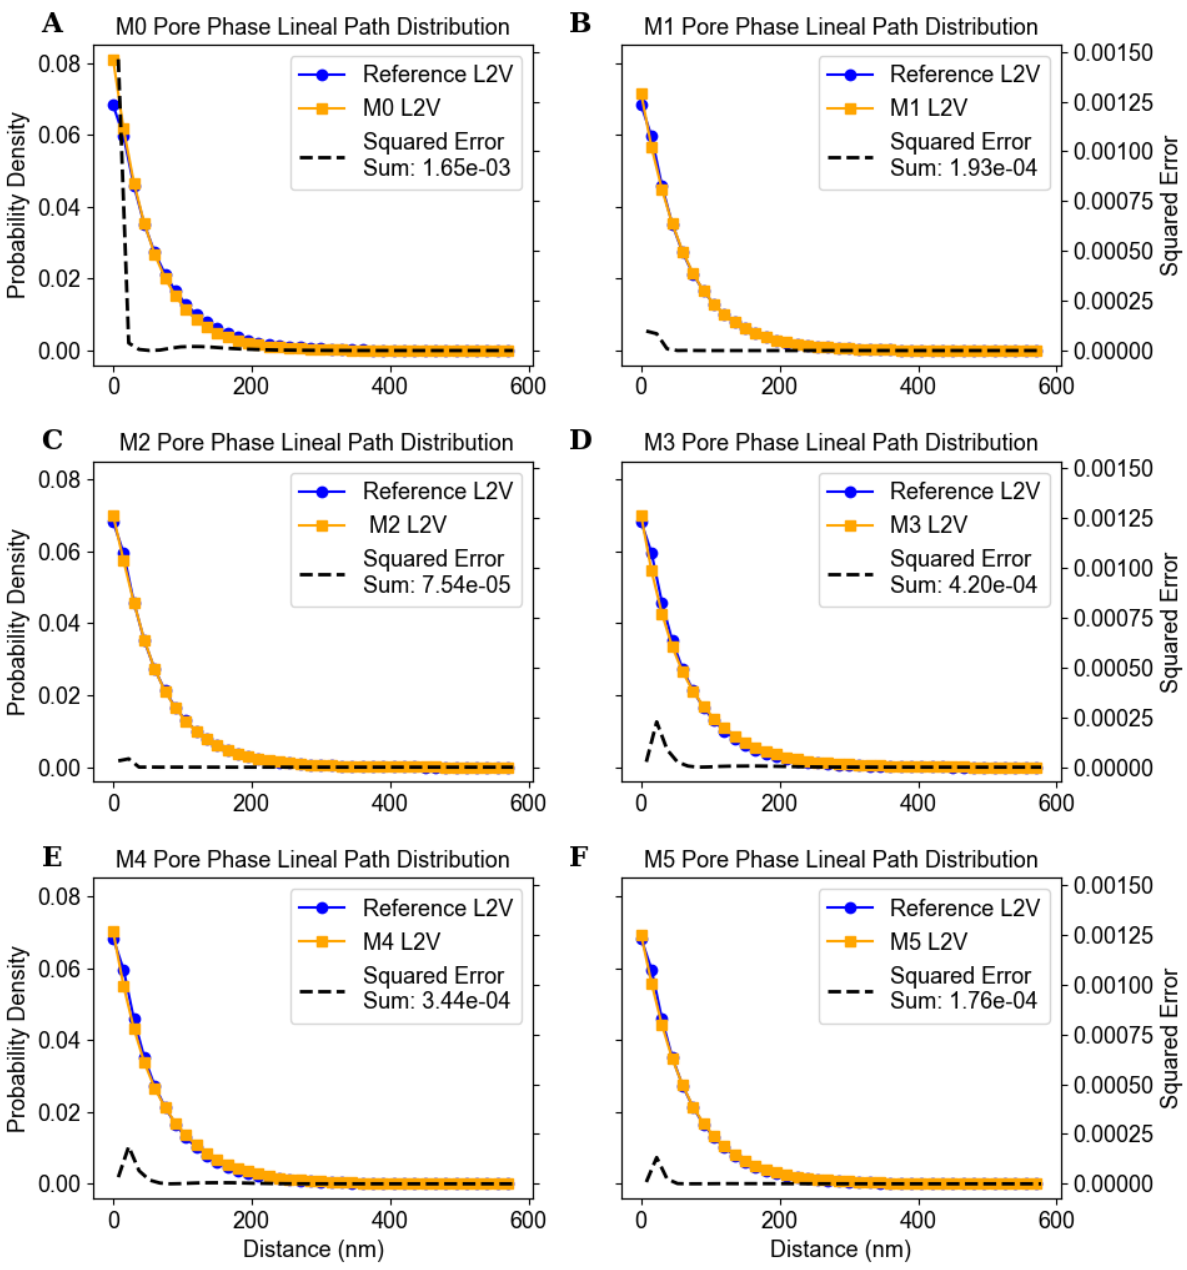


Figure 4.16: Lineal path distribution in the solid phase for the reconstructions. The squared error is represented as a black dashed line, and the integral of the squared error is stated in the legend.

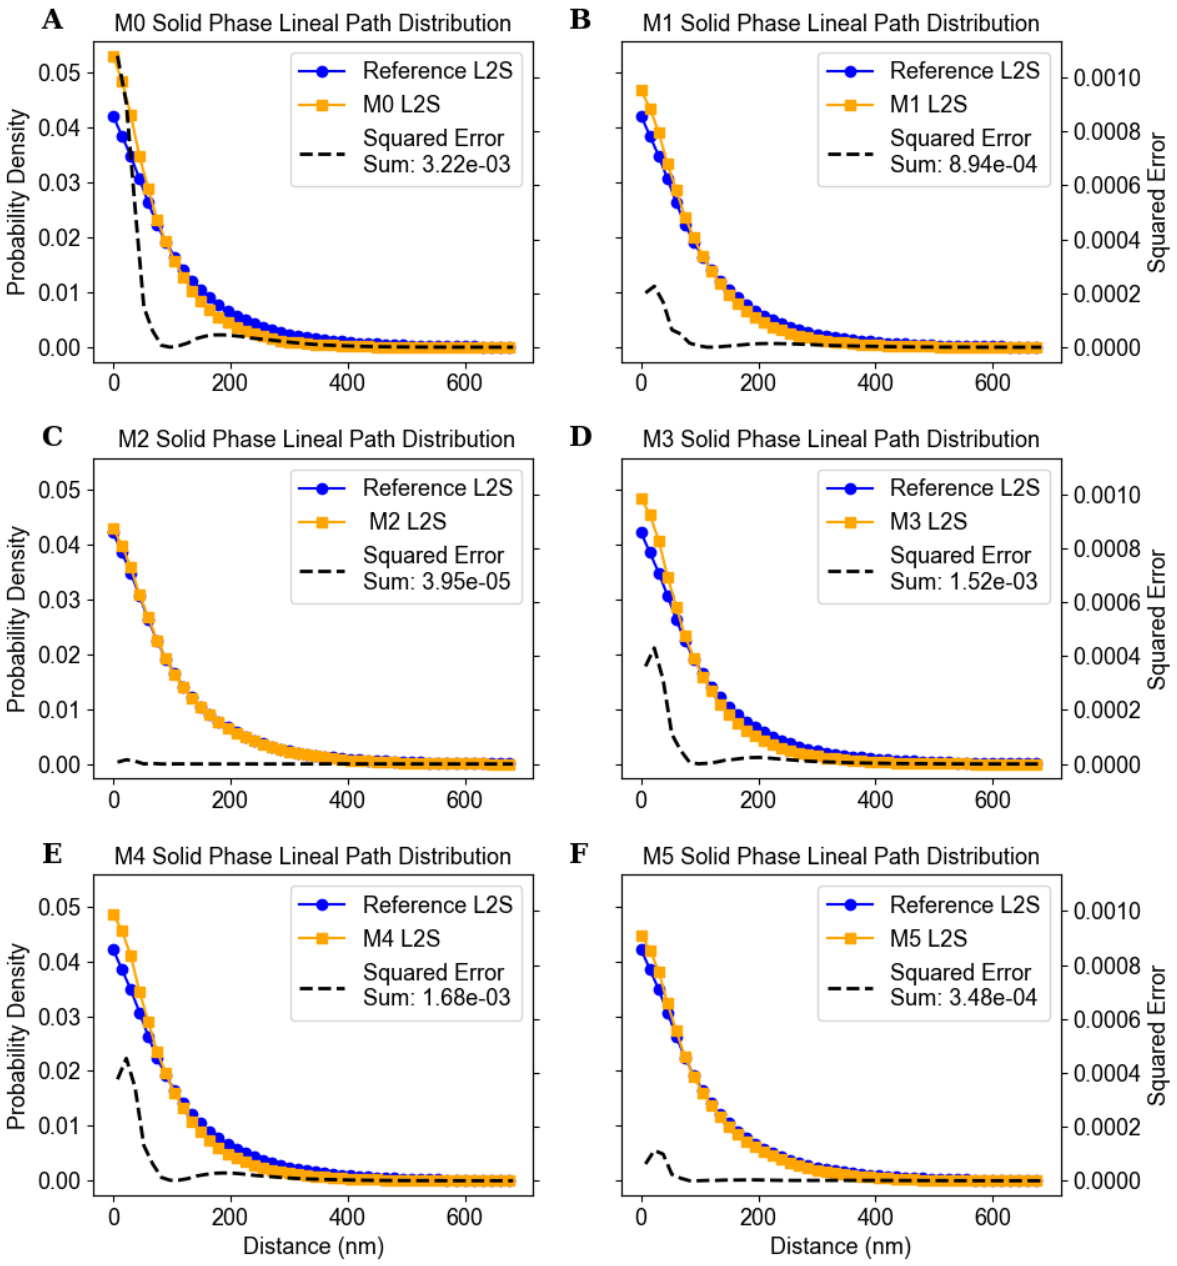


Figure 4.17: Lineal path distribution in the pore phase for the reconstructions. The squared error is represented as a black dashed line, and the integral of the squared error is stated in the legend.

Figure 4.18 shows the pore size distributions of each of the reconstructions. It is evident that optimization of the pore size distribution (M3, M4 and M5) greatly improves the accuracy over the M0-M2 reconstruction methods with an SSE reduction of up to 150-fold for M5 as compared to a 2.5 times reduction in error for M1 and barely any change when using M2. On the other hand, the error in the lineal path distributions for M3 and M4 (which optimize the pore size distribution but omit the lineal path distribution) is comparatively small.

Figure 4.19 shows the particle size distribution of each of the reconstructions. Once again, it is observed that M3, M4 and M5 greatly improve the accuracy of the particle size distribution as compared to M1 or M2. The error reduction is not as impressive as in the pore size distribution, but that is likely because the starting configuration SSE is lower for the particle size distribution. However, it should also be noted that reconstruction methods M1 and M2 actually *increase* the SSE in the particle size distribution over the starting configuration by approximately 2x and 5x, respectively.

A second observation regarding the particle size distribution is that the reconstructions all appear to have a distribution which is shifted to slightly higher radii. This can be attributed to the way the local thickness algorithm fits spheres into the domain to obtain the pore size distribution. Since the solid phase is composed of perfect spheres (albeit overlapping), the algorithm is able to perfectly fit spheres into the solid phase, which results in a maximization of the distribution. On the other hand, the reference image is not composed of overlapping spheres, and so any non-spherical elements will have their effective diameter underestimated. This could potentially be remedied by implementing the local thickness filter using equivalent diameter spheres (spheres inserted based on pore volume) instead of using inscribed spheres (spheres inserted must fit completely in the pore space, which will underestimate the volume of non-spherical shapes).

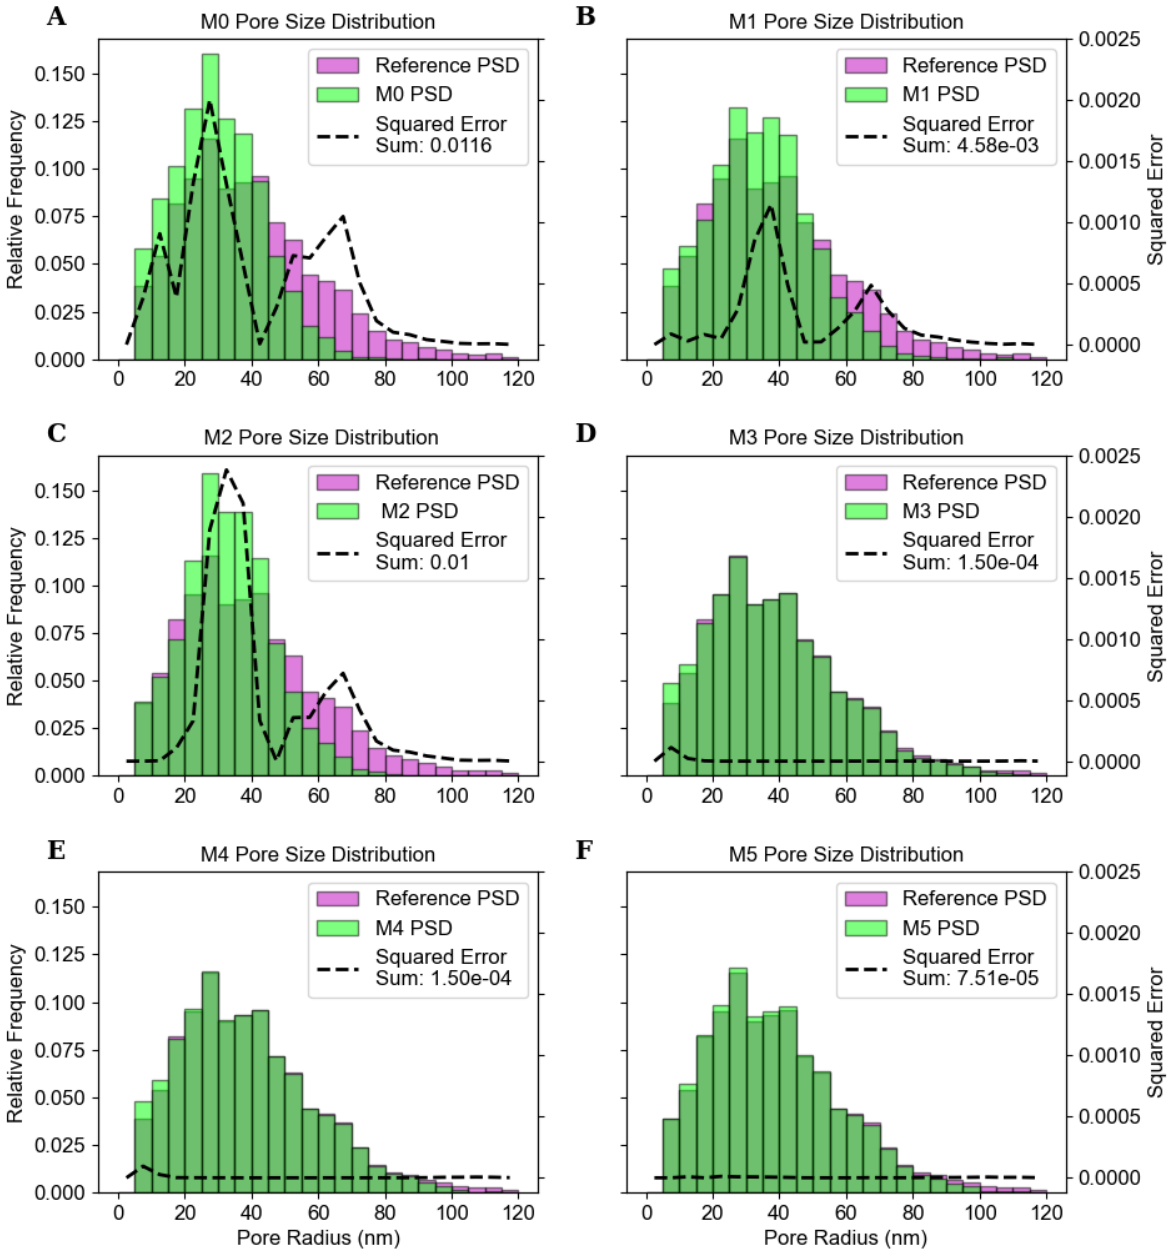


Figure 4.18: Pore size distribution of each of the reconstructions. The squared error is represented as a black dashed line, and the integral of the squared error is stated in the legend.

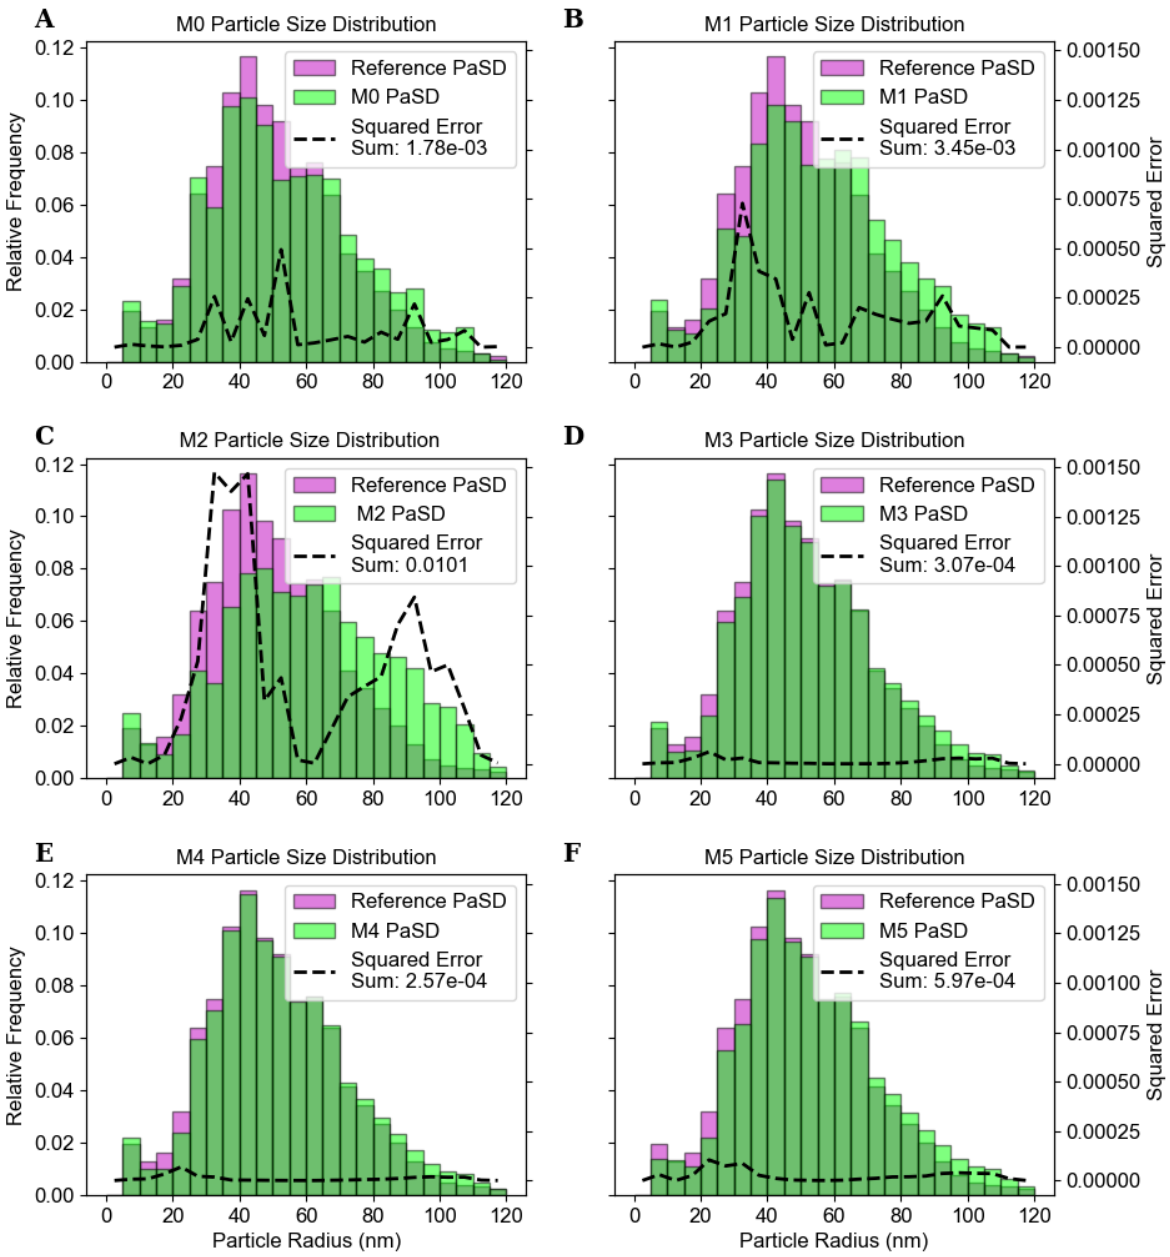


Figure 4.19: Particle size distribution for each of the reconstructions. The squared error is represented as a black dashed line, and the integral of the squared error is stated in the legend.

While this section has so far presented each of the correlations for each of the reconstructions and showed their plots, for visualization purposes it is also helpful to directly compare the sum of squared errors of each correlation and between reconstructions. Note that every correlation has been individually normalized such that the sum of its values is unity. In this manner the errors between different correlations can be compared as is shown in Figure 4.20. From this visualization it is even more evident that M1 and M2 do a comparatively poor job of optimizing the pore size and particle size distributions. In contrast, M3 and M4 do a much better job of optimizing the lineal path distributions, despite not including them in the simulating annealing energy computation. As expected, M5, which optimizes all 5 correlations, performs the best in all categories.

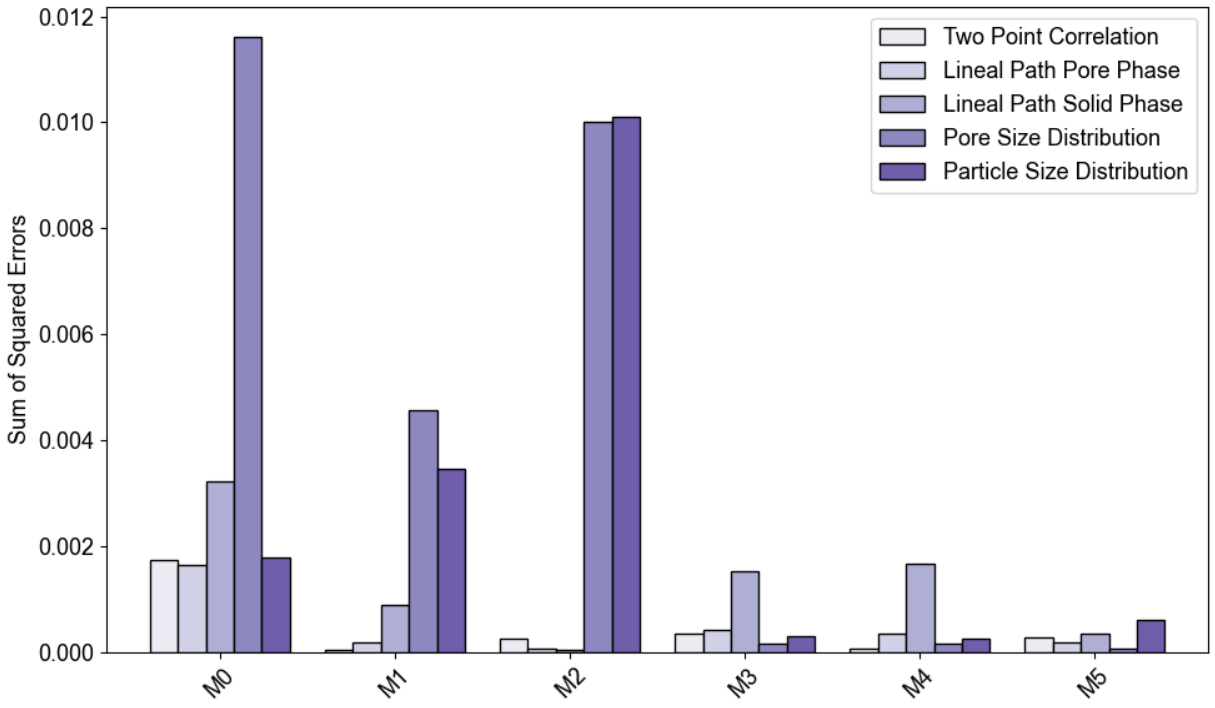


Figure 4.20: Sum of squared errors of each correlation for each reconstruction.

4.1.2.5 Porosity

Figure 4.21 shows the final computed porosities of each reconstruction. As was the case in the reconstructions from Part 1 (section 4.1.1), the porosities of all the reconstructions that use simulated annealing (M1-M5) is matched to that of the reference to the 5th decimal place with less than 0.001% error. This indicates that matching of the porosity is relatively trivial if it is included in the simulated annealing process as was done for all of the simulated annealing reconstructions.

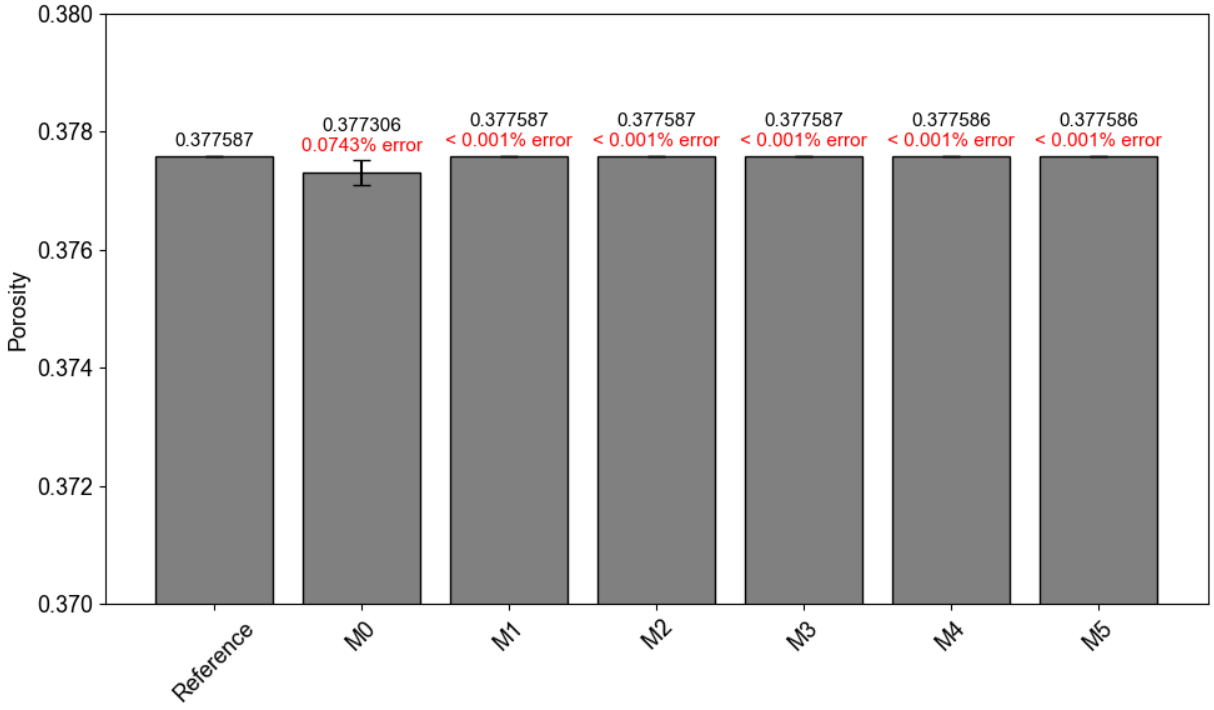


Figure 4.21: Porosity of reconstructions as compared to the reference.

4.1.2.6 Specific Surface Area

Figure 4.22 shows the specific surface area for each of the reconstruction methods. It is evident that the purely sphere-packing method (M0) greatly overestimates the specific surface area compared to the simulated annealing methods. Between the simulated annealing reconstructions, M5 performed very well, achieving a relative error of only 0.59%. The sample variability, represented by the black error bars is relatively small in all cases, indicating good reproducibility of the results.

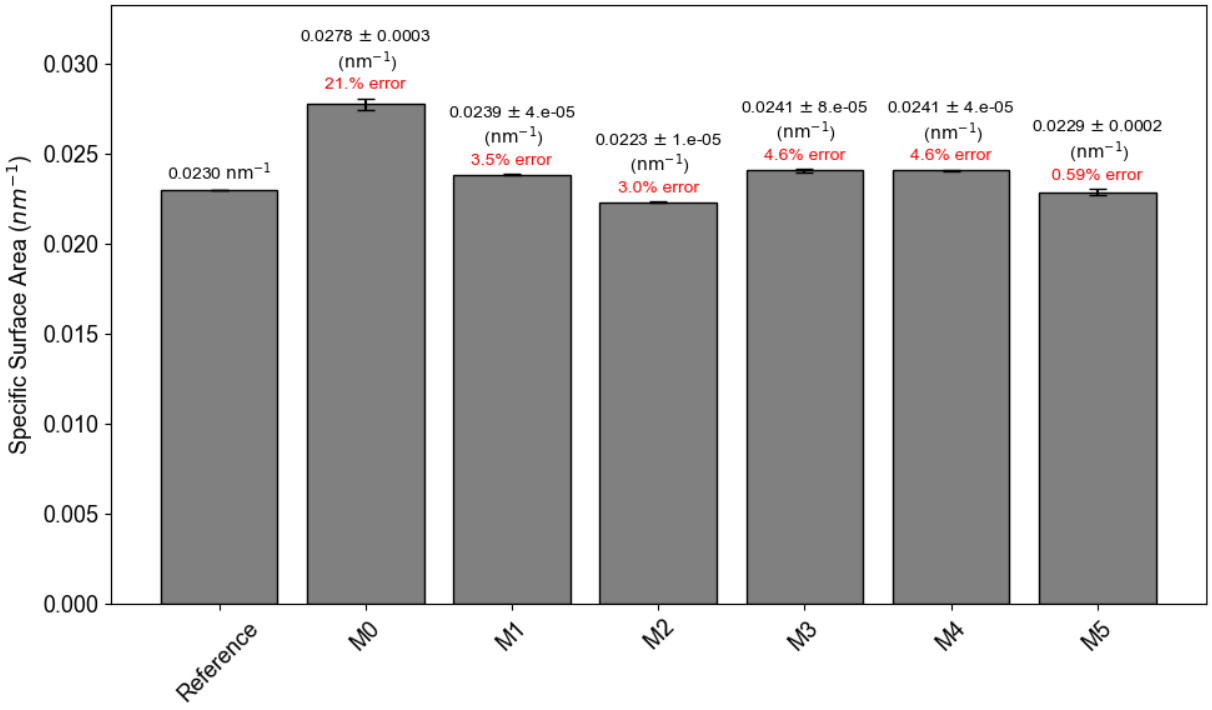


Figure 4.22: Specific Surface Area for reconstructions M0-M5. Sample error is represented by the standard deviation and shown as a black bar. Relative error compared to the reference is shown in red text.

4.1.2.7 Mean Pore and Solid Phase Diameter

Figure 4.23 and Figure 4.24 show the mean pore and particle diameters of the reference and reconstruction methods. As expected, inclusion of the pore size distribution in the simulated annealing optimization has greatly improved the accuracy of the average pore and particles sizes in the reconstructions from 14-26% relative error down to around 4% error for the mean pore size. The mean particle size is actually most optimized after just the sphere-packing stage (M0); however, it can be assumed that this is mostly a coincidence since the actual particle size distribution (Figure 4.19) has significant error. Moreover, this same reconstruction has a corresponding 26% relative error in the mean pore size. Other than this anomaly, the mean particle size is most accurate with M3 and M4. M2 and M5, which both include the lineal path distribution, show a significant increase in solid phase clumping. The inclusion of the pore and particle size distributions in M5 means that this increase is somewhat dampened though.

Based on these results, it appears that omission of the lineal path distribution and inclusion of the pore and particle size distributions in simulated annealing gives the most accurate mean pore and particle diameters.

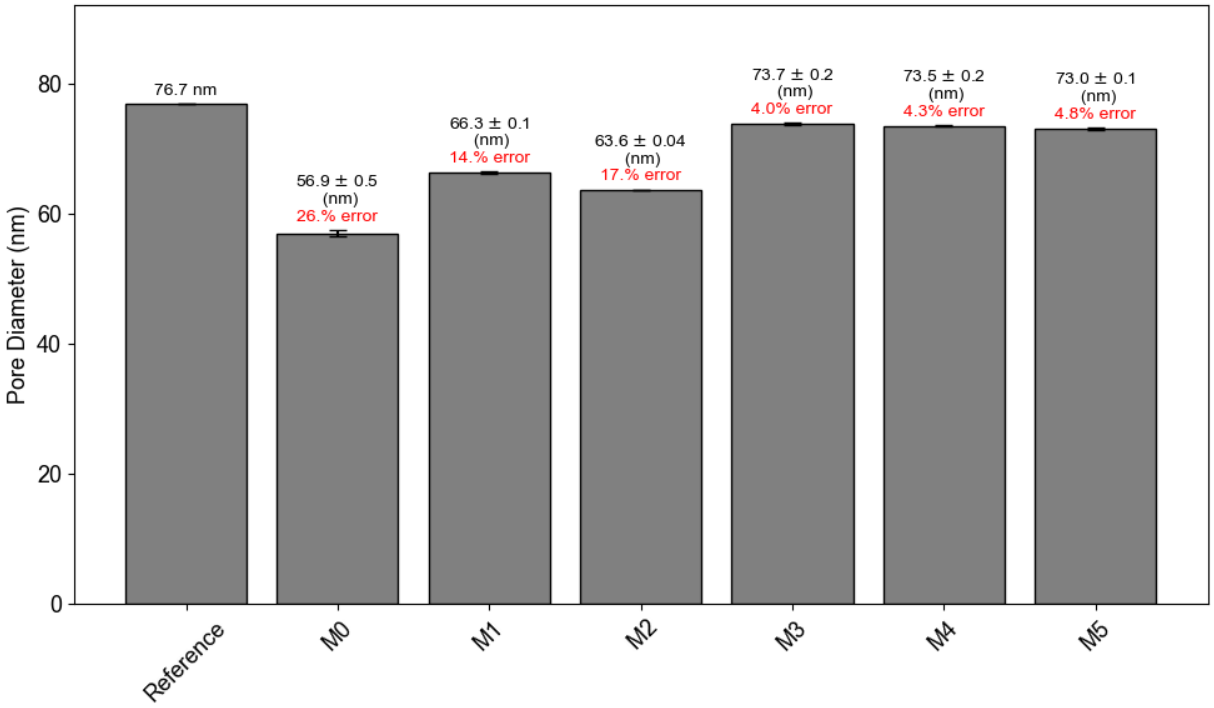


Figure 4.23: Mean pore diameter for the reference and reconstructions.

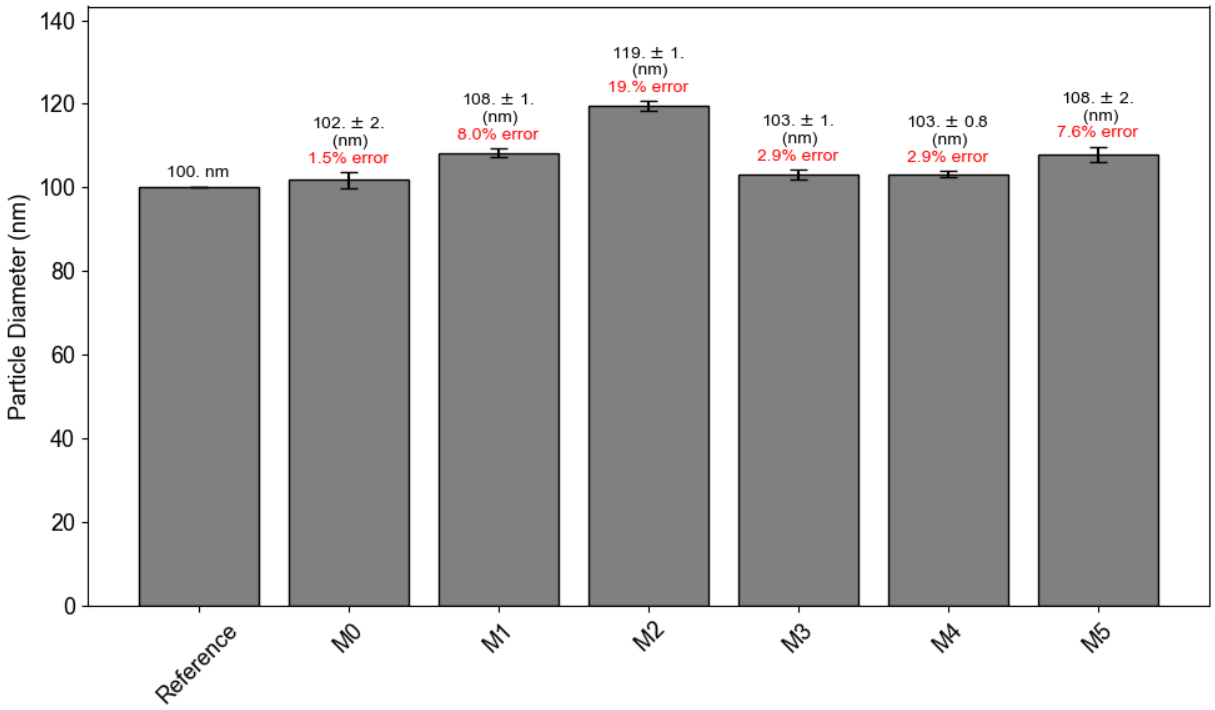


Figure 4.24: Mean Particle Diameter for the reference and reconstructions.

4.1.2.8 Mean Chord Lengths

Figure 4.25 and Figure 4.26 show the pore and particle phase mean chord lengths for the reference and each of the reconstructions. In contrast to the results and analysis of the mean pore and particle diameters, for the mean chord lengths, M2 performs most accurately, followed by M5. This means that in order to obtain an accurate linear characteristic length, the lineal path distribution *should* be included. It therefore does not appear possible to use the sphere-based simulated annealing method to reconstruct a catalyst layer that has both an accurate pore/particle diameter and an accurate pore/particle chord length. This leads to the question of which of these metrics is more important in determining the accuracy of the transport properties. This question is considered in the second part of this thesis.

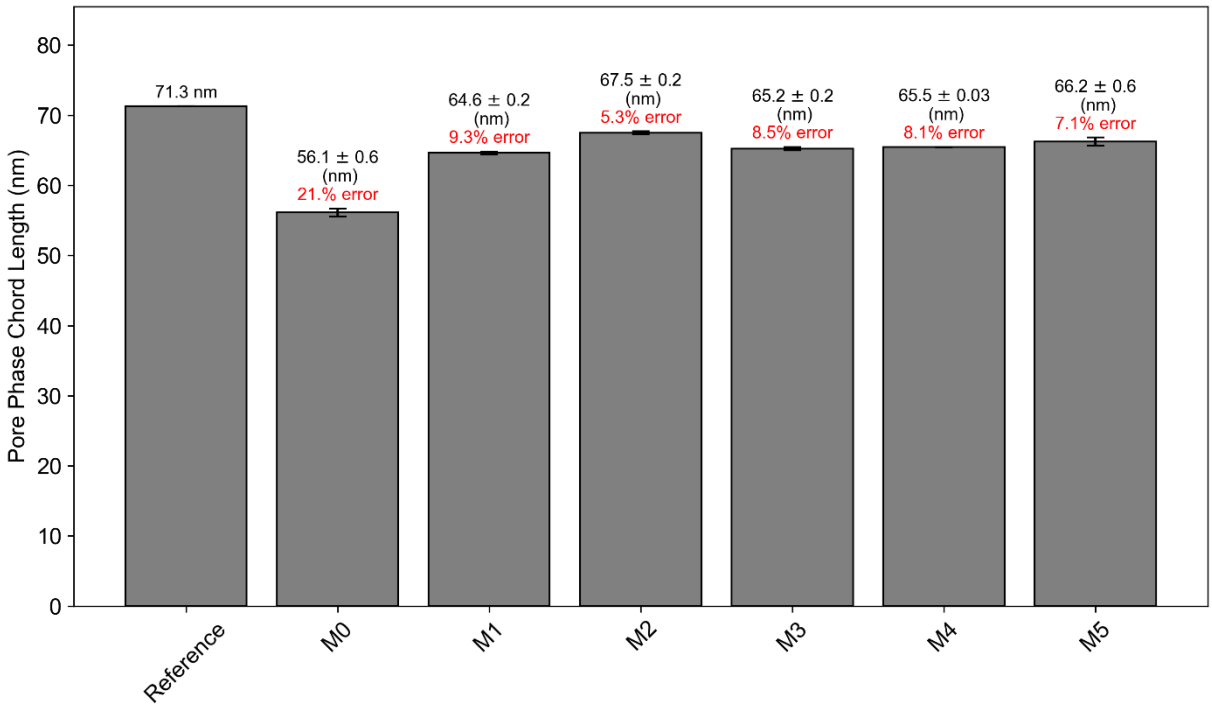


Figure 4.25: Mean pore phase chord lengths for the reference and reconstructions.

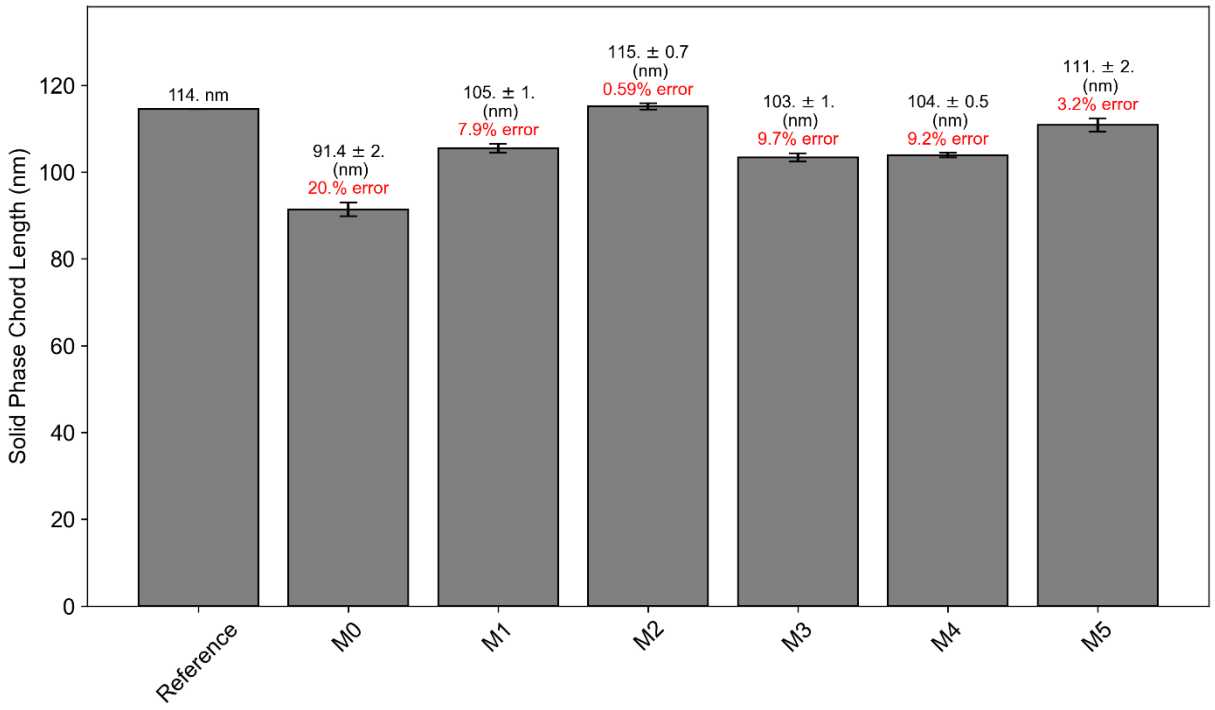


Figure 4.26: Mean particle phase chord length for the reference and reconstruction

4.2 Transport Simulation of Reconstructions

For this part of the thesis, simulations were run on each of the reconstructed catalyst layers using methods M0-M5. In the same fashion as the metrics presented in section **4.1.2**, each reconstruction method has 3 associated reconstructions, so that a measure of variability can be determined for each of the effective transport properties. The transport simulations were repeated on each of the reconstructions to obtain a sample standard deviation, which is represented on the bar charts as black error bars. The relative error is a measure of how accurate the reconstruction effective transport property was to the reference.

In order to conduct the simulations on the reference image, the reference FIB-SEM stack needs to have the same resolution in all three directions. This was accomplished by a zeroth order spline interpolation in the z-

direction. This essentially just meant repeating each of the slices 6 times to obtain a nominal 5nm resolution in every direction, which results in a blockier structure as seen in Figure 4.27.

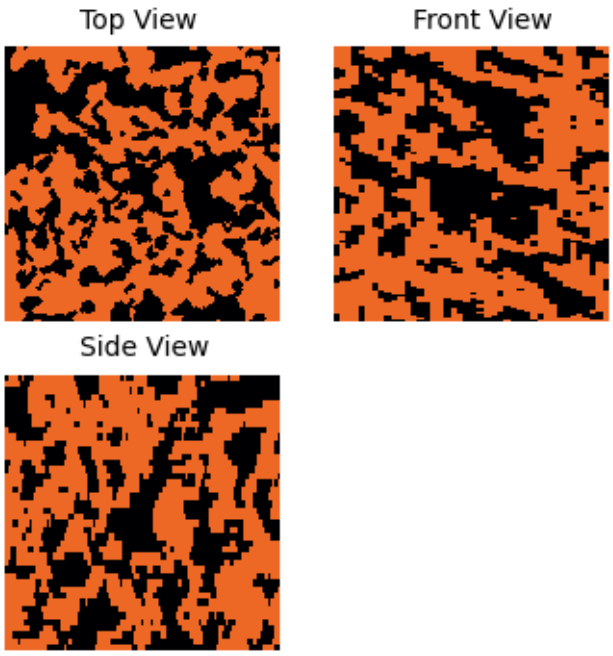


Figure 4.27: Views of the reference image after zeroth order interpolation so that the resolution is 5 nm in all three cartesian directions.

4.2.1 Effective Diffusion and Tortuosity

The effective diffusivity of the reconstructions was computed using the modified random walk method as described in section 3.2.2. This method accounts for the Knudsen effects by allowing random walker molecules to move in a straight line for the length of their mean free path. Simulations were conducted using a temperature of 298K and pressure of 1 atm. 50000 walkers were run in parallel for 200,000 steps each.

Figure 4.28 shows the normalized effective diffusivity for each reconstruction method and the reference image. Figure 4.29 shows the equivalent pore tortuosity for each reconstruction. It is evident that there is a discrepancy between the diffusivity of the reconstructions and that of the reference image. Initially, this was thought to be due to the poor resolution in the z-direction of the reference image as seen in Figure 4.27.

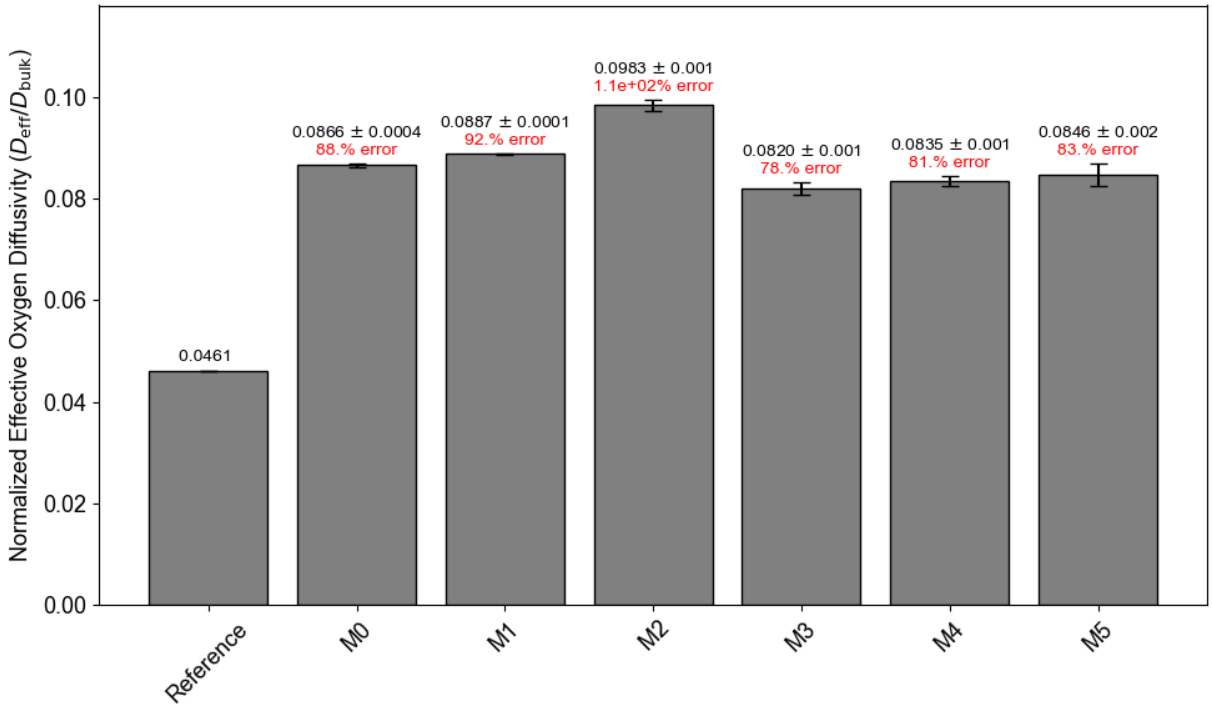


Figure 4.28: Effective Oxygen Diffusivity for each of the reconstruction methods. Simulations were conducted with a temperature of 298K and pressure of 1 atm.

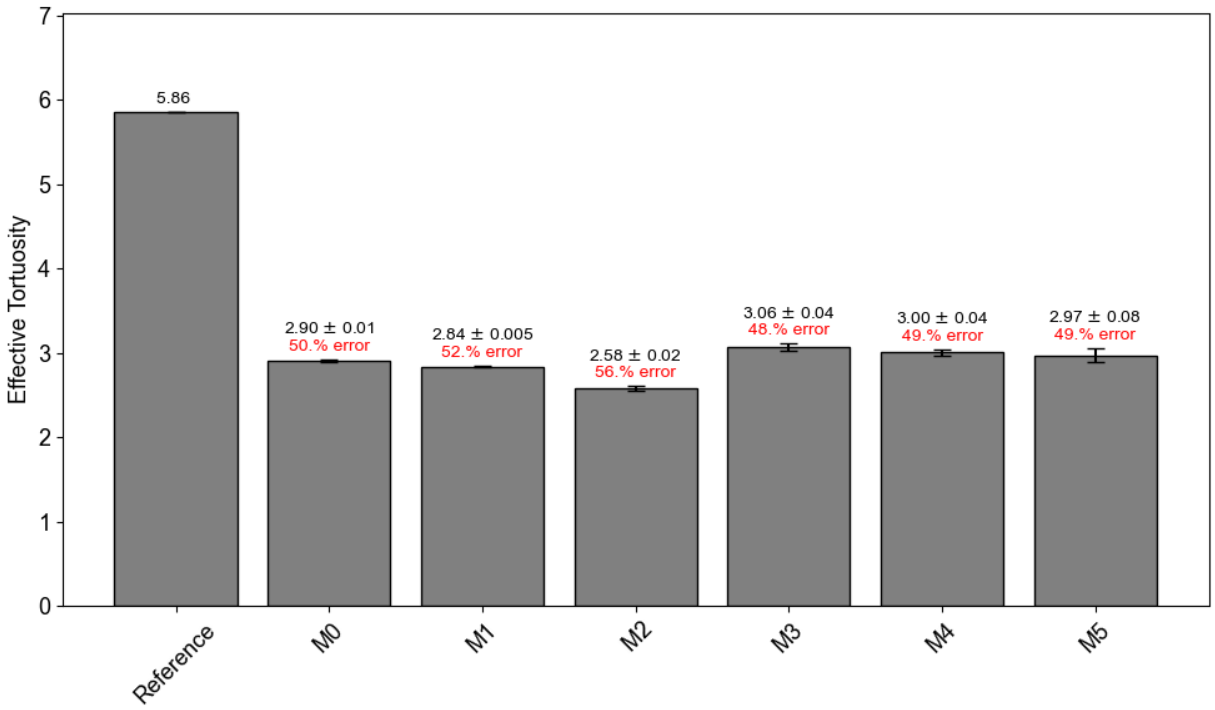


Figure 4.29: Pore tortuosity of reconstruction methods.

In order to test whether the resolution was the issue, a single reconstruction was performed on an artificially generated 3D structure similar to that of the reference image. This reference was created using *porespy*'s [55] "blobs" generator function which works by applying a gaussian blur to randomly generated noise. Diffusion calculated in this reconstruction showed similar results as those presented in

Figure 4.28: namely the diffusion in the sphere-based reconstructions was roughly twice that of the artificially generated reference image. Since this test eliminates the resolution and FIB-SEM cutting process as a possible cause for this error, an alternative explanation is that the sphere-based nature of the reconstructions results in less tortuous pathways in the pore phase. This difference is depicted in Figure 4.30 in 2D images generated artificially for demonstration purposes only. It is evident that the sphere-based image

produces a structure with much less tortuous pathways than the “blobs” image which more closely represents what is seen in FIB-SEM images. Additionally, in the “blobs” image, more dead-end pores appear which will also increase the overall tortuosity of the structure. Dead-end pores are less common in a sphere-based reconstruction due to how spheres overlap with each other. They leave ‘v’-shaped spaces which will mean there it becomes more difficult to create long, narrow dead-end pore tunnels in a sphere-packing.

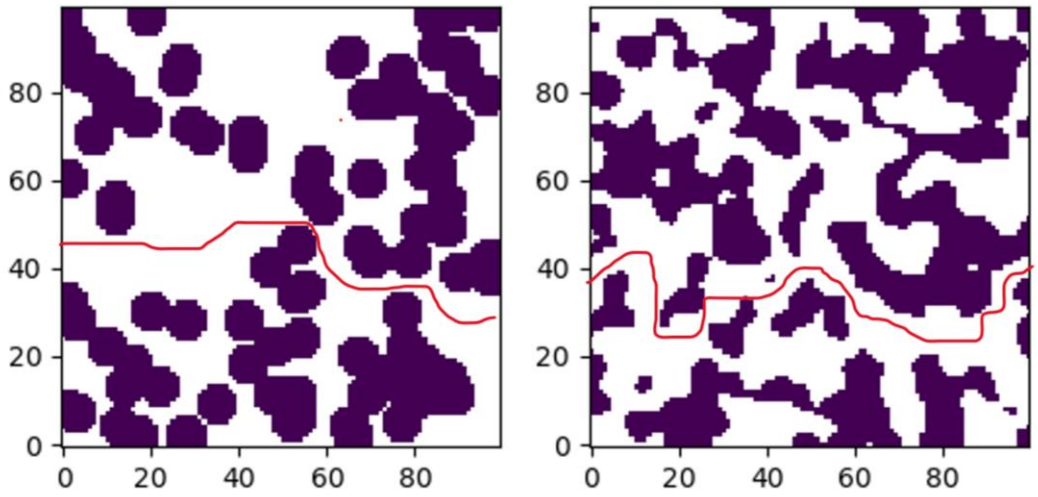


Figure 4.30: Comparison of pathways for artificially generated sphere-based image (left) and a structure similar to the reference FIB-SEM image generated by applying a Gaussian blur to random noise (right). It is evident that the sphere-based image produces a structure with less tortuous pathways.

The above analysis then might call into question the use of sphere-based simulated annealing and sphere-packing methods in reconstructions of the PEMFC catalyst layer. However, it is also important to note that the catalyst layer is composed of carbon black particles and so it might be that the sphere-based geometry is still

correct, and that the actual geometry observed in the SEM images is inaccurate which could occur from the process of binarizing the images by thresholding the greyscale values in the SEM image.

In order to gain more insight into whether the reference image is more correct or if the sphere-based reconstructions are more accurate, it is useful to compare the results from these simulations to those from the literature (both experimental and numerical studies).

The Bruggeman approximation, given a porosity of 0.378 and an exponent of 1.5 predicts a normalized effective diffusivity 0.23, which is approximately 2-3 times larger than the diffusivity predicted by the reconstructions, and 5 times larger than that of the reference image.

Das and Li [33] developed an empirical correlation for diffusion in catalyst layers, which would predict a normalized effective diffusivity of 0.28 with the same parameters, which is even higher than even the Bruggeman approximation.

Meanwhile, Zhao et al [49] performed experimental measurements of the catalyst layer effective diffusivity using a modified Loschmidt cell and found an effective diffusivity of 4.19e-7 m^2/s. If the bulk oxygen diffusion at ambient temperature and pressure is taken as 2.19e-5 m^2/s, this would be equivalent to a normalized effective diffusivity of 0.019, an entire order of magnitude smaller than what was predicted by either of the empirical correlations.

Given the large range of reported effective diffusivities it is difficult to ascertain whether or not the reference image or the sphere-based reconstructions are more realistic, and more research may be necessary.

4.2.2 Effective Thermal Conductivity

The effective thermal conductivity of each of the reconstructions was calculated using the PuMA solver in Python, detailed in section 3.2.1. The finite volume method is employed, and periodic boundary conditions are

implemented on the side boundaries, while Dirichlet boundary conditions are used for the end in the direction of the applied gradient. Symmetrical boundary conditions are used for the reference image, as these images are not periodic. The solution is obtained using a biconjugate gradient stabilized solver. Bulk phase conductivities are listed in Table 3.2.

The normalized effective thermal conductivities of each reconstruction and the reference are shown in Figure 4.31, in which sample standard deviations are used as a measure of sample variability, shown as the absolute error and the black error bars. The relative error as compared to the reference is annotated in red text as a percentage. The thermal conductivities are normalized using a $k_{\mathrm{bulk}} = 0.3 \frac{\mathrm{W}}{\mathrm{cm \cdot K}}$ (the maximum thermal conductivity assuming the entire structure is solid carbon). It is apparent that the thermal conductivity is not very sensitive to the structure of the catalyst layer as there are only small deviations dependent on which reconstruction method is employed. This makes sense since the initial sphere pack has a more randomly distributed packing of spheres. Since the solid phase materials (carbon and ionomer) have bulk thermal conductivities an order of magnitude greater than that of air, a more tortuous solid pathway will reduce the overall effective thermal conductivity. However, the fact that heat is still transferred through the air phase helps to dampen the differences in structure which is why the thermal conductivities for each of the reconstruction methods are relatively similar.

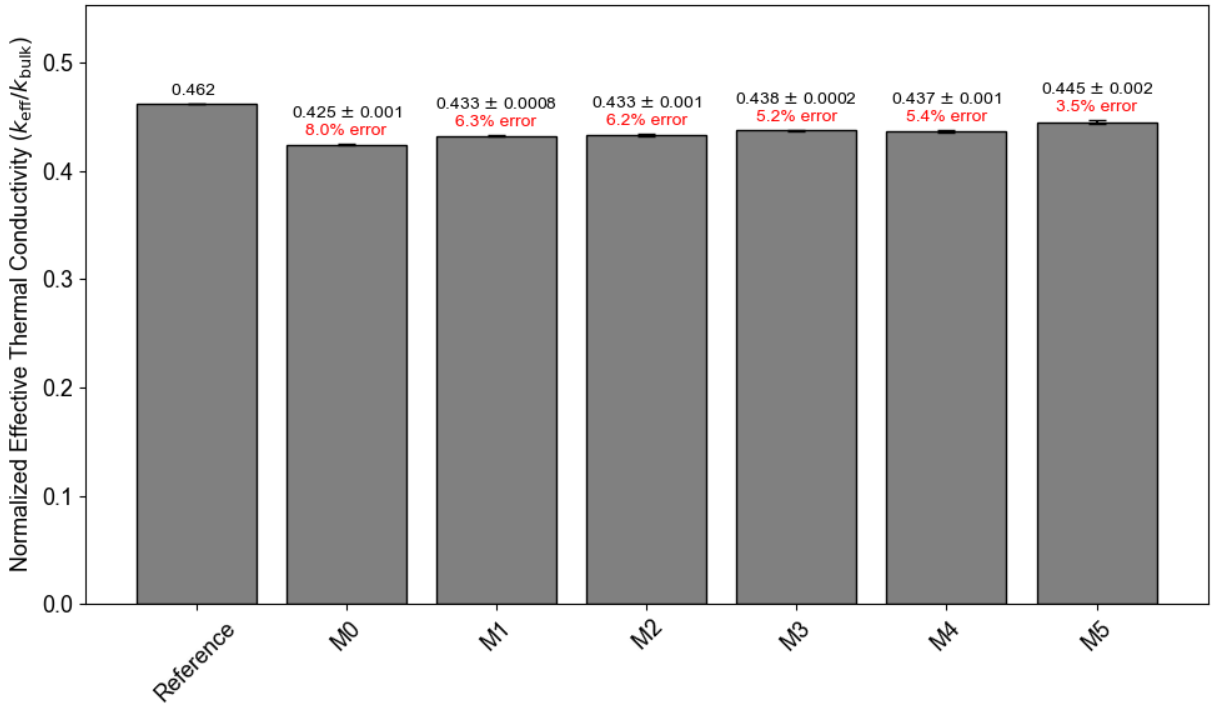


Figure 4.31: Thermal conductivity of reference and reconstructions. Relative error compared to the reference is indicated in red text. Sample variability is represented as error bars and is the standard deviation.

While each of the reconstructions exhibit similar thermal conductivities, the variability between reconstructions of the same method, calculated as the standard deviation and depicted with the black error bars, is very small, so it is reasonable to conclude that the reconstruction method does have a small, but finite effect on the effective thermal conductivity. Based on the relative error (shown in red text on the figure) the reconstruction using Method 5 had the most accurate result with a 3.5% relative error to the reference. This is compared to the purely sphere-packing method which had an 8% relative error, so it can be concluded that some form of simulated annealing is preferred, and Method 5 gives the best results.

Lange et al [19] conducted a study of the effective transport properties in a catalyst layer reconstructed using a sphere-packing approach most similar to the Method 0 presented here. Their simulations resulted in normalized effective thermal conductivities ranging from 0.33 to 0.46, which is in good agreement with the results presented here.

For a qualitative experimental verification of the effective thermal conductivity, the results presented here can be compared against the experimental measurements conducted by Burheim et al [94] where they measured the thermal conductivity of the catalyst layer under varying degrees of compaction pressure in both dry and wet conditions. They found that under dry conditions, the thermal conductivity was approximately 0.07 W/cm-K, but it increased to 0.14 W/cm-K when humidified. The wet thermal conductivity is close to the results presented here, but the simulations conducted did not account for water, so it is possible that the thermal conductivity of both the reconstructed catalyst layers and the reference is being overestimated.

4.2.3 Effective Ionic Conductivity

The effective ionic conductivity of the reconstructed catalyst layers was computed in a similar manner to calculation of the thermal conductivity. The bicgstab solver was again employed to conduct the finite volume simulation. The bulk ionic conductivity of the ionomer is assumed to be 0.03 S/cm.

The normalized effective ionic conductivities of the reconstructions are presented in Figure 4.32. More variability is observed between the effective ionic conductivity of each of the reconstruction methods than seen for the effective thermal conductivities. The initial sphere pack (M0) has a high relative error of 17%, while the best results are for Method 2 at 3.2% relative error and Method 5 at 2.8% relative error. Both Method 2 and Method 5 incorporate the lineal path distribution in the simulated annealing process, which appears to play a dominant role in generating a structure that accurately represents the ionomer phase of the catalyst layer.

This is likely because inclusion of the lineal path distribution tends to cause clumping of the solid phase into larger agglomerates. In turn, this means there is less surface area available for the ionomer to attach to during the segmentation process which means that the ionomer film thickness will be bigger, thus reducing the tortuosity of the ionomer phase and increasing its conductivity. This observation can be confirmed by returning to the specific surface area (Figure 4.22), which clearly indicates that M5 and M2 have the most accurate specific surface area (least relative error compared to the reference).

For a numerical comparison, the results presented here can be compared to those of Lange et al [20], where they computed the effective ionic conductivity on their reconstructed catalyst layers. Their reconstructions were generated using a sphere-packing approach most similar to Method 0 used in this thesis. For an ionomer volume fraction of 25% (the same as in the reconstructions presented here) they found a normalized effective ionic conductivity of 0.075. This is in close agreement with the effective ionic conductivity of the reconstruction using Method 0. However, if the reference ionic conductivity is taken to be the true value, it is evident that both Method 0 and the sphere-packing approach used by Lange et al results in a structure which underestimates the true ionic conductivity. The most accurate effective conductivity is obtained using reconstruction method 5.

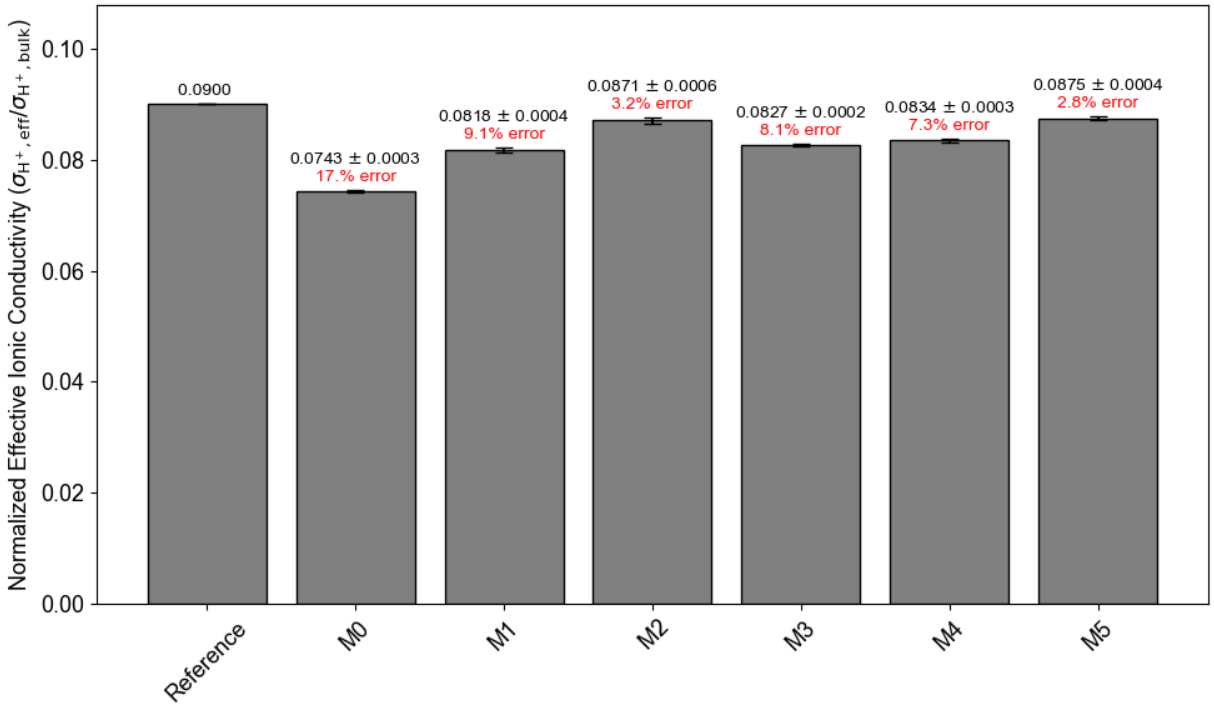


Figure 4.32: Ionic Conductivity for the reference and reconstruction methods.

4.2.4 Effective Electrical Conductivity

Figure 4.33 shows the normalized effective electrical conductivities of the reference image and reconstructed catalyst layers. In contrast to the results of the effective gas diffusivity, in this case the reconstructions are all *under*estimating the effective electrical conductivity of the reference. This indicates that the solid phase in the reconstructions is not as well connected as in the reference, whereas the pore phase is too well connected, causing the reconstruction to severely overestimate the effective oxygen diffusivity and severely underestimate the effective electrical conductivity. It may be possible to improve the accuracy of the electrical conductivity by allowing the spheres to clump even more, forming larger agglomerates but this would also then increase the average pore sizes which will further increase the predicted effective diffusivity away from that of the

reference. Indeed, similarly to the results of the ionic conductivity simulations, it appears that Method 2 and Method 5 are able to achieve values for the effective electrical conductivity closest to those of the reference, but the relative errors are still large at 26% and 23%, respectively.

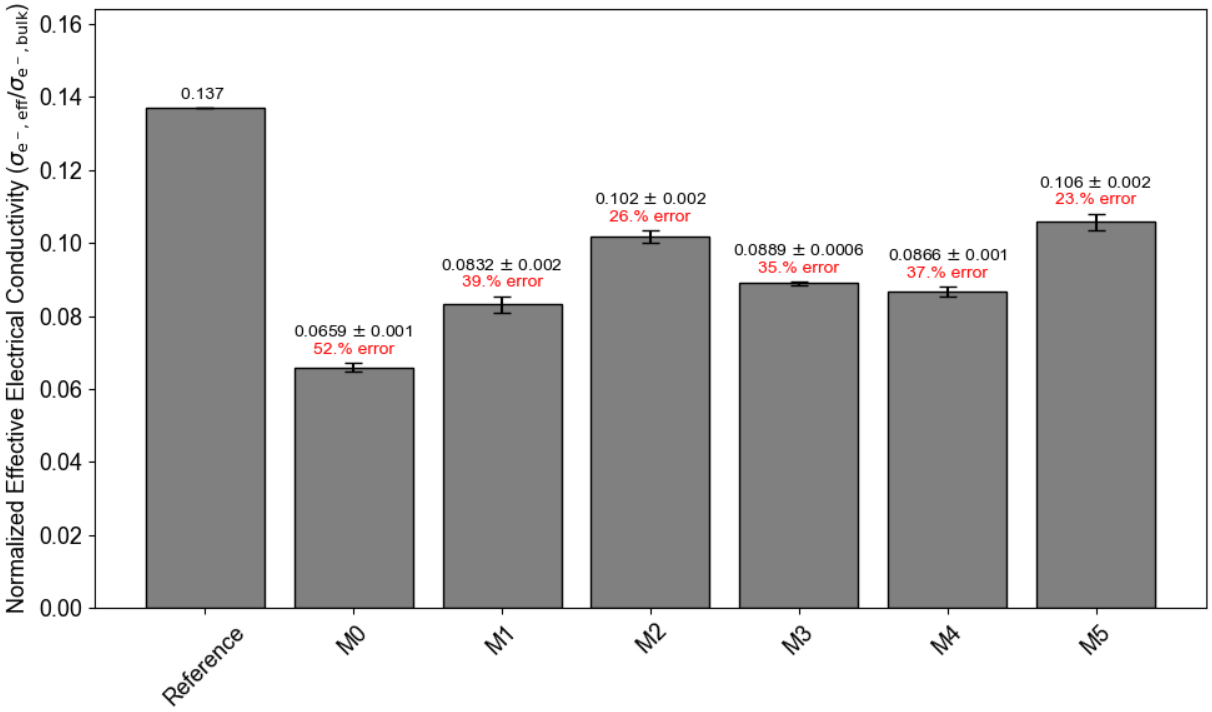


Figure 4.33: Electrical Conductivity of reference and reconstructions

The current results can be compared against those of the sphere-packing reconstruction of Lange et al [20]. They found that the effective electrical conductivity depended heavily on the amount of overlap that was allowed between spheres. This makes sense, as increased overlap means more pathways will be available for the electrons to travel. They reported normalized effective electrical conductivities that ranged from 0.015 up

to 0.060 for a carbon volume fraction of 0.38. The upper end of their range correlates well with the M0 reconstruction method presented in this thesis, which makes sense since both methods are similar. However, neither is very close to the effective electrical conductivity of the reference image, indicating that the carbon phase in the reference image is significantly better connected than in the sphere-based reconstructions. Additionally, these results are similar to those found in the study by Wu et al [23] which also used a sphere-based simulated annealing method similar to M1 in this thesis, where they report a very high carbon phase tortuosity of 7.83 with a carbon volume fraction of 0.3, indicating a normalized effective electrical conductivity of 0.038. Their reconstructions are also sphere-based which is further evidence that the discrepancy between the reference image and the reconstruction methods presented in this thesis is due to the spherical nature of the reconstructions.

4.2.5 Summary of Conductivity Data

Table 4.5 shows all of the normalized conductivities and diffusivities of each of the reconstruction methods side by side, along with their calculated standard deviation. Table 4.6 shows the relative error of each conductivity and diffusivity, with respect to the value calculated on the reference image. Once again, it is evident that the sphere-based simulated annealing reconstruction as implemented in this thesis provides an accurate geometry for the purposes of determining the ionic conductivity and the thermal conductivity, but it falls short in measuring the gas phase diffusivity and the electrical conductivity. In these cases, the gas diffusivity is significantly overestimated, and the electrical conductivity is significantly underestimated which is due to the high connectivity of the void phase and the narrow constrictions in the connections between spheres in the solid phase.

The average standard deviation for every conductivity measurement is also worth noting, and ranges from 0.2% for the thermal conductivity up to 1.7% for the electrical conductivity. This seems to indicate that both

the implemented random walk method and the finite volume methods for calculating the effective conductivities are comparably reliable.

Table 4.5: Summary of normalized conductivities measured for each reconstruction method, reported with calculated sample standard deviation from three samples of each reconstruction method. Each conductivity is normalized by its bulk value as given in Table 3.2.

	Electrical Conductivity	Ionic Conductivity	Thermal Conductivity	Oxygen Diffusivity
REF	0.137	0.0900	0.462	0.0461
M0	0.066 ± 0.001	0.0743 ± 0.0003	0.425 ± 0.001	0.0866 ± 0.0004
M1	0.083 ± 0.002	0.0818 ± 0.0004	0.433 ± 0.001	0.0887 ± 0.0001
M2	0.102 ± 0.002	0.0871 ± 0.0006	0.433 ± 0.001	0.0983 ± 0.0010
M3	0.089 ± 0.001	0.0827 ± 0.0002	0.438 ± 0.000	0.0820 ± 0.0012
M4	0.087 ± 0.001	0.0834 ± 0.0003	0.437 ± 0.001	0.0835 ± 0.0010
M5	0.106 ± 0.002	0.0875 ± 0.0004	0.445 ± 0.002	0.0846 ± 0.0022

Table 4.6: Error of each conductivity measurement relative to the reference value (calculated RE = $\frac{\mathbf{k}-\mathbf{k_{ref}}}{\mathbf{k_{ref}}}$) where RE is the relative error and k is the generalized conductivity measurement.

	Electrical Conductivity Relative error	Ionic Conductivity Relative Error	Thermal Conductivity Relative Error	Oxygen Diffusivity Relative Error
REF	-	-	-	-
M0	52.%	17	8.0%	88.%
M1	39.%	9.1	6.3%	92.%
M2	26.%	3.2%	6.2%	1.1E2%
M3	35.%	8.1%	5.2%	78.%
M4	37.%	7.3%	5.4%	81.%
M5	23.%	2.8%	3.5%	83.%

4.2.6 Chapter Summary

This chapter presented the results of the two novel reconstruction methods: In part 1, the focus was on using the lineal path distribution to reconstruct the catalyst layer structures. While this method showed improvement over the conventional method, it was evident that the pore size and particle size distributions still did not match very well.

Therefore, in Part 2, the novel computation of the pore size distribution using the medial axis method was used to generate several more reconstructions, showing significant improvements in the pore/particle size distributions, in addition to the specific surface area.

Further analysis of the reconstructions was then conducted using transport simulations to obtain the effective transport properties of each reconstruction.

Chapter 5

Conclusions and Recommendations for Future Work

5.1 Conclusions

This thesis work has put forward two new methods for reconstructing the proton exchange membrane fuel cell catalyst layer using a sphere-based simulated annealing technique and statistical correlations extracted from FIB-SEM images, in which a sphere-packing method is used to initialize the structures (M0). The conventional simulated annealing technique (M1) optimizes the two-point correlation only. While this is the simplest, the structural degeneracy inherent in the two-point correlation means that simulated annealing optimization of the two-point correlation is insufficient to accurately reconstruct the catalyst layer microstructure. Therefore, this thesis first introduces a second reconstruction technique (M2) that optimizes both the two-point correlation *and* the lineal path distributions (for the void and solid phases), which characterize the linear connectedness of the structure. M2 shows improvements over the conventional M1 technique in the accuracy of the specific surface area, and mean chord lengths. However, inclusion of the lineal path distributions causes excessive clumping of each phase which means that the radial characteristic size is inaccurate.

In order to improve the accuracy of the pore size distribution in the reconstruction, a fast implementation of the pore size distribution computation was developed and implemented in Python so that the pore size distribution can be directly optimized in the simulated annealing process. Three additional reconstructions were generated using a) the pore size and particle size distributions (M3); b) combining the two-point correlation with the pore/particle size distributions (M4); c) combining the two-point correlation, lineal path distributions and pore/particle size distributions (M5). All simulated annealing methods (M1-M5) also optimize the porosity.

Finally, transport modeling of each of the reconstructions was undertaken. Effective transport properties were obtained for the thermal, ionic and electrical conductivity using a finite volume solver. Effective diffusivity for oxygen was computed using a random walk simulation technique in order to account for Knudsen effects.

Thermal conductivity: All simulated annealing reconstruction techniques tested were relatively accurate in predicting the effective thermal conductivity, ranging from a 3.5% relative error (for M5) up to a 6.3% error (M1), indicating that the prediction of thermal conductivity is fairly insensitive to the microstructure of the reconstruction.

Ionic conductivity: The most accurate reconstruction method is still M5 with a relative error of 2.8%, while the least accurate reconstruction method was M1 with a relative error of 9.1%. However, the second most accurate reconstruction method was M2 with a relative error of 3.2%, which seems to indicate that optimization of the lineal path distribution is the most important for an accurate prediction of ionic conductivity.

Electrical conductivity: All simulated annealing reconstructions underpredicted the electrical conductivity by a relative error of 23% (M5) up to 39% (M1). The high degree of error is attributed to the use of spheres in reconstruction which tends to cause narrow constrictions where spheres touch each other and also causes the formation of many dead-end channels where a sphere is attached on the edge of an agglomerate.

Oxygen diffusivity: In contrast to the results of the electrical conductivity simulations, all of the reconstructions *over*predict the oxygen diffusivity by relative errors of 78% (M3) up to 110% (M2). As diffusion occurs in the void phase (diffusing around the spherical particles), it is evident that the use of spheres for the solid phase reduces the impedance to diffusion in the void phase. Use of spheres as a base structural unit means that there will be fewer dead-end pores present in the void phase than would be expected from the experimental SEM images.

5.2 Future Work

To alleviate the issues caused by the spherical structuring units in sphere-based simulated annealing, a few recommendations can be made. One possibility is to conduct the sphere-based simulated annealing as described in this thesis up to some arbitrary point (say, 10% of initial configuration energy), at which point the algorithm switches to a different-phase neighbour pixel-swapping simulated annealing technique. This process would begin with a reconstruction with high sphericity, but by abandoning the spherical structuring unit, and using a pixel instead, the structure would gradually have its sphericity reduced.

A second avenue to investigate is the use of a morphological closing (dilation followed by erosion of the solid phase). This would help 'smooth' out some of the inherent sphericity in the reconstruction by closing the small gaps between spheres and widening the 'bridges' between spheres where the carbon-carbon connection may be initially very narrow. By connecting the solid phase better, the process would incidentally also reduce some of the additional pathways available in the pore phase. In this way, the morphological closing would theoretically both increase the simulated effective electrical conductivity and reduce the simulated effective oxygen diffusivity, in such a way that both are closer to that of the reference. This process may also more closely resemble the actual physical structure formed by the ionomer around the spherical carbon particles, as in reality, it is unlikely that the ionomer forms a uniform thickness film around the carbon spheres and agglomerates. A morphological closing will likely have the secondary, undesired effect of reducing the porosity, so post-closing adjustments may need to be made in order to maintain the correct porosity. The simplest way to accomplish this would be to find random solid voxels on the interface edge and convert them back to pores.

Finally, it may be possible to combine the sphere-based simulated annealing method with two above recommendations, where the process could look like:

1. Sphere-packing initialization of structure.
2. Sphere-based simulated annealing optimization using M5.
3. Morphological closing of solid phase.
4. Pixel-based simulated annealing using the different-phase neighbour method proposed by Pant et al [24].

Steps 1 and 2 are the same process as was conducted for this thesis. Step 3 uses the morphological closing to reduce the sphericity of the reconstruction but may adversely affect the correlations and metrics that were optimized in step 2. Therefore, step 4 would fine-tune the structure once again, using the pixel as the base structuring unit in order to avoid the high sphericity inherent in reconstruction generated using sphere-based simulated annealing.

References

[1] X. Li, K. Jiao, in: B.D. Fath, S.E. Jorgensen (Eds.), Managing Air Quality and Energy Systems, 2nd ed., CRC Press, 2020, pp. 593–605.

[2] M. ELSayed Youssef, Moataz H.Khalil, Khairia E.ALNAdi, INTERNATIONAL JOURNAL ON Advances in Information Sciences and Service Sciences 2 (2010) 61–66.

[3] D. Banham, J. Zou, S. Mukerjee, Z. Liu, D. Yang, Y. Zhang, Y. Peng, A. Dong, J Power Sources 490 (2021) 229515.

[4] L. Chen, G. Wu, E.F. Holby, P. Zelenay, W.Q. Tao, Q. Kang, Electrochim Acta 158 (2015) 175–186.

[5] X. Li, Principles of Fuel Cells, 2006.

[6] G. Inoue, K. Park, M. So, N. Kimura, Y. Tsuge, Int J Hydrogen Energy 47 (2022) 12665–12683.

[7] S.H. Kim, H. Pitsch, J Electrochem Soc 156 (2009) B673.

[8] K.J. Lange, P.C. Sui, N. Djilali, J Power Sources 208 (2012) 354–365.

[9] S. Vierrath, F. Güder, A. Menzel, M. Hagner, R. Zengerle, M. Zacharias, S. Thiele, J Power Sources 285 (2015) 413–417.

[10] M.H. Shojaeefard, G.R. Molaeimanesh, M. Nazemian, M.R. Moqaddari, Int J Hydrogen Energy 41 (2016) 20276–20293.

[11] M. So, K. Park, T. Ohnishi, M. Ono, Y. Tsuge, G. Inoue, Int J Hydrogen Energy 44 (2019) 32170–32183.

[12] K. Jiao, X. Li, Prog Energy Combust Sci 37 (2011) 221–291.

[13] J.P. Owejan, J.E. Owejan, W. Gu, T.A. Trabold, T.W. Tighe, M.F. Mathias, J Electrochem Soc 157 (2010) B1456.

[14] Q. Li, Z. Liu, Y. Sun, S. Yang, C. Deng, Processes 9 (2021) 1–21.

[15] J. Zhao, A. Ozden, S. Shahgaldi, I.E. Alaefour, X. Li, F. Hamdullahpur, Energy 150 (2018) 69–76.

[16] R. Singh, A.R. Akhgar, P.C. Sui, K.J. Lange, N. Djilali, J Electrochem Soc 161 (2014) F415–F424.

[17] C. Ziegler, S. Thiele, R. Zengerle, J Power Sources 196 (2011) 2094–2097.

[18] S. Thiele, R. Zengerle, C. Ziegler, Nano Res 4 (2011) 849–860.

[19] K.J. Lange, P.C. Sui, N. Djilali, J Power Sources 196 (2011) 3195–3203.

[20] K.J. Lange, P.-C. Sui, N. Djilali, J Electrochem Soc 157 (2010) B1434.

[21] M.G. Politis, E.S. Kikkinides, M.E. Kainourgiakis, A.K. Stubos, Microporous and Mesoporous Materials 110 (2008) 92–99.

[22] M. So, K. Park, Y. Tsuge, G. Inoue, J Electrochem Soc 167 (2020) 013544.

[23] W. Wu, F. Jiang, Int J Hydrogen Energy 39 (2014) 15894–15906.

[24] L.M. Pant, S.K. Mitra, M. Secanell, Phys Rev E Stat Nonlin Soft Matter Phys 90 (2014) 1–13.

[25] L.M. Pant, S.K. Mitra, M. Secanell, Phys Rev E Stat Nonlin Soft Matter Phys 92 (2015) 1–12.

[26] L.M. Pant, Stochastic Characterization and Reconstruction of Porous Media, Doctor of Philosophy, University of Alberta, 2016.

[27] J. Kang, K. Moriyama, S.H. Kim, J Power Sources 325 (2016) 752–761.

[28] L. Mosser, O. Dubrule, M.J. Blunt, Phys Rev E 96 (2017).

[29] A. Gayon-Lombardo, L. Mosser, N.P. Brandon, S.J. Cooper, NPJ Comput Mater 6 (2020) 1–11.

[30] R. Shams, M. Masihi, R.B. Boozarjomehry, M.J. Blunt, J Pet Sci Eng 186 (2020) 106794.

[31] P.C. Sui, X. Zhu, N. Djilali, Electrochemical Energy Reviews 2 (2019) 428–466.

[32] Von D. A. G. Bruggeman, Ann Phys (1935).

[33] P.K. Das, X. Li, Z.S. Liu, Appl Energy 87 (2010) 2785–2796.

[34] R. Zhang, P. He, F. Bai, L. Chen, W.Q. Tao, Int J Green Energy 18 (2021) 1147–1160.

[35] M.S. Ismail, D.B. Ingham, L. Ma, K.J. Hughes, M. Pourkashanian, Energy 122 (2017) 420–430.

[36] D.Y.C. Leung, J. Xuan, Micro & Nano-Engineering of Fuel Cells (Sustainable Energy Developments), 2015.

[37] B. Jacob, Theory and Applications of Transport in Porous Media Modeling Phenomena of Flow and Transport in Porous Media, 2018.

[38] D.A.G. Bruggeman, Ann Phys 416 (1935) 636–664.

[39] G. Inoue, K. Yokoyama, J. Ooyama, T. Terao, T. Tokunaga, N. Kubo, M. Kawase, J Power Sources 327 (2016) 610–621.

[40] P. Abraham B, K. Murugavel K, Electrochim Acta 389 (2021) 138793.

[41] S. Torquato, Random Heterogeneous Materials, Springer, Princeton, 2002.

[42] P. Čapek, M. Veselý, B. Bernauer, P. Sysel, V. Hejtmánek, M. Kočiřík, L. Brabec, O. Prokopová, Comput Mater Sci 89 (2014) 142–156.

[43] B. Lu, S. Torquato, Phys Rev A (Coll Park) 45 (1992) 922–929.

[44] B. Lu, S. Torquato, Phys Rev A (Coll Park) 45 (1992) 7292–7301.

[45] P. Čapek, V. Hejtmánek, J. Kolafa, L. Brabec, Transp Porous Media 88 (2011) 87–106.

[46] P. Čapek, V. Hejtmánek, L. Brabec, A. Zikánová, M. Kočiřík, Transp Porous Media 76 (2009) 179–198.

[47] B. Lu, S. Torquato, J Chem Phys 98 (1993) 6472–6482.

[48] M.M. Tomadakis, S. V. Sotirchos, AIChE Journal 39 (1993) 397–412.

[49] J. Zhao, S. Shahgaldi, I. Alaefour, S. Yang, X. Li, Int J Hydrogen Energy 43 (2018) 3776–3785.

[50] J.R. Nimmo, in: Encyclopedia of Soils in the Environment, Elsevier, 2005, pp. 295–303.

[51] H. Giesche, Particle & Particle Systems Characterization 23 (2006) 9–19.

[52] H.L. Ritter, L.C. Drake, Industrial & Engineering Chemistry Analytical Edition 17 (1945) 782–786.

[53] E.W. Washburn, National Academy of Sciences 7 (1921) 115–116.

[54] Yu.M. Volfkovich, V.S. Bagotzky, J Power Sources 48 (1994) 327–338.

[55] J. Gostick, Z. Khan, T. Tranter, M. Kok, M. Agnaou, M. Sadeghi, R. Jervis, J Open Source Softw 4 (2019) 1296.

[56] M. Eikerling, J Electrochem Soc 153 (2006) E58.

[57] J. Zhao, Catalyst Layers in Polymer Electrolyte Membrane Fuel Cells: Formation, Characterization and Performance, 2019.

[58] P. Adler, Porous Media Geometry and Transports, Butterworth-Heinemann, Boston, 1992.

[59] S. Brunauer, P.H. Emmett, E. Teller, J Am Chem Soc 60 (1938) 309–319.

[60] W.E. Lorensen, H.E. Cline, Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1987 21 (1987) 163–169.

[61] T. Lewiner, H. Lopes, A.W. Vieira, G. Tavares, Journal of Graphics Tools 8 (2003) 1–15.

[62] J. Zhao, S. Shahgaldi, A. Ozden, I.E. Alaefour, X. Li, F. Hamdullahpur, Int J Energy Res 43 (2019) 3011–3019.

[63] S. Thiele, R. Zengerle, C. Ziegler, ECS Meeting Abstracts MA2010-01 (2010) 1419–1419.

[64] A.G. Star, T.F. Fuller, Chem Eng Sci 201 (2019) 309–318.

[65] H. Schulenburg, B. Schwanitz, N. Linse, G.G. Scherer, A. Wokaun, J. Krbanjevic, R. Grothausmann, I. Manke, Journal of Physical Chemistry C 115 (2011) 14236–14243.

[66] T. Hutzenlaub, J. Becker, R. Zengerle, S. Thiele, J Power Sources 227 (2013) 260–266.

[67] M. Klingele, R. Zengerle, S. Thiele, J Power Sources 275 (2015) 852–859.

[68] S. Ghosh, H. Ohashi, H. Tabata, Int J Hydrogen Energy 40 (2015) 15663–15671.

[69] W.K. Epting, J. Gelb, S. Litster, Adv Funct Mater 22 (2012) 555–560.

[70] R. Lin, M. Dong, S. Lan, M. Lou, Energy 262 (2023).

[71] W. Song, H. Yu, L. Hao, Z. Miao, B. Yi, Z. Shao, Solid State Ion 181 (2010) 453–458.

[72] C.Y. Jung, S.K. Kim, S.J. Lee, S.C. Yi, Electrochim Acta 211 (2016) 142–147.

[73] Z. Yu, R.N. Carter, J Power Sources 195 (2010) 1079–1084.

[74] A. Berson, H.W. Choi, J.G. Pharoah, Phys Rev E Stat Nonlin Soft Matter Phys 83 (2011) 1–12.

[75] M. Sabharwal, L.M. Pant, N. Patel, M. Secanell, J Electrochem Soc 166 (2019) F3065–F3080.

[76] N.A. Siddique, F. Liu, Electrochim Acta 55 (2010) 5357–5366.

[77] U.R. Salomov, E. Chiavazzo, P. Asinari, Computers and Mathematics with Applications 67 (2014) 393–411.

[78] Y. Gao, W. Qu, R. Zhu, J Power Sources 482 (2021) 228917.

[79] T. Sousa, C.M. Rangel, Journal of Solid State Electrochemistry 20 (2016) 541–554.

[80] C.L.Y. Yeong, S. Torquato, Phys Rev E 58 (1998) 224–233.

[81] L. Ingber, Math Comput Model 18 (1993) 29–57.

[82] L.M. Pant, M. Sabharwal, S. Mitra, M. Secanell, ECS Trans 69 (2015) 105–120.

[83] Y. Zhang, Stochastic Reconstruction and Morphological Studies of Catalyst Layers of Proton Exchange Membrane Fuel Cells, University of Waterloo, 2021.

[84] X. Liu, K. Park, M. So, S. Ishikawa, T. Terao, K. Shinohara, C. Komori, N. Kimura, G. Inoue, Y. Tsuge, Journal of Power Sources Advances 14 (2022) 100084.

[85] N.P. Siegel, M.W. Ellis, D.J. Nelson, M.R. Von Spakovsky, J Power Sources 115 (2003) 81–89.

[86] Y. Wang, S. Wang, S. Liu, H. Li, K. Zhu, Electrochim Acta 318 (2019) 770–782.

[87] T. Berning, N. Djilali, J Power Sources 124 (2003) 440–452.

[88] F. Barbir, in: Elsevier Science & Technology, United States, 2012.

[89] K. Jiao, Experimental and Modelling Studies of Cold Start Processes in Proton Exchange Membrane Fuel Cells, 2011.

[90] P.P. Mukherjee, C.-Y. Wang, J Electrochem Soc 154 (2007) B1121.

[91] R. Barbosa, B. Escobar, U. Cano, J. Ortegon, V.M. Sanchez, Int J Hydrogen Energy 41 (2016) 19399–19407.

[92] A.P. Saab, F.H. Garzon, T.A. Zawodzinski, J Electrochem Soc 149 (2002) A1541.

[93] S.G. Kandlikar, Z. Lu, Appl Therm Eng 29 (2009) 1276–1280.

[94] O.S. Burheim, H. Su, H.H. Hauge, S. Pasupathi, B.G. Pollet, Int J Hydrogen Energy 39 (2014) 9397–9408.

[95] V.M. Barragán, M.A. Izquierdo-Gil, J.C. Maroto, P. Antoranz, S. Muñoz, Int J Heat Mass Transf 176 (2021).

[96] M. Ahadi, A. Putz, J. Stumper, M. Bahrami, J Power Sources 354 (2017) 215–228.

[97] Z. Yu, R.N. Carter, J. Zhang, (2012) 557–565.

[98] S. Salari, J. Stumper, M. Bahrami, Int J Hydrogen Energy 43 (2018) 16704–16718.

[99] G. Inoue, M. Kawase, J Power Sources 327 (2016) 1–10.

[100] J. Becker, C. Wieser, S. Fell, K. Steiner, Int J Heat Mass Transf 54 (2011) 1360–1368.

[101] M. Sabharwal, L.M. Pant, A. Putz, D. Susac, J. Jankovic, M. Secanell, Fuel Cells 16 (2016) 734–753.

[102] W. Zheng, S.H. Kim, J Electrochem Soc 165 (2018) F468–F478.

[103] Y. Gao, X. Zhang, Electrochim Acta 218 (2016) 101–109.

[104] L.B. (Leonard B. Loeb, The Kinetic Theory of Gases., Second edition., McGraw-Hill book company, inc., New York and, 1934.

[105] W.G. Pollard, R.D. Present, Physical Review 73 (1948) 762–774.

[106] Y. Gao, Int J Heat Mass Transf 88 (2015) 122–132.

[107] J.C. Ferguson, A. Borner, F. Panerai, S. Close, N.N. Mansour, Comput Mater Sci 203 (2022) 111030.

[108] J. Gostick, M. Aghighi, J. Hinebaugh, T. Tranter, M.A. Hoeh, H. Day, B. Spellacy, M.H. Sharqawy, A. Bazylak, A. Burns, W. Lehnert, A. Putz, Comput Sci Eng 18 (2016) 60–74.

[109] M. El Hannach, M. Prat, J. Pauchet, Int J Hydrogen Energy 37 (2012) 18996–19006.

[110] M. Aghighi, J. Gostick, J Appl Electrochem 47 (2017) 1323–1338.

[111] M. Aghighi, M.A. Hoeh, W. Lehnert, G. Merle, J. Gostick, J Electrochem Soc 163 (2016) F384–F392.

[112] S. Jung, M. Sabharwal, A. Jarauta, F. Wei, M. Gingras, J. Gostick, M. Secanell, J Electrochem Soc 168 (2021) 064501.

[113] T. Kruger, H. Kusumaatmaja, A. Kuzmin, O. Shardt, S. Goncalo, E.M. Viggen, The Lattice Boltzmann Method, Principles and Practice, 2017.

[114] K. Suankulova, Correlation Functions in Determining the Structure of Chromosomes, 2013.

[115] E. Patelli, G. Schuëller, Comput Mater Sci 45 (2009) 536–549.

[116] E. Patelli, G.I. Schuëller, Comput Optim Appl 53 (2012) 903–931.

[117] R. Okuta, Y. Unno, D. Nishino, S. Hido, C. Loomis, in: Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-First Annual Conference on Neural Information Processing Systems (NIPS), 2017, pp. 1–7.

[118] C.J. Gommes, Y. Jiao, S. Torquato, Phys Rev E Stat Nonlin Soft Matter Phys 85 (2012) 1–16.

[119] R.P. Dougherty, K.-H. Kunzelmann, in: Microscopy & Microanalysis 2007 Meeting, Ft. Lauderdale, 2007.

[120] T.Y. Zhang, C.Y. Suen, Commun ACM 27 (1984) 236–239.

[121] T.C. Lee, R.L. Kashyap, C.N. Chu, CVGIP: Graphical Models and Image Processing 56 (1994) 462–478.

[122] S. van der Walt, J.L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J.D. Warner, N. Yager, E. Gouillart, T. Yu, PeerJ 2 (2014) e453.

[123] S.K. Lam, A. Pitrou, S. Seibert, in: Proceedings of LLVM-HPC 2015: 2nd Workshop on the LLVM Compiler Infrastructure in HPC - Held in Conjunction with SC 2015: The International Conference for High Performance Computing, Networking, Storage and Analysis, Association for Computing Machinery, 2015.

[124] D.A. Coker, S. Torquato, J Appl Phys 77 (1995) 6087–6099.

[125] C.R. Harris, K.J. Millman, S.J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N.J. Smith, R. Kern, M. Picus, S. Hoyer, M.H. van Kerkwijk, M. Brett, A. Haldane, J.F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, T.E. Oliphant, Nature 585 (2020) 357–362.

[126] P. Virtanen, R. Gommers, T.E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S.J. van der Walt, M. Brett, J. Wilson, K.J. Millman, N. Mayorov, A.R.J. Nelson, E. Jones, R. Kern, E. Larson, C.J. Carey, İ. Polat, Y. Feng, E.W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E.A. Quintero, C.R. Harris, A.M. Archibald, A.H. Ribeiro, F. Pedregosa, P. van Mulbregt, A. Vijaykumar, A. pietro Bardelli, A. Rothberg, A. Hilboll, A.

Kloeckner, A. Scopatz, A. Lee, A. Rokem, C.N. Woods, C. Fulton, C. Masson, C. Häggström, C. Fitzgerald, D.A. Nicholson, D.R. Hagen, D. v. Pasechnik, E. Olivetti, E. Martin, E. Wieser, F. Silva, F. Lenders, F. Wilhelm, G. Young, G.A. Price, G.-L. Ingold, G.E. Allen, G.R. Lee, H. Audren, I. Probst, J.P. Dietrich, J. Silterra, J.T. Webber, J. Slavič, J. Nothman, J. Buchner, J. Kulick, J.L. Schönberger, J.V. de Miranda Cardoso, J. Reimer, J. Harrington, J.L.C. Rodríguez, J. Nunez-Iglesias, J. Kuczynski, K. Tritz, M. Thoma, M. Newville, M. Kümmerer, M. Bolingbroke, M. Tartre, M. Pak, N.J. Smith, N. Nowaczyk, N. Shebanov, O. Pavlyk, P.A. Brodtkorb, P. Lee, R.T. McGibbon, R. Feldbauer, S. Lewis, S. Tygier, S. Sievert, S. Vigna, S. Peterson, S. More, T. Pudlik, T. Oshima, T.J. Pingel, T.P. Robitaille, T. Spura, T.R. Jones, T. Cera, T. Leslie, T. Zito, T. Krauss, U. Upadhyay, Y.O. Halchenko, Y. Vázquez-Baeza, Nat Methods 17 (2020) 261–272.

[127] P.E. Barrett, P.L. Shopbell, M.C. Britton, R. Ebert, P. Barrett, J. Hunter, J.T. Miller, J.-C. Hsu, P. Greenfield, Matplotlib-A Portable Python Plotting Package Astronomical Python View Project Magnetic CVs View Project Matplotlib-A Portable Python Plotting Package, 2005.

[128] J.C. Ferguson, F. Panerai, A. Borner, N.N. Mansour, SoftwareX 7 (2018) 81–87.

[129] H.A. van der Vorst, SIAM Journal on Scientific and Statistical Computing 13 (1992) 631–644.

[130] P.E. Khizhnyak, A. V. Chechetkin, A.P. Glybin, Journal of Engineering Physics 37 (1979) 1073–1075.

[131] R. Sigwadi, M.S. Dhlamini, T. Mokrani, F. Ṇemavhola, P.F. Nonjola, P.F. Msomi, Heliyon 5 (2019).

[132] Y.A. Çengel, J.M. Cimbala, A.J. (Afshin J. Ghajar, Fundamentals of Thermal-Fluid Sciences, n.d.

[133] T.A.A. Broadbent, The Mathematical Gazette 41 (1957) 231–231.

[134] Journal of the Society of Chemical Industry 57 (1938) 901–901.

[135] X. Michalet, Phys Rev E 82 (2010) 041914.

[136] A. Rahman, Physical Review 136 (1964) A405–A411.

[137] S. Chandrasekhar, Rev Mod Phys 15 (1943) 1–89.

[138] E.W. Weisstein, MathWorld--A Wolfram Web Resource (2023).

[139] P. Ruzzante, X. Li, Energy (2023) 128200.