

**FUNDAMENTAL DEVELOPMENTS IN VOLUMETRIC
ADDITIVE MANUFACTURING**

by

CHARLES M. RACKSON

B.S., B.S., & B.A., Seattle University, 2015

M.S., University of Colorado Boulder, 2018

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Electrical, Computer, and Energy Engineering
2022

Committee Members:

Robert McLeod, Chair

Maxim Shusteff

Stephanie Bryant

Rafael Piestun

Shu-Wei Huang

RACKSON, CHARLES M. (Ph.D.)

FUNDAMENTAL DEVELOPMENTS IN VOLUMETRIC ADDITIVE MANUFACTURING

Thesis directed by Prof. Robert McLeod

Abstract

The relatively new field of Volumetric Additive Manufacturing (VAM) performs printing into a fixed volume of resin in a single lithographic step. This has multiple advantages over multi-step additive manufacturing methods. Notably, non-contiguous parts become possible, layering effects are eliminated, print speeds are improved by orders of magnitude, and the lack of material movement mid-print enables high-viscosity or even solid resins, opening the doors to a host of new material properties. However, prior to this work, VAM suffered from multiple limitations that both restricted its scope of application and prevented its practical use as a manufacturing technique.

Firstly, there was a lack of tomographic calculation techniques to provide sufficient optical dose control for a variety of VAM applications. From printing 3D parts without shape error, to controlling grayscale optical dose for e.g. printing optics or functionally graded materials, control over dose is critical in VAM. We present a simple and effective approach to VAM image computation that significantly improves over prior methods. We demonstrate the flexibility of the approach by extending it to include material response models, optical models including beam shape and occlusions for overprinting, as well as to the problem of controlling grayscale dose throughout a print.

Secondly, VAM was limited to printing into small cylindrical sample packages. We present a tomosynthetic VAM printing geometry that allows for printing into flat surfaces. This lends itself to applications such as the manufacturing of flat microfluidic chips, printing into well plates, or any other printing application requiring beam access through a surface. It further allows for the

patterning of arbitrarily large sheets. We discuss theory and image computation for this printing geometry, and we present experimental proof of concept prints.

Thirdly, VAM ubiquitously suffered from large striations – similar in appearance to layering, and on the order of print feature size – impacting the homogeneity and shape-accuracy of printed parts, and presenting a significant cosmetic defect. We present evidence that the material nonlinearity upon which VAM relies also drives striations via a self-writing waveguide (SWW) effect. We demonstrate a simple and effective method of dramatically reducing striations via a latent-cure, flood-exposure step. We additionally show this to drastically shorten the period from initial gelation to print completion. This mitigates the problem of partial-print sinking in low power or low viscosity prints, thus further expanding resin options and increasing the efficacy of low-cost VAM printers.

We discuss the fundamentals of VAM optics and materials, and the constraints on both that inform the design approach of VAM printers and resins. We present measurement methods for material selection, and we detail the design and alignment of a particular printer.

Dedication

To Andres Takach, Dutch Siedentopf, and Woo-Joong Kim. And to my wife, my sister, and my parents, for their support.

Acknowledgements

I would like to express my sincere thanks to everyone who helped me during my PhD. I am especially grateful to three people. First, to my advisor, Robert McLeod, for his guidance, encouragement, support, and for everything he taught me along the way. Second, to Amy Sullivan; she was effectively a second advisor in the group, and her enthusiasm, expertise, professionalism, and kindness were inspiring, were incredibly helpful, and made the lab a good place to be. And third, to Maxim Shusteff for his support, for his dedication to building community and collaboration, and for leading a welcoming and encouraging group at Lawrence Livermore National Laboratory.

I would like to thank the rest of my committee members for their guidance: Stephanie Bryant, Rafael Piestun, and Shu-Wei Huang.

Next, I wish to thank my collaborators and colleagues, without whom I would have accomplished far less, and with much less enjoyment. They include: Kyle Champley, Joseph Toombs, Erika Fong, Vishal Bansal, Hayden Taylor, Martin De Beer, Caitlyn Cook, Bryan Moran, Johanna Schwartz, Chi Chung Li, John Karnes, Todd Weisgraber, James Oakdale, Allison Kaczmarek, Yihan Huang, Daniel Mittelberger, Dominique Porcincula, Doran Harrop, Vanessa Huizar, Joe Dragavon, Dylan Domaille, Roland Faller, Kelvin Wagner, and Parag Shah.

I would like to thank my lab mates: Jacob Friedlein, David Glugla, Martha Bodine, Archish Muralidharan, Marvin Alim, Arlo Marquez-Grap, Andrew Sias, Gabe Seymour, Megan Renny, Darren Liu, Callie Higgins, Johnny Hergert, David Miller, Camila Uzcategui, and Jamie Kowalski.

Thanks to those who have helped me, and made my time in graduate school a good one: my family, Michael Grayson, Shane Verploegh, Bohan Zhang, Josep Fargas, NV Vyjayanthi, Elizabeth

Strong, Jan Bartos, Philip Zurek, Greg and Christina Morrin, Ryan Richardson, Scott Johnson, Brant Rumberger, Fabrizio Giorgetta, Esther Baumann, William Raseman, Guillermo Nevett, Steve Guard, Julia Clarson, Chloë Zehr, Eric Rosenthal, Kshitij Sadasivan, Polly Murray, Sean Horvath, Vesta Zhelyaskova, Mikaela DeRousseau, Maxwell Duffy, Jeff Budrys, Emily Hannah, Zev Underwood, Todd Graveson, Daniel Pyle, Benjamin Carberry, Julia Burak, and Alex Podschwit.

Lastly, thanks to my funding sources: CU Boulder, The U.S. Department of Energy, The National Institutes of Health, and the National Science Foundation.

Contents

Chapter	
1	Introduction 1
2	Printer Design 6
2.1	Optical Design Requirements and Fundamental Constraints 6
2.1.1	Depth of Focus 6
2.1.2	Number of Resolvable Spots 8
2.1.3	Power and Etendue 10
2.2	Optical Design Approach 10
2.2.1	Laser and Fiber Light Source 11
2.2.2	Pattern Generation 12
2.2.3	DMD Overview and Diffraction Efficiency 15
2.2.4	Relay Optics 18
2.2.5	System Overview 19
2.2.6	Alignment Methods 20
2.3	Summary 23
3	Resin Design 25
3.1	Absorption 26
3.2	Photoinitiator Concentration and Molar Absorptivity 27
3.3	Viscosity 28

3.4	Measuring Molar Absorptivity	29
3.5	Resin Preparation & Reuse	33
3.6	Post Processing	34
3.7	Photo Product Fluorescence	35
3.8	Conclusions	37
4	Image Computation	38
4.1	Binary-VAM Image Computation	40
4.1.1	FBP and The Problem of Negative Intensity	41
4.1.2	A Binary VAM Problem Statement and Prior Approaches	44
4.1.3	Object Space Image Computation for Binary VAM Printing	47
4.1.4	Scaling Between Independent Reconstruction Slices	61
4.1.5	Scaling to Partial Rotations	67
4.1.6	Scaling Tomographic Reconstructions to Real Optical Units	68
4.1.7	Controlling Print Dynamics	69
4.1.8	Including Material Response	73
4.1.9	Binary Printing with Inhibition	73
4.1.10	Additional Printing Geometries and Optical Models for Binary VAM	74
4.1.11	Image Computation for More Complex Projection Models	74
4.1.12	Experimental Demonstrations of OSMO	82
4.2	Image Computation for Grayscale-Dose VAM Printing	86
4.3	FBP Enabled by Inhibition	91
4.3.1	Compensating an Image-Space Shift in Object-Space	92
4.3.2	Shifted FBP Versus OSMO for Binary VAM	97
4.4	Future Work	98
4.4.1	Relaxed Voxel Constraints	99
4.4.2	Accounting for Nonreciprocity	100

4.4.3	Accounting for Diffusion	103
4.4.4	OSMO Extrapolation	105
4.4.5	More Projection Directions	106
4.5	Conclusions	107
5	Tomosynthesis	109
5.1	Tomosynthetic Tomography Fundamentals	110
5.2	Applying OSMO to Tomosynthesis	113
5.3	Tomosynthesis Experimental Demonstration	117
5.4	Future Work	119
5.4.1	Light Coupling	119
5.4.2	Tomosynthesis Optics	121
5.4.3	Computation for Stitched VAM	123
5.4.4	Grayscale Tomosynthetic VAM	124
5.5	Conclusions	125
6	Striations and Latent Image VAM	126
6.1	Striations in VAM	128
6.2	Latent Cure Method	131
6.2.1	Dose Contrast Reduction Upon Latent Cure	132
6.2.2	Timing the Stopping of VAM Prints for Latent Cure	135
6.2.3	Details of Illumination and of Example Prints	135
6.3	Latent Cure to Avoid Partial-Print Sinking	137
6.4	Conclusions	138
7	Conclusions & Summary of Future Work	140

Bibliography	143
---------------------	------------

Appendix

A OSMO Sample Code	149
---------------------------	------------

Figures

Figure

- 1.1 Tomographic VAM printing of a 3D geometry. **A:** Rotating cylindrical vial containing liquid resin. The surrounding rectangular cuvette holds a refractive index matching bath. 405 nm writing light is incident on the resin. **B:** Schematic of a 2D slice-region of the printing setup. **C:** After about one minute, a part solidifies due to accumulated optical energy. It is removed from the surrounding liquid resin, then it is washed and post-cured. 2

- 2.1 Illustration of the dependence of voxel size on beam width and divergence. For this illustration, the intensity distribution across the width of the beams is ignored, and only beam width is considered. For a more detailed analysis, see the chapter on image computation. **A:** Schematic of a 3D tomographic back projection showing the intersection of non diverging beams projected from square DMD pixels. The intersection of these beams forms a voxel whose shape is approximately cylindrical. As the number of projection angles approaches infinity, the voxel shape approaches an exact cylinder. **B:** Two dimensional view of a single slice-region of the reconstruction showing the intersection of non diverging beams. The circular (in 2D) voxel that is formed is outlined in black. **C:** Diverging Gaussian beams whose focus is set to the center of the vial. Voxels formed near the center of the vial are small in diameter - equal to the waist width of the focused beam. **D:** The voxels near the edge, however, are larger and not necessary circular, as they are formed by the intersection of defocused beams. 7
- 2.2 Number of resolvable spots at a wavelength of 405 nm that can fit across the width of a vial when the DOF is equal to the vial diameter. **Left:** Since the vial is round, when the DOF is equal to the depth through the diameter of the vial, it is also equal the width of the vial. A DOF equal to the vial diameter yields an $f\#$. That $f\#$ also defines minimum beam spot size $2\omega_0$. Thus, the number of resolvable spots N_{SPOTS} across the width of a vial is equal to $\text{DOF}/(2\omega_0) = \sqrt{\pi/2}\sqrt{\text{DOF}/\lambda}$. **Right:** For 405nm light, we can plot this relationship (blue plot) to see that N_{SPOTS} goes as $\sqrt{\text{DOF}}$ or, equivalently, as $\sqrt{\text{vial diameter}}$. The minimum spot size of a beam follows the same trend, as is plotted in orange. 9
- 2.3 A burnt fiber core caused by focusing a 1 W, 405 nm laser onto a $70 \times 70 \mu\text{m}$ fiber with a polymer cladding. The laser ran for many hours without burning the fiber, suggesting dust as a cause of failure. 12

- 2.4 DMD micromirror layout, tilt directions, and light input and output directions. **A:** A Cartesian grid of micro-mirrors. The line about which each mirror pivots is sketched with a black dashed line. Each mirror can tilt by $\pm 12^\circ$, where 0° would correspond to the plane of the mirror being parallel to the grid of mirrors. An input and output beam form a plane that is orthogonal to the mirror pivot planes. An end-on view of this plane is shown in red. That is, beams input to the DMD lie in a 45° plane that is orthogonal to the plane formed by the locus of micro mirrors. **B:** A sketch of light incident on and exiting a DMD *on* state pixel. The dashed blue lines are Cartesian axes, with one axis orthogonal to the plane of the DMD. The input and output beams, shown in purple, lie in a 45° plane to the DMD. The angle between an incident beam and the normal to the surface of the DMD is variable. The power efficiency of reflected beam depends on this angle, θ_{inc} . **C:** Micro mirrors act as a blazed grating. The path-length difference between beams going to different mirrors causes a relative phase delay, and thus interference. When choosing an input angle θ_{inc} , we seek to minimize destructive interference, thus maximizing power in the 0^{th} diffraction order. 16
- 2.5 Simulation of DMD diffraction. A linear phase ramp is applied to the DMD by light incident at θ_{inc} . Scanning through a range of input angles lets us see the effects on diffraction efficiency. Here, we assumed 405 nm light, and a micromirror pitch matching our Vialux V-9001 DMD ($7.6 \mu m$). **A:** A poorly chosen input angle of $\theta_{inc} = 37^\circ$. **B:** The resultant diffraction pattern. The 0^{th} order contains only about a quarter of the total power - a poor efficiency. Instead of the majority of the power going into a single order, most of it is split into four orders. Such an example from the lab is shown. **C:** The DMD phase that results from an optimum choice of $\theta_{inc} = 23.3^\circ$. **D:** The resultant diffraction pattern. Here, the power into a single order is maximized, with about 80% of the power going into the 0^{th} order. This is also observed in the lab, as shown. 17

- 2.6 The Design of a VAM printer used in this work. The laser is a multi-mode 3 W, 405 laser from Civil Laser. The fiber is a $100 \times 100 \mu m$ square-core fiber made by Ceramtec Optics. Note that a 0.22 NA, $70 \times 70 \mu m$ may also be used for larger sample depth of focus, although that comes at the cost of reduced power. The mirrors, M_1 and M_2 are AR-coated for high reflectivity at 405 nm. L_1 is a $20\times$ microscope objective with an effective focal length of 9 mm. The focal lengths of lenses $L_2 - L_4$ are: 90 mm, 51mm, and 100mm, respectively. A helium-neon laser is used to monitor print formation; its 633 nm wavelength does not affect the resin. Lens L_5 expands the beam such that its width spans the vial diameter. The lens L_6 then focuses that expanding beam into a screen. These shadowgram focal lengths are not critical, as they are used for qualitative print monitoring. 19
- 2.7 VAM printer details. **A:** Light path from fiber, to DMD, then towards print plane. Note that the light angles up towards the DMD. **B:** Print vial with resin, immersed in a bath of index-matched liquid. A rotation stage spins the vial within the index-match bath. Here, the index-match liquid was resin without photoinitiator, making it non-reactive to the 405 nm light. **C:** A photograph of a part just after print completion. The part is still immersed in resin, but its increased density causes increased refractive index which, for this resin, allows it to be seen by eye. **D:** A photograph of the shadowgram screen just after print completion. The shadowgram makes refractive index changes noticeable, which aids in determining the point in time at which the part is complete. 21

- 2.8 Simulated effects of misaligning the center of projections with the axis of rotation.
- A:** A 2D tomographic reconstruction (using FBP with 720 projection angles) of a target geometry. **B:** The reconstruction that results from laterally misaligning the axis of rotation by 0.4% the width of the projections. This corresponds to a lateral misalignment of only $4\ \mu\text{m}$ for a 1 cm print, but noticeably degrades reconstruction sharpness. **C:** The reconstruction that results from laterally misaligning the axis of rotation by 0.6% the width of the projections. Further reconstruction degradation is apparent. 23
- 3.1 Initial rough measurements of molar absorptivity for four photoinitiators that were eventually discounted from our search for an initiator with $\epsilon \approx 1$ at the candidate laser wavelengths of 405 nm and 473 nm. **A:** Irgacure 1173 was not absorptive enough. Note that the isobestic point - where initial and bleached absorptivities are equal - is near the 473 nm laser line. **B:** Irgacure 819 was appropriately sensitive only near 450 nm - not a candidate wavelength. No bleached data was collected for this sample. Note that the flat region of the plot between about 330 and 440 nm is saturated data. **C:** Irgacure 379 was too sensitive at 405 nm, and not sensitive enough at 473 nm. **D:** Irgacure 184 was too low in sensitivity at visible wavelengths. 31

- 3.2 Molar absorptivity data for Irgacure 907 collected with assistance from Darren Liu. **A:** Molar absorptivity computed from different concentration ranges of I907 in ethanol. The solid and dashed lines show data for unbleached and photobleached initiator, respectively. **B:** The data from (A) was narrowed to exclude wavelengths where measurements were saturated or very noisy, then the data was averaged to plot molar absorptivity of both photobleached and unbleached I907. The molar absorptivity of the radical photoinitiator, TEMPO, is also plotted in red (data from [1]). **C and D:** raw data showing absorbance and I907 concentration for two optical wavelengths. For these, the line fit was constrained to pass through the origin. Both data sets were from the "High Concentrations" set of dilutions. All points should ideally perfectly fit a line. Especially for longer wavelengths, spectrometer noise significantly degrades the data accuracy. However, we found the final level of accuracy of this data (B) to be sufficient for an initial photoinitiator selection. 32
- 3.3 Microscope images of I907 in PEGDA250 placed between glass slides, then gelled into 2D shapes with a 405 nm laser. **Top:** 575-625 nm emissions via a 561 nm stimulation source. Local stimulation intensity was approximately proportional to the 405 nm writing dose. **Bottom:** The same microscope, a Nikon AR1, was used in differential interference contrast mode to examine the refractive index change in the gelled structures. Thanks to Joseph Dragavon for his assistance with collecting this data. 36
- 4.1 Overview of tomographic VAM concept. **A:** Computed images, optically projected into a vial of photosensitive resin. **B:** Cross-section of accumulated optical dose distribution (normalized units of dose). **C:** Perfectly-gelled cross section of a "thinker"; resin gels where the optical dose is above a threshold, forming a 3D printed part that is removed from remaining liquid resin. Image from [2]. 41

- 4.2 Illustration of FBP and how circular projection geometries non-uniformly sample the frequency-space of an object. **A:** A 2D object comprised of four squares of constant value, with forward projection illustrated. **B:** Different projection angle sample different data from the object. **C:** Simply back-projecting the projection data yields a poor reconstruction of the original object. **D:** A spatial Fourier transform of the object in (A). Via the Projection Slice Theorem, each projection image from (B) is equal to the inverse Fourier transform of a line passing through the origin of (D), perpendicular to the projection direction. **E:** The Fourier-space linear ramp filter that must be applied to the data in (B) to correct for non-uniform sampling. **F:** The object in (A) is correctly reconstructed after filtering. This filtering and back-projecting is the FBP algorithm. 43
- 4.3 Dose reconstruction via forward projecting a binary model of a target geometry, then back projecting the resultant set of images. The reconstruction is blurred due to suppressed high spatial frequencies inherent to forward then back projecting. Instead of optimizing image intensities in projection-space to improve the reconstructed dose, we will instead optimize the model to improve dose. This first reconstruction serves as an initial step in the model-optimization algorithm shown in the next figure. Note that since all projection values are non-negative in this case (as no filtering nor thresholding was applied), they can be represented by intensities with no modification. Image from [2]. 48

4.4 Object space model optimization (OSMO) algorithm for binary-gelation VAM printing. **A:** One model-update iteration of the OSMO loop. Each model in the loop produces a dose-reconstruction by the process shown in (B) forward projection, clipping any resultant negative values to zero, then back-projecting and normalizing to form a dose-reconstruction. Negative values must be clipped to convert projections to physically realizable intensities. Note that, unlike all subsequent models, the initial model (Fig. 4.3) has no negative values in its projections, and thus does not require a clipping step. **A, Left Column:** A given i^{th} model (in this case, the initial model is pictured) produces a reconstruction. **A, Step 1:** Unwanted extra dose in the out-of-part regions of the reconstruction are computed and subtracted from the i^{th} model to form an intermediate model. Only the extra dose that lies above a lower threshold, D_l , is included. **A, Middle Column:** The intermediate model is then used to compute a dose-reconstruction. **A, Step 2:** The prior step of subtracting out-of-part values from the model, inadvertently reduces some desired dose in the in-part regions. This missing in-part dose is calculated by comparing the in-part reconstruction voxels to an upper threshold, D_h . These values are then added to the intermediate model to form a $(i+1)^{st}$ model, completing one model-update iteration. **B:** The model produced by m iterations, the dose reconstruction it produces, and the resultant perfectly gelled part. Computationally, the gel threshold is modeled by setting all voxels above a gelation-dose to unity, and setting all voxels below that dose to zero. Image from [2]. 50

4.5 Reconstruction Quality Metrics. **A:** A target geometry to be printed. **B:** An optimized dose-reconstruction, shown in normalized units of dose. The parameter values were $D_l = 0.6$ and $D_h = 0.8$. **C: Voxel Error Rate (VER)** converging to zero. **D:** The distribution of doses in the reconstruction for the in-part and out-of-part regions. Since the histograms do not overlap, VER is 0, and thus the gel threshold is above all out-of-part voxels, and below all in-part voxels. Therefore, this dose-distribution will print a part that theoretically exactly matches the target geometry. The **Process Window** is the distance between histograms, and represents room for error with optics or materials, without a change in the final gelled shape. As long as dose is applied such that the gel threshold lies within the process window, the part will gel in the correct shape. The **In-Part Dose Range (IPDR)** is the range of doses, as a fraction of maximum dose, that in-part voxels receive. Image from [2]. 53

4.6 Trade-off between process window size, and IPDR. **A, Top:** Binary VAM reconstruction with $D_l = 0.51$, and $D_h = 0.59$. **A, Bottom:** Histograms of in-part and out-of-part dose. **B, Top:** Binary VAM reconstruction with $D_l = 0.95$, and $D_h = 0.98$. **B, Bottom:** Histograms of in-part and out-of-part dose. Reconstruction-A has a process window over 3 times larger than that of Reconstruction-B, but an IPDR that is over 10 times larger. Note that choosing parameter values that try to simultaneously accomplish a small IPDR and a large process window (e.g. $D_l = 0.7$, and $D_h = 0.98$) does not work – trying to achieve such a reconstruction results in histogram overlap, and a non-zero VER. Image from [2]. 55

- 4.7 Frequency-filtering a target to generate an improved initial model. This reduces the number of algorithm iterations necessary to converge to good solution, compared to simply using the target as the initial model. For this target, with $D_l = 0.85$ and $D_h = 0.9$, required iterations were reduced by 25%. **A:** A target geometry. **B:** An initial reconstruction that results from using the target geometry as the initial model, a process illustrated in Fig. 4.4. **C:** The target geometry with a 2D frequency ramp filter applied. **D:** The initial reconstruction that results from using the filtered target as a model. Both initial models converge to qualitatively identical final dose reconstructions, with identical quality-metric values. Image from [2]. 57
- 4.8 Examples of OSMO convergence from different initial guesses: the original target geometry, the filtered target as described above, an array of ones, and an array of random numbers between zero and one. **A and B:** Convergence for the mesh structure shown in Fig. 4.7A with 720 projection angles. For the plot shown in (A), D_l was 0.85 and D_h was 0.9. For (B), D_l was 0.65 and D_h was 0.75. This is an unusual case in which filtering slows convergence to $VER = 0$. **C:** A different target geometry, yielding the convergence shown in **D** with $D_l = 0.65$ and $D_h = 0.75$ 58

- 4.9 Convergence comparison example with the previously-published CAL2019 method. **A:** Binary target geometry. **B:** VER convergence for the first 30 iterations of the OSMO and CAL2019 methods. Note that the initial VER value for OSMO is slightly smaller than for CAL2019, due to different initial guesses: 2D model filtering versus clipped FBP, respectively. While CAL2019 plateaus to a small VER near zero (10^{-5}), further iterations of OSMO drive the reconstruction beyond a VER of zero, to a positive process window size. The CAL2019 reconstruction had an IPDR of 0.12 and a process window size of -0.23%. **C and D:** OSMO dose reconstruction and histogram. The OSMO parameter values were $D_l = 0.83$, and $D_h = 0.9$, $VER = 0$, and the process window size was 5.9% of the maximum dose with an IPDR of 0.11. **E and F:** CAL2019 dose reconstruction and histogram. The histogram overlap yields a non-zero VER value with no positive process window size. Since $VER > 0$, a print would have voxels with gelation error. While these voxels would be few, the lack of a process window would result in further error in a physical implementation, due to imperfect precision in optics and materials. Image from [2]. 59
- 4.10 Modeled effects of total dose errors upon gelation for OSMO and CAL2019. Reconstructions shown in Fig. 4.9. **A:** Modeled gelation when the OSMO reconstruction (shown in Fig. 4.9C) is overexposed by 2.95%. Since this error falls within the reconstruction process window, no gelation errors occur, and the final simulated part matched the target geometry. **B:** CAL2019 gelation upon the same 2.95% overexposure. **C:** The extra gelation that occurs in (C) due to this error. **D:** Modeled gelation when the OSMO reconstruction (shown in Fig. 4.9C) is underexposed by 2.95%. Since this error falls within the reconstruction process window, no gelation errors occur. **E:** CAL2019 gelation upon the same 2.95% underexposure. **F:** the missing gelation in (E) that occurred due to underexposing the CAL2019 reconstruction by 2.95%. Inset shows the missing gelation size relative to small print features. Image from [2]. 61

- 4.11 Modeled effects of total gelation errors due to projector non-uniformity, with peak intensity at the center of the projections. Apodization functions were applied to the image sets computed in Fig. 4.9. before back projecting them to form reconstructions. **A:** Gaussian apodization (0.84 at edges, and unity at center) applied to each projector image computed in Fig. 4.9. **B:** Resultant gelation for CAL2019 - the unwanted and missing gelation voxels are detailed in **C**. For (B/C), VER was 0.16, and the threshold that minimized the number of gelation voxel errors was 0.87. **D:** Resultant gelation for OSMO with the apodization in (A) applied to all projections. **E:** Extra and missing gelation from (D). Here, VER was 1.8E-5, and the threshold that minimized the number of gelation voxel errors was 0.89. Image from [2]. 62
- 4.12 Modeled effects of total gelation errors due to projector non-uniformity, with peak intensity offset from the center of the projections. Apodization functions were applied to the image sets computed in Fig. 4.9. before back projecting them to form reconstructions. **A:** An apodization function, with a maximum value of 1 offset by 20% the total width of the function, with a value of 0.95 at the left edge, and a value of 0.83 at the right edge. **B:** Resultant gelation of applying (A) to the CAL2019 image set. **C:** The resultant extra and missing gelation voxels in (B). Here, VER was 0.15, and the threshold that minimized the number of gelation voxel errors was 0.87. **D:** Resultant gelation from applying (A) to the OSMO image set computed in Fig. 4.9. **E:** The resultant extra and missing gelation voxels in (D). Here, VER was 7.8E-5, and the threshold that minimized the number of gelation voxel errors was 0.89. Image from [2]. 63

- 4.13 Global dose histograms from different slice-scalings for reconstruction of a Thinker geometry. **A:** Scaling the independent 2D slices to have the same maximum reconstruction dose results in a process window of 2.8%. **B:** Scaling the slices to the minimum in-part dose increases the process window size to 3.3%. This is due to the increased process-window overlap caused by scaling the slices to have the high end of each process window equal in dose, thus increasing the global process window. 65
- 4.14 Method of plotting reconstruction dose details to better visualize slice-scaling choices. For each slice (vertical axis), a red dot indicates the maximum out-of-part dose. The length of the blue line is the IPDR. The left end of the a given line is the minimum in-part dose for that slice, and the right end is the maximum dose. 66
- 4.15 Per-slice reconstruction details from re-scaling slices of a Thinker geometry using the three criteria presented in this section. For each slice: the red dot is the largest out-of-part dose, the left end of the horizontal blue line is the smallest in-part dose, the right end of the blue line is the largest dose, and the green line represents the gelation threshold that is at the midpoint between the rightmost red dot and the leftmost line-end. 67
- 4.16 Re-binning pixels across multiple rotations enables simultaneous part formation. **A:** Binary target geometry. **B:** When a constant image set (gray squares) is projected across 10 rotations, the object starts forming (surpassing the gel threshold) after 9 rotations. When the intensity of pixels is re-binned such that the majority of the dose is delivered at the end of the build (orange circles), the object does not form until the final rotation. **C:** Rate of gelation for a constant image set versus the image-set with value redistributed as per Eqn. 4.6. **C, Top Row:** Dose reconstruction after the 8th, 9th and 10th rotation for redistributed image set. **C, Bottom Row:** Dose reconstruction after the 8th, 9th and 10th rotation for the constant image set. Image from [2]. 70

4.17 Experimental demonstration of redistributing light intensities. **A:** model 3D part. **B:** Calculated proportion of gelled voxels after each rotation when the intensity is either constant for each rotation (gray squares) or redistributed across 30 rotations such that the majority of the dose is delivered at the end of the build (orange circles). **C:** Snapshots from videos of parts forming without (constant) and with redistributing the light intensities at the time points denoted in (B) by pink stars. At t_1 no object is visible in either the constant nor the redistributed light intensities. At t_2 the edges of the cube are visible for the constant case, but no part is visible in the redistributed case. At t_3 the full part is visible in both cases. **D:** Extracted parts for the constant (left) and redistributed (right) light intensities. Scale bar: 2 mm. Experimental parts were printed using a custom VAM printer equipped with a 405 nm LED light engine (CEL5500, Digital Light Innovations), with a maximum intensity of $55\text{mW}/\text{cm}^2$ illuminating a resin vial fixtured to a rotation mount (HDR50, ThorLabs). [3] Bisphenol A glycerolate diacrylate (BPAGDA), poly(ethylene glycol) diacrylate (Mn 250) (PEGDA), and - Methyl-4'-(methylthio)-2-morpholinopropiophenone and Irgacure 907 photoinitiator were purchased from Sigma Aldrich and used without further purification. The photosensitive resin mixture is a 75:25 wt% mixture of BPAGDA and PEGDA, respectively, with 40 mM of Irgacure 907. [4]. 72

4.18 Dose reconstruction from m^{th} OSMO model, applied to inhibited resin. This process is used during model optimization, and to compute a final image set. **A:** m^{th} model. **B:** One of N (e.g. 360) forward projections. **C:** Image set is clipped not at zero, but at $g_{min} < 0$. **D:** Image set is shifted by $+|g_{min}|$, yielding non-negative intensities. **E:** Image set is back projected, forming a dose reconstruction. **F:** The dose is applied to an inhibited resin. The ratio of $-g_{min}$ to 1 is the same as the ratio of the inhibition period to the conversion period of the resin. In terms of conversion, when all units are normalized, this has the effect of subtracting a factor of $|g_{min}|$ from the resultant conversion distribution. For mathematical convenience when optimizing a model, adding $|g_{min}|$ from step (C) to (D), and subtracting $|g_{min}|$ at the final step, can be omitted. By the linearity of the Radon Transform, the addition and subtraction of $|g_{min}|$ cancel each other out (under normalization), yielding the same final conversion distribution. 75

4.19 The effects of including material inhibition when computing VAM image-sets. **A:** A target geometry. **C:** A conversion profile computed under the assumption of no material inhibition. **C:** The voxel distribution from the reconstruction shown in (B). A 4% process window size allows for some material and optical imperfections, without affecting the final shape of the gelled part. **D:** A conversion profile computed assuming a 10% inhibition period. **E:** The voxel distribution from the reconstruction shown in (D). The process window size is four times larger than what was achieved in (B)/(C) without inhibition, resulting in a reconstruction with four times more tolerance to material and optical errors. A high-contrast reconstruction such as this would be desirable for applications with large optical or material uncertainties, or for applications where the initial-gelled part must be driven to higher degrees of polymeric conversion before removal from remaining resin, and post-cure. Image from [2]. 76

- 4.20 **A:** Gaussian point spread function (PSF). **B:** Back projection with Gaussian-pixel-PSF projector. **C:** Optimized reconstruction with $D_l = 0.5$, $D_h = 0.6$, and 150 iterations. Optical parameters: wavelength = 405 nm, numerical aperture = 0.1, digital micromirror device pixel pitch = $10.8 \mu m$, and magnification = 0.36. **(D)** Thresholded reconstruction, threshold = 0.545. Here, VER was 0.074, IPDR was 0.512, and process window size was -15.5%. Image from [2]. 77
- 4.21 Details for the Gaussian beam reconstruction shown in Fig. 4.20. **Top:** Algorithm Convergence, measured in VER. **Bottom:** Histograms of dose for Gaussian beam reconstruction. Image from [2]. 78
- 4.22 **A:** Target and occluding geometries. **C:** Back projection with custom projector; occluded regions are shown as black regions adjacent to the occluding geometries. **C, D:** Optimized raw and thresholded reconstructions with $D_l = 0.35$, $D_h = 0.60$, threshold = 0.515, and 120 iterations. Image from [2]. 79
- 4.23 Details for the occlusions reconstruction shown in Fig. 4.20. **Top:** Algorithm Convergence, measured in VER. **Bottom:** Histograms of dose for the occlusions reconstruction. Image from [2]. 80
- 4.24 **A:** Dense target and occluding geometries. **B and C:** Optimized raw and thresholded reconstructions with $D_l = 0.73$, $D_h = 0.74$, threshold = 0.687. **D:** Algorithm convergence and histograms of dose in B. Here, VER was 0.0756, IPDR was 0.4341, and Process Window was -38.3%. Image from [2]. 81

- 4.25 **A:** A 2D dose reconstruction yielding a VER of zero, made by a set of projections over only a 130° range, with $D_l = 0.78$ and $D_h = 0.79$. **B:** A gelation threshold applied to the dose shown in (A) yields a perfectly gelled slice of a thinker geometry. **C:** A tomosynthetic printing geometry, with projection direction vertically angled by 40° from the standard printing geometry. **D:** A slice of a 3D thinker-geometry dose-reconstruction with VER near zero ($1.7E-4$), made with the tomosynthesis writing geometry shown in (C), with $D_l = 0.74$ and $D_h = 0.75$. The process window size was -0.70% , and the IPDR was 0.26. Note that (D) shows a 2D slice of a 3D reconstruction, not a 2D reconstruction. In this case, 360 angles, all at a 50° to the axis of rotation, were used. Image from [2]. 82
- 4.26 Algorithm Convergence and Histograms of Dose for reconstructions shown in Fig. 4.25. **A and B** Dose histograms and algorithm convergence for tomosynthetic reconstruction shown in Fig. 4.25D. Here, the algorithm converged to a VER of $1.7E-4$. **C and D** Dose histograms and algorithm convergence for limited angular range reconstruction shown in Fig. 4.25A. Here, the algorithm converged to a VER of zero. Image from [2]. 83
- 4.27 VAM print using OSMO algorithm, yielding a PW of 5%. The resin was 3:1 BPA-GDA:PEDGA250 with 80 mM of I907. The print time was 48 s with a maximum optical intensity of 233 mW/cm^2 . **A:** SEM image of the print taken by Archish Muralidharan. **B-C:** Photographs of the part. **D-F:** Corresponding views of the target STL file. 84
- 4.28 VAM print using OSMO algorithm, yielding a PW of 4.8%. The resin was 3:1 BPA-GDA:PEDGA250 with 80 mM of I907. The print time was 72 s with a maximum optical intensity of 233 mW/cm^2 . **A-B:** Photographs of the part. **C-D:** SEM image of the print taken by Archish Muralidharan (scale bar: 5 mm). **E-H:** Corresponding views of the target STL file. 85

- 4.29 VAM prints using OSMO algorithm. The box was printed into a 3:1 BPA-GDA:PEDGA250 resin with 80 mM of I907. The print time was 37 s with a maximum optical intensity of 426 mW/cm^2 . The octet strut was printed into a 3:1 BPA-GDA:PEDGA250 resin with 20 mM of I907. The print time was 184 s with a maximum optical intensity of 647 mW/cm^2 . **A-B**: SEM images of the prints taken by Archish Muralidharan (scale bars: 5 mm). **C-D**: Corresponding views of the target STL file. 86
- 4.30 Prints of a hollow box made up of mesh walls. **A**: Projected side view of the target. **B-C**: Print immersed in resin and after a 10 min ethanol wash, respectively - severe swelling and distortion occurred. **D**: A dose reconstruction profile of the top of the box. **E-F**: The same box target printed as in index structure into a 2-stage holographic photopolymer [5] prepared by Gabe Seymour. The printed boxes had 8 mm side lengths. 87
- 4.31 Grayscale tomographic reconstructions by non-negative projection sets. **A**: A target dose to reconstruct and its histogram. **B**: A model, M , formed by frequency-filtering the target. **C**: The reconstruction and histogram yielded from the model. **D-F**: Reconstructions and their histograms yielded by adding a constant to the model. The more the model is shifted, the higher the dose values that reconstruct the object in the necessarily reduced-contrast reconstructions. There is a trade-off between uniformity and contrast. E.g. the average spread of each dose-region in the reconstruction in (F) is smaller than those in (C-E). However, the reconstructions shown in (C-E) exhibit a larger dose contrast than the reconstruction shown in (F). Image from [2]. 89

4.32 Example of reconstructing a gradient index (GRIN) lens profile via the object space method presented in this section. **A:** A target GRIN profile of the form $sech(gr)$, where g is a constant and r a radial coordinate. **B:** The reconstruction resulting from object-space method described in this section, with no constant shift added. **C:** The difference between the target and the reconstruction. **D:** Cross sections through the centers of the target and the reconstruction. **E:** The central portion of (D). Even though the approximation is poor at the outer radius of the reconstruction, it is much better in the central 50% of the print. The central region could be used to print a good approximation to a GRIN lens. 90

4.33 Example of a reconstruction made via FBP, expect with the image set shifted to be non-negative. **A:** The reconstruction formed by a filtered image set that was shifted to have a minimum value of zero. In arbitrary units, the background dose is 18, and the in-part dose is 19.25 (7% higher than the background). **B:** A conversion profile that would be produced if the dose in (A) were applied to a resin with a sufficiently large inhibition period, followed by a linear conversion vs dose response. Such a conversion curve is sketched in (D). **C:** Sketch of a resin response without a radical inhibitor added. Oxygen adds a very small inhibition period, so applying the dose in (A) results in a very small conversion contrast. That is, if the intensity profile in (A) is applied such that the oxygen in the out-of-part regions is exactly depleted, the total dose applied will be small since the oxygen inhibition period is small. Since the value 7% larger than this small dose is small, the in-part regions are driven only a small way up the conversion curve. This is shown with the blue arrows. The black arrows show an alternative use for the dose. **D:** An inhibition period equal to a dose of 18 results in a conversion profile (B) with a background of approximately zero, and an in-part conversion corresponding to a dose of 1.25. Since the in-part dose is 7% greater than the background dose, further increasing the inhibition period would allow for higher in-part conversion while keeping the background conversion at zero. More inhibition would allow for an further increased in-part conversion. The blue arrows show an example conversion range, likely most applicable as endpoints to grayscale reconstruction's dynamic range. The black arrows shown an example with a slightly shorter exposure time. This might be used for a binary print in which we want to ensure no out-of-part gelation. 94

- 4.34 Example showing how discretely sampled FBP reconstructions vary with angular sampling density. This is relevant to VAM since spatial resolution is finite (projected beams are not infinitely small), as is angular resolution. In the continuous limit, a reconstruction would have two voxel values: zero and one (after normalization). Instead, we see broad distributions of voxel values for both in-part and out-of-part voxels. **A, B, C**: FBP reconstructions of a binary target, along with dose histograms using 0.5, 1, and 1.3 images per azimuthal degree, respectively. Further increasing angular resolution did little to improve the result. This is likely because spatial resolution - and not angular resolution - is the limiting factor to reconstruction quality. Unfortunately, while it is conceptually easy to increase angular resolution in VAM, spatial resolution is fundamentally optically limited. Any such deviations from theory could have some effect on value ranges of a shifted FBP image set. . . . 96
- 4.35 Example of the very low dose contrast that results from using shifted FBP on a particular target geometry: a resolution target. **A**: Image set histogram from FBP, producing **B**, a dose reconstruction. **C**: Histogram of the shifted FBP image set; note that non-negative value range. **D**: Shifted FBP dose reconstruction from (C), with dose histograms shown in **E**. Note the very low dose contrast. **F**: Sketch of a hypothetical inhibited resin. If we wanted an in-part conversion that would be produced by a dose value of 1 without the presence of inhibition, we would need the inhibition period to span 225 units of dose. That is, since the fractional difference between the in and out-of-part doses (E) is about 225:1, then so must be the ratio of the dose to clear inhibition and the dose to affect conversion. 97

- 4.36 Simulated reconstruction showing the impact of nonreciprocity, assuming a simple model for effective dose: $t \times I^{0.86}$. **A-B**: OSMO dose reconstruction with a process window of 1.6%. Suppose we assume that the resin responds linearly to an effective dose of $t \times I^{0.86}$, where t is time and I is optical intensity. Applying this scaling exponent to the project images, then back projecting, results in a reconstruction of effective dose, which will be proportional to resin response. Applying this to the image set that generated (A-B), we get **C-D**, the simulated effect of material nonreciprocity. A negative process window occurs, along with a large VER. 102
- 4.37 OSMO extrapolation strategy. **A**: Illustration of the value of a single voxel in an OSMO model. A single iteration is performed by an OSMO step. Extrapolating model values can speed OSMO convergence. However, we have not yet found an extrapolation speed that, for all targets, always speeds convergence while never slowing it. **B**: An initial OSMO model. **C**: A final OSMO model. 106
- 4.38 Increased projection direction concepts for VAM. **A**: The addition of projections parallel to the axis of rotation would significantly increase options for dose delivery. **B**: Tilting the axis of resin rotation during printing would allow access to a full range of tomosynthesis angles in a single print. This would likely require index matching, although compensation for image distortion is possible. This approach could also be applied to flat samples. 107
- 5.1 Potential applications of tomosynthetic printing geometry. **A**: 2D pattern stitching scheme. Note that if the glass slides are part of the final package, we should be careful to only include resin voxels in the optimization process. That is, we do not care what dose is delivered to glass, air, etc., so we should take care to not over-constrain the optimization problem by including surrounding out-of-part voxels that physically will lie outside of the resin. **B**: Example of printing into a well plate - something that conventional VAM would be unable to do. 110

- 5.2 Tomosynthesis PSF computed via the Projection Slice Theorem. **Left Column:** 2D slices of frequency space filled by sweeping a plane through the origin about all azimuthal angles. The plane angles are perpendicular to the real space projection directions, where $\theta = 0^\circ$ is standard tomography. **Middle Column:** The PSF of the projection geometry, computed by taking the Fourier transform of the filled Fourier space. **Right Column:** Enlarged view of the PSFs. (A) is not far from standard tomography at 10° , so we see the PSF (C) similar in shape to a cylindrical voxel. At a very steep tomosynthesis angle (G) of 80° , the PSF (I) is elongated vertically; we lose resolution in the vertical direction. 112
- 5.3 Example of a tomosynthetic OSMO reconstruction with a positive (0.5%) process window. D_l was 0.86, D_h was 0.9, just over 1k iterations were used, and the tomosynthesis angle was 20° . **A:** Slice of the target geometry (shown in F), made through the center of the shape and normal to the axis of rotation. **B:** The same slice, but through the 3D dose reconstruction. Note that the gaps between the pillars are high in dose - this drives the PW to be small. **C & D:** Slices through the center of the part along different coordinate axes. In both, the axis of rotation lies vertically along the center of the reconstruction slice. The tomosynthesis angle can be seen in the streaks of dose. **D:** Dose reconstruction histogram. **F:** The target geometry. 114
- 5.4 Example of a tomosynthetic OSMO reconstruction with a VER of 0.2%. D_l was 0.87, D_h was 0.9, 5k iterations were run, and the tomosynthesis angle was 35° . **A:** An octet-strut target geometry. **B:** Reconstruction dose histograms. $VER > 0$ implies histogram overlap. **C-E:** Slices through the 3D dose Reconstruction, parallel to the axis of rotation. (C) is a central slice. **F-H:** Slices perpendicular to the axis of rotation. (F) is a central slice. 115

- 5.5 Example of a tomosynthetic OS MO reconstruction with PW of 3%. D_l was 0.87, D_h was 0.9, 2k iterations were run, and the tomosynthesis angle was 35° . **A**: A knot geometry. **B**: Reconstruction dose histograms. **C-E**: Slices through the 3D dose Reconstruction, perpendicular to the axis of rotation. (C) is a central slice. **F-G**: Slices parallel to the axis of rotation. (F) is a central slice. 116
- 5.6 A tomosynthetic OS MO reconstruction with PW of 2.1%. D_l was 0.85, D_h was 0.9, 1.5k iterations were run, and the tomosynthesis angle was 25° . **A & B**: A dental aligner target geometry. **C**: Reconstruction dose histograms. **D**: Central slice through the 3D dose Reconstruction, parallel to the axis of rotation. (E) Central slice perpendicular to the rotation axis. **F**: Central slice, normal to the prior slices. 117
- 5.7 Experimental setup used to print with the 25° tomosynthesis image set discuss in the previous figure. **A**: Side view. The rotation axis was tilted by 25° , achieving the desired tomosynthesis angle with the writing light. **B**: View of vial with resin (3:1 BPA-GDA:PEGDA250, 20 mM I907) immersed in index-match fluid (the same monomer blend, but without photoinitiator). **C**: Image of the printed part immediately after print completion. The maximum print intensity was 230 mW/cm^2 , the print time was 109 s, and the part was about 1 cm long. 118
- 5.8 Photos of the part printed in Fig. 5.7 for a qualitative comparison to the target geometry. The perspectives are different (camera-part distance, and simulated distance to the STL file), making any but the most crude qualitative comparison difficult. However, the shape looks generally correct, we believe this to be the first experimental demonstration of tomosynthetic printing, and it proves that the method is, indeed, possible. 119

- 5.9 A few conceptual approaches to coupling light into flat samples. **A:** Immersing the flat sample in an index-match bath. **B:** Refracting light directly into the sample. Avoiding a back-reflection at the exiting air-glass interface is important for avoiding stray light in the sample. An optical absorber, such as black polydimethylsiloxane, laminated to the back surface, prevents this. **C:** A prism with index-match oil between the prism and the sample. 120
- 5.10 Experimental attempt at coupling light into a flat sample surface with a prism and index-match oil. **A:** Side view. The sample holder was a glass slide glued to a tube. In terms of the coupling interface, this is equivalent to a sandwich-style sample. **B:** View through the prism of the sample. Where oil does not fill the gap between sample and prism, a reflection is seen instead. The glass slide was not perfectly normal to the tube's central axis, causing the introduction of air into the oil film during sample rotation. 121
- 5.11 Illustration of tomosynthesis optics challenge, ignoring the problem of coupling light into a sample that we discussed in the previous section. **A:** To print the same sized part as with a standard tomography setup, tomosynthetic VAM requires an extended DOF. This requirement becomes even more noticeable at low tomosynthesis angles. Unfortunately, it is these lower angles that we would prefer to use, as they result in improved resolution (see Fig. 5.2). **B:** One approach could be to tilt the DMD. This will result in a tilted image whose degree of tilt goes as the imaging magnification. This, however, creates imaging aberrations. **C:** A diffractive optical element (DOE), such as a blazed transmission grating, could be used to angle the writing light, while keeping the plane of focus parallel to the DOE. This is a particularly attractive solution as it decoupling DOF from the length and width of the sample. 122

- 6.1 **A–B:** A conventional VAM print process. **C:** Resultant dental aligner print exhibiting typical striations. As detailed in Section 6.1, initial local gelation causes focusing, leading to a self-written-waveguide (SWW) effect, manifesting as striations (scale bars: 3 mm). **D–E:** The latent-image process proposed here. **D:** Patterning is stopped just before gelation occurs, leaving a 3D latent image of polymer conversion that is higher in the desired print region than in the surrounding resin. **E:** The latent image is developed across the gelation threshold via diffuse, uniform LED illumination, driving only the desired region to gelation. **F:** The resultant print exhibits dramatically reduced striations, with a smooth surface and improved refractive index homogeneity. Image from [6]. 127
- 6.2 Scanning electron microscopy (SEM) scans of typical striations in conventional VAM prints. **A:** A cone printed with the axis of rotation aligned with the central axis of the cone. A roughly $60 \mu m$ striation pitch was measured both near the tip of cone, where the focused beam size was measured to be $11 \mu m$, and near the base of the cone, where the beam size was $23 \mu m$. A single projector pixel was measured to have a focused size of 6 and $18 \mu m$ at the center and at the edge of the print, respectively. **B:** A slab print – the same as in Fig. 6.4A, B, F, G – exhibiting irregular striations. Scale bars: 1 mm. Shadowgram images of the cone print formation are shown in Fig. 6.3. See Section 6.2.3 for further print details. Image from [6]. 129
- 6.3 Shadowgram showing the development of the cone print from Fig. 6.2. The formation of striations is visible upon the initial gelation. Image from [6]. 130

- 6.4 Striation mitigation by the latent cure method proposed in this paper. A–D: Shadowgrams and E: SEM scan of conventional VAM prints of a slab, mesh, and tilted mesh showing striations on the order of feature size upon completion. F–J: The same geometries printed using the latent cure method. Striations are largely eliminated, both in surface roughness (J), and in refractive index uniformity shown via shadowgrams of prints still immersed in resin (F–J). A–B and F–G illustrate the improved shape fidelity yielded by the method. Scale bars: 1 mm for A–J, with 500 μm for zooms. The shadowgrams were captured with the basic collimated LED shadowgraph setup described in [5]. Image from [6]. 132
- 6.5 In-part and out-of-part dose histograms, in normalized units, of VAM reconstructions. The top image shows a process window [2] of 9.1%. Adding a latent flood exposure to the entire volume shrinks that to 8.2%, slightly reducing the error tolerance of the printing process (e.g. timing, illumination uniformity, etc.). When a VAM print is stopped just prior to the gelation of any voxels, the in-part histograms highest values are just under the gelation threshold. Thus, to push these in-part voxels to gelation, a dose equal to the width of the in-part histogram (the In Part Dose Range [2]), must be added. Adding this value of dose to all voxels, both in and out of part, reduces the contrast of the reconstruction. However, as long as the histograms are separated by a positive value, the gelation threshold can be chosen to lie between them, thus gelling all in-part voxels without inadvertently gelling any out-of-part voxels. The latent VAM technique would likely work well in combination with the image-set re-binning technique described in Section 4.1.7. A re-binned image set would require a shorter flood exposure period, and the images in the final rotation would better match a uniform exposure than would a standard image set. Both these factors would minimize the extent to which the process window size was reduced. Image from [6]. 133

- 6.6 Sketch of a resin response requiring a large process window to avoid gelation errors. Consider a resin that requires a degree of conversion beyond gelation in order to survive post processing. The additional dose required to attain this conversion requires a process window of at least that dose in order to avoid unwanted out-of-part gelation. Image from [6]. 134
- 6.7 **A:** Sketch of latent-cure illumination setup. An LED with diffuser is aimed at the corner of the VAM index match bath; it is also aimed upwards so that it is not directly primarily orthogonally to the rotation axis. **B:** Image of a VAM print in progress. As soon as the first hint of gelation appeared on a shadowgram, the writing illumination was turned off. **C:** Image of the same print during the latent-cure phase. Similar setups have been used on other VAM printers to good effect. The particulars of the flood-exposure setup seem to not be important; so far they have all worked. . . 136
- 6.8 Latent cure as a method of speeding the GP in a low power VAM print. **A–D:** Conventional VAM process, suffering from low print power and low viscosity resin. The partially-gelled regions sink before the part is complete. **D:** Severely distorted final part. **E–H:** VAM print with latent cure used to quickly develop the entire print through its GP before significant sinking occurs. Even though the patterning process and viscosity were unchanged from A–D, the rapid gelation via latent cure results in, H: a significantly improved print. The GP could be easily further shortened by increasing the power of the latent exposure. In this case, a GP of only a few seconds was difficult to control, so the latent-cure power was reduced so that gelation could be easily observed, and the latent exposure stopped before unwanted gelation. Image from [6]. 138
- A.1 Sample OSMO code implemented in Matlab 150

Chapter 1

Introduction

Multi-step additive manufacturing (AM) is being pursued for many diverse applications [7] like: printing optics [8, 9, 10], microfluidics [11, 12, 13], composites requiring overprinting [14, 15], regenerative medicine [16, 17, 18, 19, 20], etc. However, multi-step processes have some inherent challenges. From extrusion-based techniques to stereolithographic [21] printing, they struggle with three main challenges: **1.** Structures must be connected and supported, **2.** new printing-material must be transported to the part during the print, and **3.** printed parts suffer from layering effects [22]. **Challenge #1** has been addressed by a variety of approaches [23], but these present additional steps that slow down printing, and the final part-options are ultimately still limited. **Challenge #2** not only limits print speed, but limits the resins that can be used for printing; e.g. printing cannot be done into solids, and the mechanical properties of parts are ultimately limited by the viscosity of the printing-resins [24, 25]. **Challenge #3** is tolerable for some applications, but for others (e.g. printing optics, printing parts that must be uniform or otherwise patterned in modulus), it can cause inhomogeneities in material properties (e.g. refractive index, modulus, chemical functionality, etc.) that limit its applications, and it is a problem that the field of multi-step AM is still working to solve.

The relatively new field of Volumetric Additive Manufacturing (VAM) performs printing in a single lithographic step [26, 27, 28, 4, 5]. While some approaches rely on two-color chemistry to print layer-by-layer in a fixed volume of resin [29], the most general VAM method to print into a single-photon absorbing photosensitive material does so tomographically. In tomographic VAM printing

(Fig. 1.1), a series of 2D optical patterns are projected into a rotating volume of photosensitive material. For printing 3D shapes, a liquid resin is used, and over the course of tens of seconds to a few minutes, the accumulated optical dose distribution polymerizes the material, resulting in an arbitrary 3D structure, and the printed part is removed from the remaining liquid resin. VAM is free of many of the constraints inherent to multi-step AM. Most notably, non-contiguous parts become possible, layering effects are eliminated, and the lack of material movement mid-print enables high-viscosity or even solid resins, with no penalty to print-time, opening the doors to a host of new material properties [28].

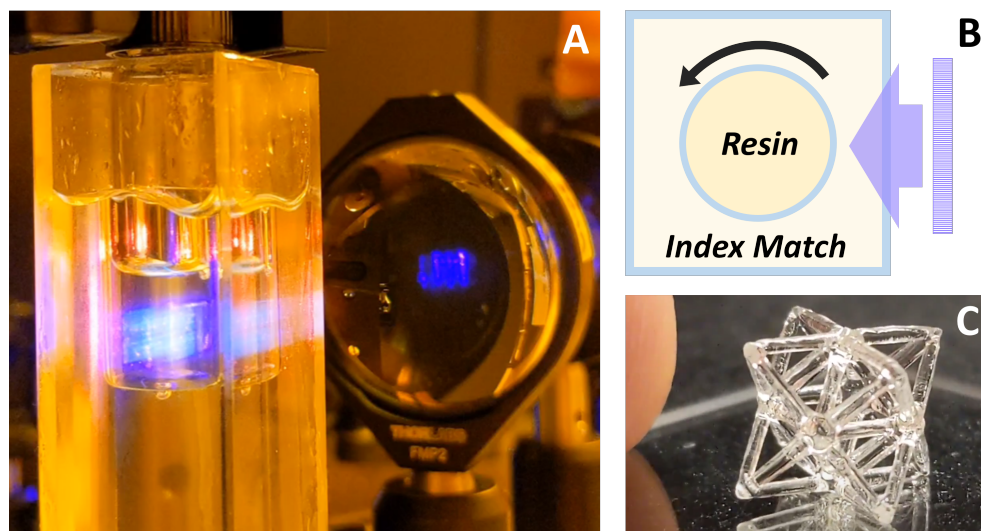


Figure 1.1: Tomographic VAM printing of a 3D geometry. **A:** Rotating cylindrical vial containing liquid resin. The surrounding rectangular cuvette holds a refractive index matching bath. 405 nm writing light is incident on the resin. **B:** Schematic of a 2D slice-region of the printing setup. **C:** After about one minute, a part solidifies due to accumulated optical energy. It is removed from the surrounding liquid resin, then it is washed and post-cured.

Compared to prior AM techniques such as stereolithography, VAM can print a broader range of geometries, it promises dramatic improvements to strength and other properties of printed parts, it offers improved uniformity of such properties, and it does so with orders of magnitude of improvement to print speed. It has also been demonstrated to be easily scalable from micron [5] to centimeter sized features with little impact on speed. This represents a significant flexibility advantage when compared with e.g. two-photon [9] beam-scanning methods which, although pro-

viding sub-micron resolution, yield print speeds on the order of years per cubic centimeter. VAM can pattern tens of megavoxels in sub-minute print times, with voxel size scalable to the desired application, making it an attractive technique across a relatively wide range of size scales.

Furthermore, VAM as a general lithographic patterning technique extends beyond the problem of printing 3D solid shapes. The ability to deliver optical dose within a fixed volume of material - even a solid - enables the tuning of 3D distributions of material properties. From writing refractive index distributions into holographic photopolymer (Fig. 4.30) for the creation of optical elements (Fig. 4.32), to the 3D patterning of chemical functionality for the manufacturing of artificial tissue [16, 17, 18], lithographic patterning enables a broad range of devices across many fields.

For all of these reasons, VAM holds promise as the future of additive manufacturing for a wide range of applications. However, prior to the work presented in this thesis, VAM suffered from **multiple limitations** that both restricted its scope of application and prevented its practical use as a manufacturing technique.

The **first** of these limitations was the lack of tomographic calculation techniques to provide sufficient optical dose control for a variety of VAM applications. VAM tomographic computation methods that insufficiently control dose delivery can cause errors in the shape of a final gelled part, either for purely mathematical reasons or because they require perfect optics and materials. Real VAM is enabled by a combination of imperfect optics, material science, and computation. While significant progress had been reported on optics and materials [26, 27, 29], there was a lack of analysis in the literature about the inverse math problem of computing the images that, when projected into resin, integrate together to print a part [30]. An image-computation algorithm yielding perfectly-shaped gelation for complex parts, even with imperfect optics and materials, could expand the range of structures printable by tomographic VAM. Increased control over dose could also allow for control over resin conversion beyond just the state of gelation, opening the doors to printing arbitrary functionally graded materials with VAM. This work presents an approach to VAM image computation that is simple, effective, and that significantly improves control over optical dose in VAM. We demonstrate the flexibility of the approach by extending it to include

material response models, optical models including beam shape and occlusions for overprinting, as well as to the problem of controlling grayscale dose throughout a print.

A **second** limitation in the field of VAM was the inability to print into other than cylindrical sample packages. Such a limitation not only restricts potential applications, but in a cylindrical printing geometry, optical and materials constraints limit print size in two of the three Cartesian dimensions. Here, we present a tomosynthetic VAM printing geometry that allows for printing into flat surfaces. This lends itself to applications such as the manufacturing of flat microfluidic chips, printing into well plates, or any other printing application with a flat sample or limited angular range of beam access. Furthermore, it removes one dimensional constraint to print size, allowing for the patterning of arbitrarily large sheets. We discuss theoretical limitations, the tomosynthetic VAM image computation problem and a solution, and we present experimental proof of concept prints.

A **third** limitation in VAM was the presence of striations in all printed parts shown throughout the literature. Although VAM is fundamentally free of inherent layering effects, it suffered from large striations – similar in appearance to layering, and on the order of print feature size – impacting the homogeneity and shape-accuracy of printed parts, and presenting a significant cosmetic defect. We examine the source of these striations, and discuss the hypothesis that the material nonlinearity upon which VAM relies also drives striations via a self-writing waveguide (SWW) effect. We demonstrate a simple and effective method of dramatically reducing striations via the addition of a uniform optical exposure applied to the entire volume of resin at the end of a print. We show that this step has the additional advantage of drastically shortening the period from initial gelation to print completion. This mitigates the problem of partial-print sinking in low power or low viscosity prints, thus further expanding resin options and increasing the efficacy of low-cost VAM printers.

Before examining these challenges and their solutions, we contextualize them by discussing the problem of VAM in general. There are three fundamental pillars of VAM: optics, materials, and image computation. These areas are highly coupled, but it is helpful to approach them as separate

problems, with each defining constraints on the others. Before an extensive discussion on image computation, we examine the fundamental constraints on VAM optics and materials, and how they inform design approaches of VAM printers and resins. The chapter on resin design discusses bounds on absorption, photoinitiator concentration and molar absorptivity, methods of measuring molar absorptivity for photoinitiator selection, resin preparation techniques, and print post processing. But first we begin with optics. We discuss the fundamental constraints and trade-offs in VAM optics, how this informs the design of VAM printers, and details of the design and alignment of the primary VAM printer used in this work.

Chapter 2

Printer Design

The first tomographic VAM printer was comprised of a rotation stage to spin a vial of resin, an index-matching bath to avoid the optical power of the curved vial, and an overhead projector as a patterning source [31]. Printer design has since significantly improved. Alternatives to index-matching have been demonstrated [30], print monitoring optical systems have been developed [32], and, most importantly, control over the spatial distribution and power of lithographic projection beams has been vastly improved [33]. This chapter discusses printer design: fundamental optical constraints, design requirements, approaches to experimentally achieving them, print monitoring methods, and alignment techniques. We begin by discussing the problem of focusing light across a vial, and how depth of focuses requirements limit the number of resolvable spots in a tomographic VAM printer.

2.1 Optical Design Requirements and Fundamental Constraints

2.1.1 Depth of Focus

In a tomographic reconstruction, a given voxel - a discrete element of the reconstruction - is created by the superposition of beams. In the most general case, every azimuthal projection angle contributes to the addition of dose at the voxel. For non-diverging beams that are uniform in intensity across their square cross section (originating from square pixels), the shape of the resultant voxel will be approximately cylindrical. As the number of azimuthal projection angles goes to infinity, the voxel shape converges to a cylinder. This is illustrated in Fig. 2.1.A.




Figure 2.1: Illustration of the dependence of voxel size on beam width and divergence. For this illustration, the intensity distribution across the width of the beams is ignored, and only beam width is considered. For a more detailed analysis, see the chapter on image computation. **A:** Schematic of a 3D tomographic back projection showing the intersection of non diverging beams projected from square DMD pixels. The intersection of these beams forms a voxel whose shape is approximately cylindrical. As the number of projection angles approaches infinity, the voxel shape approaches an exact cylinder. **B:** Two dimensional view of a single slice-region of the reconstruction showing the intersection of non diverging beams. The circular (in 2D) voxel that is formed is outlined in black. **C:** Diverging Gaussian beams whose focus is set to the center of the vial. Voxels formed near the center of the vial are small in diameter - equal to the waist width of the focused beam. **D:** The voxels near the edge, however, are larger and not necessary circular, as they are formed by the intersection of defocused beams.

Voxel size is bounded by the width and the shape of the projection beams. In the case of uniform, non-diverging beams, voxels throughout the reconstruction will have a diameter equal to the beam width (Fig. 2.1B). If, however, the beams are not uniform in width, and instead diverge, the upper bound on voxel size will vary throughout the reconstruction volume. The smallest voxels will occur where the narrowest parts of the beams intersect, with the largest voxels occurring where the widest parts of the beams intersect. The case in which the beam focus is set half way through the the vial is illustrated in Fig. 2.1.C,D. Note that factors such as beam intensity profile, the relative intensity of different overlapping projection beams, the particular set of projection angles contributing to a given voxel, etc. all affect the size and shape of voxels. However the feature size and resolution of a reconstruction are bounded by the maximum width of the writing beams. Thus, in the most general case, global print precision is optimized when the maximum width of writing beams is minimized. This give us our first optical system constraint: in general, **the depth of focus (DOF) of projected beams should equal the diameter of the build volume.**

2.1.2 Number of Resolvable Spots

Suppose we wish to set the DOF of a Gaussian beam to a particular value - to the diameter of the build volume. Depth of focus has been defined in various ways: by Rayleigh range, by double the distance from the focus of the beam to where it expands by a factor of $\sqrt{2}$, and by other arbitrary metrics. For the purposes of this analysis, let us take the DOF of a Gaussian beam to be double the distance, in the direction of beam propagation, from the beam waist to where the beam expands in width by a factor of $\sqrt{2}$, where beam widths are defined by a $1/e^2$ drop in intensity in the radial direction (orthogonal to the propagation direction). By this definition, $\text{DOF} = (8\lambda/\pi)(f\#)^2$ [34], where λ is the optical wavelength, and $f\#$ is the focal length divided by the illuminated diameter of the focusing lens. Thus, we can calculate the $f\#$ required for the DOF to equal the diameter of our print vial.

The $f\#$ also determines the smallest spot size of the beams - the width of focused beam waist. This spot size can be computed as $2\omega_0 = (4\lambda/\pi)f\#$ [34]. Since the depth through the vial is equal

to the DOF, and the depth through the vial equals its width (since the vial is round), we have that the width of the vial is equal to the $DOF = (8\lambda/\pi)(f\#)^2$. The number of resolvable spots N_{SPOTS} that can fit across the width of the vial is then given by $DOF/(2\omega_0)$. Or, equivalently, $N_{SPOTS} = \sqrt{\pi/2}\sqrt{DOF/\lambda} = \sqrt{\pi/2}\sqrt{(\text{vial diameter})/\lambda}$. The minimum spot size of a beam also goes as the root of DOF: $\text{Spot Size} = \sqrt{DOF} 2\lambda/\pi$. This is illustrated in Fig. 2.1.

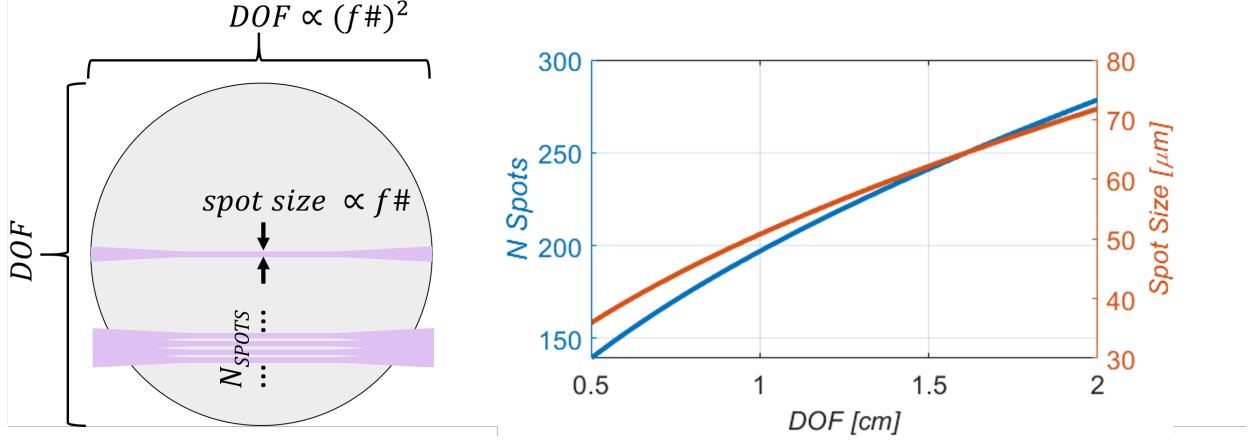


Figure 2.2: Number of resolvable spots at a wavelength of 405 nm that can fit across the width of a vial when the DOF is equal to the vial diameter. **Left:** Since the vial is round, when the DOF is equal to the depth through the diameter of the vial, it is also equal the width of the vial. A DOF equal to the vial diameter yields an $f\#$. That $f\#$ also defines minimum beam spot size $2\omega_0$. Thus, the number of resolvable spots N_{SPOTS} across the width of a vial is equal to $DOF/(2\omega_0) = \sqrt{\pi/2}\sqrt{DOF/\lambda}$. **Right:** For 405nm light, we can plot this relationship (blue plot) to see that N_{SPOTS} goes as \sqrt{DOF} or, equivalently, as $\sqrt{\text{vial diameter}}$. The minimum spot size of a beam follows the same trend, as is plotted in orange.

Understanding the limit on the number of resolvable spots is useful for designing the optical and resin systems, and for predicting print resolution. For example, we know that the limit on N_{SPOTS} for a 1 cm diameter vial with 405 nm light is approximately 200. This means that increasing the resolution of a pattern generator beyond about 200 pixels will likely minimally improve print fidelity. It would also mean that if print times and resin viscosity allow for about half of a micron (1 cm / 200) of diffusion distance, the mobility of chemical species during printing is likely not the limiting factor of print fidelity. While such limits are not absolute, and clever improvement beyond them may be possible, they do provide rough constraints that inform the design of the

overall printing system.

2.1.3 Power and Etendue

It is desirable for VAM projections to be high in optical power to minimize print times. As we will discuss later, increased optical power relaxes resin design constraints. The bound on $f\#$ described in the previous subsection has implications for which light sources can efficiently provide high DOF writing beams, and which cannot. This is most simply understood via the concept of etendue - a geometrical measure of how much some given light spreads over area and angle [34].

Consider an LED light source. The integrated product of its area and the solid angle from which light is emitted is a relatively large value - it has a large etendue. This value is large relative to the same product for the high DOF light of VAM writing beams. However, etendue is a conserved quantity [34]. Thus, in the absence of blocking some of the light, the VAM print plane will have the same etendue as the LED, regardless of the imaging path that was taken from LED source to patterned projection beams. This means that without throwing away light, an LED will produce writing beams with a relatively high divergence - with a low DOF. A typical method of imaging with high DOF is to cut out high angle light with an iris, however this cuts out power. The exact amount of power lost is dependent upon the angular power distribution of the light source, but making the approximation that power is uniformly distributed over angle, the power lost goes as the square of the ratio of the numerical apertures of the writing beams and of the LED - a very large loss. This yields the formulation of another design constraint: the light source for our VAM printer should be of high power and low etendue.

2.2 Optical Design Approach

In the previous section we introduced some of the basic design parameters for VAM printers, and discussed a resultant theoretical resolution limitation. In particular, we concluded that the VAM writing beams should have a large DOF, equal to the diameter of the print volume, and that the printer light source should be high in power and low in etendue. In this section, we

discuss a design approach to achieve desirable patterned illumination, some resultant challenges and solutions, and the design of a particular printer.

2.2.1 Laser and Fiber Light Source

The obvious choice for light source high in power and low in etendue is a laser. However, the output of a laser is not ideal for illuminating a pattern generator for multiple reasons. A laser intensity profile tends to be brightest in the center, and to fall off radially. While a grayscale pattern generator could perhaps compensate for this field non-uniformity, using the grayscale dynamic range of a pattern generator for this purpose is certainly not ideal. Lasers also have noticeable coherence effects. Interference patterns, speckle, etc. all detract from the ideal scenario of uniform projected pixel values. Fortunately, the use of an optical fiber alleviates many of these problems for VAM applications.

The idea of using a fiber for illumination homogenization, first applied to VAM by Loterie Et al. [33], is to use a laser to fill the modes of a fiber. The output end of the fiber then serves as a relatively low etendue, well-defined optical source, with each point on the core emitting light over the NA of the fiber. Filling the modes of a circular-core fiber is possible by e.g. launching a Gaussian beam into one end at the correct angle. However, more convenient is to use a square core fiber, not only for the ease of mode filling and mixing, but also for the convenient square shaped light exiting the fiber that nicely matches the square format of many pattern generators.

The etendue of light exiting a mode-filled fiber is appropriate for VAM. For example, consider a $70 \times 70 \mu m$, 0.22 NA fiber. If that $70 \mu m$ dimension is eventually magnified to an order centimeter print plane, the NA of the light will be proportionally reduced, resulting in a order centimeter DOF without requiring a loss of power.

Coupling light into a fiber without damage can be challenging due to the high intensities involved. For example, a 3 W laser focused onto a $70 \times 70 \mu m$ fiber core produces an intensity of about $60 \text{ GW}/\text{cm}^2$. If a piece of dust lands on the fiber core - perhaps guided by the optical beam - it will likely burn, damaging the core in the process. One approach to minimize the chance of

burning, other than simply avoiding dust in the lab, is the spread the laser focus over the entire area of the fiber core, minimizing intensity. Another problem arises, however, if significant power is spread beyond the core; the polymer cladding that often surrounds square core fibers is susceptible to burning (e.g. Fig. 2.3). It is much more convenient to purchase a laser that comes fiber coupled; the alignment is already done, and typically the input end of the fiber is sealed from dust.



Figure 2.3: A burnt fiber core caused by focusing a 1 W, 405 nm laser onto a $70 \times 70 \mu\text{m}$ fiber with a polymer cladding. The laser ran for many hours without burning the fiber, suggesting dust as a cause of failure.

A fiber provides an opportunity to mitigate the effects of laser speckle. Speckle is visible in the output of a fiber. However, moving the fiber changes exactly how the modes are filled, resulting in movement of the speckle pattern. Sources of movement such as a piezoelectric actuator, an electric toothbrush, or even a tapping finger, move the speckle pattern quickly enough that it is washed out on the order 25 millisecond timescale of a single VAM image projection. That is, the average intensity over a VAM projection time is effectively uniform. Note that this does not negate all coherence effects, but it nicely solves the problem of speckle for laser illumination in VAM.

2.2.2 Pattern Generation

To turn the uniform illumination output by the fiber into tomographically-computed writing beams, we need some sort of pattern generator. For a rotationally symmetric target print, the projected pattern would be invariant with angle, and the required static pattern could be made with

something like a chrome mask. However, for all other general printing, the pattern changes with each new projection angle, requiring variable pattern generation. Here, we discuss the parameters that make a variable pattern generator desirable for VAM: resolution, grayscale bit depth, refresh speed, and damage threshold.

Resolution requirements for VAM patterns are modest. To avoid being the limiting factor to print fidelity, the linear pixel count across the diameter of a VAM print need only be a few hundred (see Section 2.1.2). The resolution along the rotation axis of a print is not fundamentally limited and is theoretically without bound. However, for printing parts with an approximately cubic aspect ratio, we only need a pattern generator with a resolution on the order of hundreds of pixels squared. E.g., printing 2×2 cm parts with 405 nm light would only require 289×289 patterns (Fig. 2.1).

Grayscale resolution requirements are more difficult to meet. As we will discuss later, we have observed simulated degradation in the quality of tomographic reconstructions when limited to 8-bits of projection grayscale resolution, with 12-bits appearing to be sufficient for no degradation. However, this degradation may manifest only as a reduction in system error tolerances, and may not necessarily impact printed part fidelity. The minimum bit depth that achieves maximum print fidelity is unclear. However, for the purposes of this section, we will state that a good VAM pattern generator should have at least 8 bits of grayscale depth, with 12 bits preferable.

Pattern refresh speed is important for VAM printing. As we will discuss later, having a high number of projection angles, on the order of at least one per degree, is necessary for maximizing tomographic reconstruction quality. The refresh period times the number of projections yields the time per rotation of the print stage. If this rotation time makes up a large percentage of the total print time, a few problems can arise.

The first problem with slow rotation is that the initial gelation during a print can scatter subsequent writing beams. If the rotation rate is fast, this scattering is averaged over a wide range of angles as the rest of the part gels. If instead, however, the rotation rate is very slow, most of this scattering will occur from one side of the print, causing scattering effects to be concentrated on one side of the part. Note that this effect becomes irrelevant if applying the latent cure technique

discussed in a later chapter.

The second problem with slow rotation has to do with tomographic reconstruction and the effects of not completing a full set of projection images. In the extreme case where the print gels before a full rotation is completed, only a limited angular range is available to contribute to a tomographic reconstruction. This limits reconstruction quality (see the chapter on image computation for further details). If the number of rotations completed is a low, non-integer number greater than one, then the dose contribution from the missing portion of the final rotation is significant. It would be possible to compensate for this by redistributing intensity among the projection images, but this would require knowing exactly how long a print would take, which would require a high level of process control. A much more convenient approach is to print with a high spin rate such that many rotations occur during a print. Any final incomplete rotation would then make up only a small percentage of the integrated optical dose, and its effects could be ignored. Thus, a high spin rate is preferred.

Since it is desirable to have a high spin rate and a high angular resolution, it is desirable for a pattern generator have a high refresh rate, with "high" being relative to print time. For example, an 80 Hz refresh rate with 360 projection images yields just over 13 rotations over a 1 minute print. The additional partial rotation results in some angular range of projections contributing about 8% more dose to the reconstruction than it should. Whether this is acceptable depends on the error tolerances of the tomographic reconstruction and of the optical and material systems. Although it would be ideal to exactly computationally compensate for incomplete rotations, this is not always practical, and thus an ideal VAM pattern generator maximizes image refresh rates. Note that it is possible to instead stop the projector at an integer number of rotations before any gelation, and then to apply the latent cure technique discussed in the chapter on latent VAM.

Lastly, it is desirable for a pattern generator to have a high optical damage threshold. An example of an instrument with a low damage threshold is a liquid crystal spatial light modulator (SLM). Unfortunately, SLMs are not rated to the multi W/cm^2 intensities of high power VAM printers, and would likely be an unreliable option, at least for high energy wavelengths like 405 nm.

It is possible that an SLM would be an appropriate choice for longer wavelength VAM systems.

Given all of the above requirements, and a 405 nm wavelength, a digital light modulator (DMD) emerges as the natural choice for image generation. DMDs have plenty of resolution; e.g. the Vialux V-9001 used for prints shown in this work has 2560×1600 pixels. Since DMDs work in reflection, their damage threshold is high - easily supporting a few W/cm^2 at a 405 nm wavelength. DMDs support relatively high image refresh rates and grayscale bit depths. Since grayscale is tuned via an on/off duty cycle, there is a trade-off between refresh rate and bit depth. For example, our V-9001 can support 8-bit images at 303 Hz, but that rate drops to 20 Hz for 12-bit images. While not perfect, a DMD spans most of the required design space, and appears to be the best option for VAM pattern generation.

2.2.3 DMD Overview and Diffraction Efficiency

A DMD presents a theoretical efficiency limitation stemming from the fact that an array of micro-mirrors behaves like a blazed grating and produces diffraction orders. Since the different orders leave the DMD at different angles, imaging them all to a VAM sample would create multiple, offset images. Instead, we block all but the 0^{th} diffraction order, leaving us with a single image as desired. Any power going into the other orders is thus lost. In this subsection we discuss the basic functionality of a DMD, and how to maximize 0^{th} order diffraction efficiency.

DMD pixels are arranged on a Cartesian grid. The most common pixel type is a small square mirror, oriented in line with the grid of pixels. That is, if the pixels run up and down; left and right; then the edges of each pixel are horizontal or vertical. A second type of DMD has each pixel rotated by 45° , forming a Cartesian grid of diamond-shaped micro mirrors. We will discuss the prior square-mirror type, as it is the layout used in the Vialux V-9001 DMD used in this work.

Regardless of the layout, a DMD micromirror pixel works by rapidly tilting back and forth, changing the angle at which light is reflected. A micromirror pivots about a diagonal line running between its opposite corners (Fig. 2.4). A DMD pixel has two tilt states: *on*, where the micromirror is tilted by $+12^\circ$ in one direction, and *off* where the micromirror is tilted in the opposite direction




Figure 2.4: DMD micromirror layout, tilt directions, and light input and output directions. **A:** A Cartesian grid of micro-mirrors. The line about which each mirror pivots is sketched with a black dashed line. Each mirror can tilt by $\pm 12^\circ$, where 0° would correspond to the plane of the mirror being parallel to the grid of mirrors. An input and output beam form a plane that is orthogonal to the mirror pivot planes. An end-on view of this plane is shown in red. That is, beams input to the DMD lie in a 45° plane that is orthogonal to the plane formed by the locus of micro mirrors. **B:** A sketch of light incident on and exiting a DMD *on* state pixel. The dashed blue lines are Cartesian axes, with one axis orthogonal to the plane of the DMD. The input and output beams, shown in purple, lie in a 45° plane to the DMD. The angle between an incident beam and the normal to the surface of the DMD is variable. The power efficiency of reflected beam depends on this angle, θ_{inc} . **C:** Micro mirrors act as a blazed grating. The path-length difference between beams going to different mirrors causes a relative phase delay, and thus interference. When choosing an input angle θ_{inc} , we seek to minimize destructive interference, thus maximizing power in the 0^{th} diffraction order.

by -12° . Light incident on the DMD is reflected towards a desired image plane by the *on* pixels, and is sent to a beam dump by the *off* pixels.

A grayscale DMD pixel value is set by choosing an *on/off* duty cycle. E.g. with a 50/50 cycle, a pixel reflects light towards the target image 50% of the time, creating a 50% gray level when averaged over time. Higher grayscale bit depths require that a larger number of mirror cycles be integrated over time. For example, an 8-bit image requires 256 mirror states. The darkest gray pixel would then be created with a single *on* state, averaged with 255 *off* states. A higher bit depth image would require a higher number of mirror states. Since the mirror switching rate is fixed, higher bit depths thus require a longer integration period for a single image. This creates a trade-off between image refresh rate and grayscale bit depth.

Understanding where DMD diffraction efficiency is maximized is relatively straightforward.

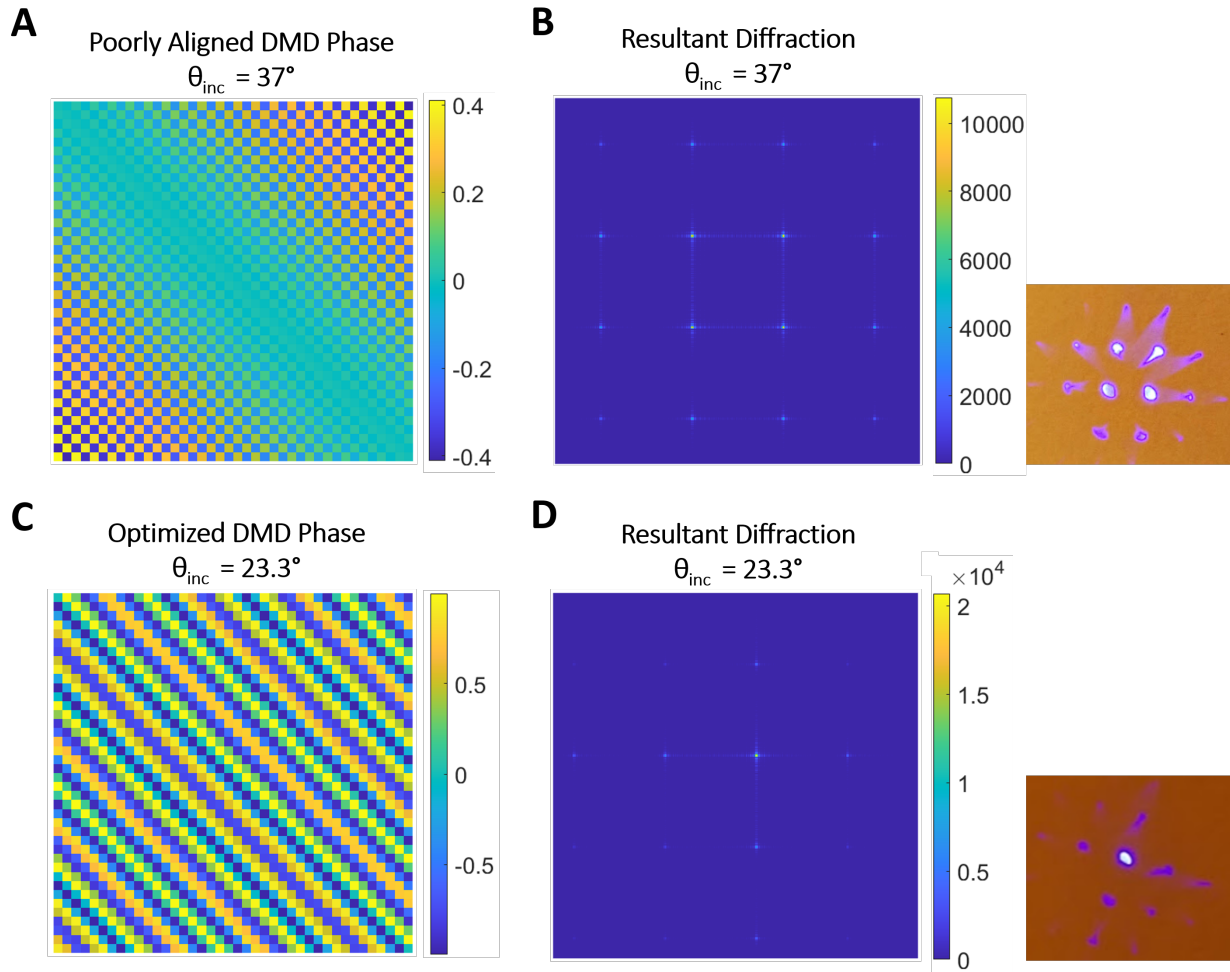


Figure 2.5: Simulation of DMD diffraction. A linear phase ramp is applied to the DMD by light incident at θ_{inc} . Scanning through a range of input angles lets us see the effects on diffraction efficiency. Here, we assumed 405 nm light, and a micromirror pitch matching our Vialux V-9001 DMD ($7.6 \mu m$). **A**: A poorly chosen input angle of $\theta_{inc} = 37^\circ$. **B**: The resultant diffraction pattern. The 0^{th} order contains only about a quarter of the total power - a poor efficiency. Instead of the majority of the power going into a single order, most of it is split into four orders. Such an example from the lab is shown. **C**: The DMD phase that results from an optimum choice of $\theta_{inc} = 23.3^\circ$. **D**: The resultant diffraction pattern. Here, the power into a single order is maximized, with about 80% of the power going into the 0^{th} order. This is also observed in the lab, as shown.

We input light such that it lies parallel to the plane shown in red in Fig. 2.4. The variable over which we optimize is then the angle of incidence to the DMD-normal within that plane, θ_{inc} . Diffraction from a DMD in such a configuration is easily modeled. We simply apply a linear phase ramp to the DMD mirrors that is a function of wavelength, of θ_{inc} , and of mirror pitch. Taking

the Fourier transform of this phase function then yields the far-field diffraction from the DMD (or equivalently, the diffraction distribution that results from taking the Fourier transform of the DMD with a lens). See Fig. 2.5. The percent power going into the 0^{th} order can then be measured, and the input angle θ_{inc} optimized.

For our particular Vialux V-9001 DMD pixel pitch ($7.6 \mu m$) and illumination wavelength (405 nm), simulation shows that 0^{th} order diffraction efficiency is maximized when $\theta_{inc} = 23.2^\circ$. In the lab, however, the highest 0^{th} order diffraction efficiency I have measured is about 75%. This may be due to imperfect alignment, imperfect mirror reflectivity, or other sources of loss at the DMD. Regardless, even a basic simulation of DMD diffraction aids in DMD alignment, as will be discussed later, and provides a useful starting point when constructing a system.

2.2.4 Relay Optics

Imaging requirements in a VAM printer are simple. First, light exiting the optical fiber must be magnified, and imaged onto the DMD. Second, the DMD must be magnified or demagnified to match print size, and imaged onto the sample plane. Imaging the output face of the square-core fiber onto the DMD is easily done with a two-lens system. One convenient approach is to use a microscope objective as the first lens, and to mount it on a 3-axis translation stage for ease of alignment with the fiber core.

Imaging from the DMD to the sample is best done with a 4F, two-lens system. Not only does this provide a Fourier plane at which to filter out all but the 0^{th} DMD diffraction order, but it also provides telecentric imaging to the sample. Thus, even if the print diameter exceeds the writing beam depth of focus, the projected patterns suffer only from defocus, and not from a change in magnification. That is, in a telecentric configuration, the beams do not change width as they go out of focus, and the beams remain parallel.

2.2.5 System Overview

Here we present a particular VAM system that realizes the design parameters outlined in this chapter. The light source is a 3 W, 405 nm multi-mode laser from Civil Laser that came coupled to a 0.22 NA, 100 μm round-core fiber. This was replaced with a 0.22 NA, 100 \times 100 μm square-core fiber. A fiber-coupled laser is far preferable to a laser that must be free-space coupled into a fiber, as it greatly reduces the chances of burning the fiber. The DMD used was a Vialux V-9001. After all system losses: fiber coupling, losses at mirrors and lenses, and DMD diffraction, the final print-plane power is about 0.5 W/cm^2 . The printer layout is shown in Fig. 2.

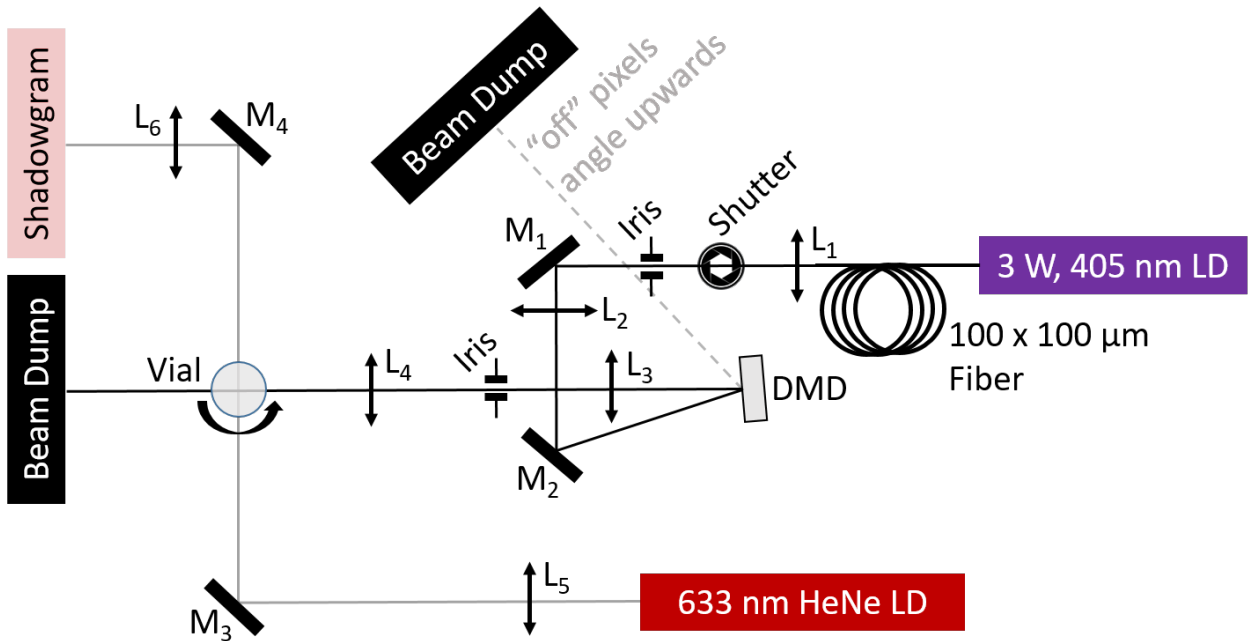


Figure 2.6: The Design of a VAM printer used in this work. The laser is a multi-mode 3 W, 405 nm laser from Civil Laser. The fiber is a 100 \times 100 μm square-core fiber made by Ceramtec Optics. Note that a 0.22 NA, 70 \times 70 μm may also be used for larger sample depth of focus, although that comes at the cost of reduced power. The mirrors, M_1 and M_2 are AR-coated for high reflectivity at 405 nm. L_1 is a 20 \times microscope objective with an effective focal length of 9 mm. The focal lengths of lenses $L_2 - L_4$ are: 90 mm, 51mm, and 100mm, respectively. A helium-neon laser is used to monitor print formation; its 633 nm wavelength does not affect the resin. Lens L_5 expands the beam such that its width spans the vial diameter. The lens L_6 then focuses that expanding beam into a screen. These shadowgram focal lengths are not critical, as they are used for qualitative print monitoring.

A 633 nm helium-neon laser is projected through the sample, and onto a screen for the purposes of monitoring the formation of gelation during printing via a shadowgram. While the 405 nm wavelength initiates polymerization, the 633 nm wavelength has no appreciable effect on the resin. Any gelation in the vial deflects the red light, and the corresponding intensity on the screen is reduced. This makes obvious small changes in refractive index in the resin - even changes so small that they are not observable by eye upon direct inspection. See Fig. 2.7 for an example of a shadowgram image. While much more sophistication is possible for in situ print metrology (see [5]), a simple shadowgram is sufficient for basic print timing.

2.2.6 Alignment Methods

Here we briefly discuss a few points about VAM system alignment: aligning a DMD for good diffraction efficiency, aligning the central axis of the print vial to the rotation axis, and aligning the axis of rotation to the center of the projected patterns

There are a few approaches to DMD alignment. The simplest is to rotate the DMD by 45° so that the lines about which the micro-mirrors pivot is vertical. Then light is input at the same height of the DMD, and leaves at that same height. In this configuration all light stays in a horizontal plane that is parallel to the optical table. Then, to align the incident light at the correct angle to the DMD (e.g., at 23.3°, as shown in Fig. 2.5), one needs only to rotate the DMD about a vertical axis. For each new candidate position, a lens can be used to take the Fourier transform of the DMD, all but the 0th diffraction order can be blocked, and the remaining power can be measured. This process can be repeated over a range of angles until 0th diffraction power is maximized. The downside of this approach is that the DMD pixels form a rotated Cartesian grid which does not conveniently match the unrotated Cartesian grid of tomographically computed images.

An alternate DMD configuration is the one shown in Fig. 2.7. Here, the DMD is nominally still horizontal, although not perfectly so. In this configuration, light is sent up at an angle towards the DMD, and the DMD is tilted forward and to the side such that the writing light emitted from the DMD is parallel to the optical table. This tilt rotates the DMD slightly, requiring images loaded




Figure 2.7: VAM printer details. **A:** Light path from fiber, to DMD, then towards print plane. Note that the light angles up towards the DMD. **B:** Print vial with resin, immersed in a bath of index-matched liquid. A rotation stage spins the vial within the index-match bath. Here, the index-match liquid was resin without photoinitiator, making it non-reactive to the 405 nm light. **C:** A photograph of a part just after print completion. The part is still immersed in resin, but its increased density causes increased refractive index which, for this resin, allows it to be seen by eye. **D:** A photograph of the shadowgram screen just after print completion. The shadowgram makes refractive index changes noticeable, which aids in determining the point in time at which the part is complete.

onto the DMD to be rotated by a few degrees so that the writing patterns aligned with the table. Alignment in this case is a bit arduous. First, a target illumination angle is known, as discussed in the section above on DMD diffraction. Some quick trigonometry then tells us the distance a beam should travel upwards towards the DMD, the distance it should travel laterally, and the distance it should travel parallel to the final writing beam direction. Calipers can then be used as a rough

guide to establish initial beam alignment. As in the previous scenario, a lens and iris are used to extract the resultant 0^{th} order diffraction, and its power is measured. To test a new angle, the mirrors before the DMD are used to scan the beam to come towards the DMD at a slightly larger or smaller angle. The tilt of the DMD must then be adjusted so that the exiting light is brought to be again parallel to the optical table. Diffraction efficiency can then be measured. This process is tedious, with many small adjustments needed to test each new candidate alignment position. Although we have achieved a diffraction efficiency of about 75% in the past, DMD alignment in this configuration is difficult enough that settling for lower efficiency, around 60%, is often the most practical choice.

Aligning the center of the resin vial to the axis of rotation of the stage is a simple process, but one that must be performed with each new print due to the low repeatability of the clamp that holds the vial. One easy method of checking for misalignment is to set the vial spinning via the stage, and to look at the shadowgram. If the axis of symmetry of the vial and the axis of rotation of the stage are not co-aligned, the shadowgram image will drift left and right as the stage spins. X-Y micrometers between the vial clamp and the stage can then be used to laterally adjust the position of the vial until this shadowgram lateral movement is eliminated.

A simple method of aligning the axis of rotation to the center of the projection patterns is as follows. First, a cross pattern is projected from the DMD, and its position marked far from the print plane (e.g., it is marked with an iris at the end of the optical table, at a far wall, etc.). Then a filled vial (filled with e.g. water) is put into the vial clamp under the rotation stage. Then the central axis of the vial is aligned to the rotation axis as per the above description. If the vial's center is not aligned to the pattern's center, the cylindrical shape of the filled vial will deflect the projected cross pattern so that it no longer lands at the first marked location. The entire stage is then moved laterally until the cross pattern lands where it originally did without the presence of the vial. As long as the alignment of the DMD projection is not changed, and the lateral position of the stage is not changed, this step will not need to be repeated between prints.

Thus, the center of the projection pattern is laterally aligned to the central axis of the vial

which is aligned with the axis of rotation. After achieving this state, alignment for a new print is simple. Upon installing a new sample into the vial clamp, the vial must be aligned with the axis of rotation. Then the index-match bath is installed, and a print can be performed. It is important that every step of lateral position alignment be executed with care, as small errors can significantly impact tomographic reconstruction quality. Fig. 2.8 presents the simulated effects of laterally misaligning the axis of rotation and the center of the projection pattern. In this example, filtered back projection (FBP) was used to tomographically reconstruct a target image using 720 projection angles. The results indicate that misalignments of only a few microns can noticeably degrade the sharpness of tomographic reconstructions for order 1 cm sized prints.




Figure 2.8: Simulated effects of misaligning the center of projections with the axis of rotation. **A:** A 2D tomographic reconstruction (using FBP with 720 projection angles) of a target geometry. **B:** The reconstruction that results from laterally misaligning the axis of rotation by 0.4% the width of the projections. This corresponds to a lateral misalignment of only $4 \mu\text{m}$ for a 1 cm print, but noticeably degrades reconstruction sharpness. **C:** The reconstruction that results from laterally misaligning the axis of rotation by 0.6% the width of the projections. Further reconstruction degradation is apparent.

2.3 Summary

In this chapter, we discussed the fundamentals of VAM printer design. We first examined fundamental optical constraints: how beam depth of focus affects the size of written voxels throughout a print, how the requirement that beam depth of focus span a print implies a limit on the number

of resolvable spots in the print, and how the depth of focus requirement can be reformulated as an optical etendue requirement to help inform our choice of light source. We presented an illumination method consistent with VAM requirements: a laser and a square-core optical fiber. We discussed methods of pattern generation, and why a DMD is a good choice for high-power VAM. We examined the unwanted diffraction that occurs with a DMD, and we discussed a theoretical and experimental approach to maximizing its efficiency. We presented the complete design of a VAM system, and discussed the importance of telecentric imaging for VAM projections. Finally, we described methods of system alignment, and we demonstrated the importance of precise pattern alignment via tomographic simulations.

Chapter 3

Resin Design

This chapter discusses basic parameters for designing VAM resins. At its most basic, a VAM resin needs to be photosensitive and sufficiently transparent for patterning beams to reach into the desired print regions. The possible changes upon optical exposure are many: e.g. physical properties such as the state of gelation, density, and diffusivity; chemical properties such as molecular weight and reactivity; and optical properties such as refractive index, absorptivity, scatter, etc. To illustrate some of the fundamentals of resin design in VAM, we will focus here on resins for binary VAM for which the property of interest is the final state of gelation. That is, resins for the traditional 3D printing problem - that of creating a solid object in a desired shape.

One of the most ubiquitous photosensitive resins, and the type that we will focus our discussions on here, are acrylate-functionalized monomers mixed with a radical photoinitiator. In our resins, the monomer is essentially transparent to the writing wavelength, and is unaffected by writing beams. When a photoinitiator molecule absorbs a photon, photolysis occurs, and the molecule is split into two pieces, with each piece carrying a radical. Each radical can initiate a propagation reaction. In the case of chain-growth polymerization - the kind we want for VAM so that prints occur quickly - monomers polymerize to form a growing oligomer that will then react with more monomer, thus growing a chain until the reaction is terminated by a radical or by another growing chain. Initially, as the local degree of polymerization in a resin increases before gelation, the index of refraction slightly increases due to an increased density and a slight change in the polarizability of acrylate bonds. This change in index is small enough that it appears to have no noticeable perturb-

ing effects on VAM writing beams. However, as conversion continues increasing, a phase-change threshold is reached, and gelation occurs. With gelation comes sudden shrinkage and a noticeable jump in refractive index. It is at this point that VAM prints are typically stopped. The gelled structure is removed from the surrounding liquid resin, the part is washed, and then the part is post-cured - often with a uniform optical exposure. The post-cure serves to polymerize unreacted monomer within the gelled structure, thus increasing conversion and corresponding properties such as mechanical modulus [35].

Such a resin system has a few important design parameters that we will discuss in this chapter. The total optical absorption through the depth of a print is fundamentally important to tomography, and is primarily determined by the concentration of the radical photoinitiator and by its molar absorptivity at the writing wavelength. Resin viscosity is another important property, and affects rates of diffusion during printing, as well the rate at which gelled regions sink in the resin before print completion. We begin by discussing bounds on optical absorption, and how a desired absorption informs other resin design choices.

3.1 Absorption

The decay of light intensity through the depth of resin is well-described by the Beer Lambert law: $I(z) = I_0 e^{-\mu z}$, where I_0 is the initial intensity applied to the resin, z is a spatial coordinate through the resin in the direction of light propagation, and μ is an attenuation coefficient. That is, with a single absorbing species, light intensity decays exponentially. The total absorption of light through a depth of resin can be written as $A = \epsilon C \ell$, where C is the concentration of the absorbing species ($mols/L$), ϵ is the *molar absorptivity* of the absorbing species ($L/(mol\ cm)$), and ℓ is the path length through the resin (cm). As we will discuss in the chapter on image computation, tomographic reconstructions for VAM become more constrained and lower in quality (as defined by the *process window* metric we will describe later) as absorption increases. However, if we were to tune absorption towards zero, the rate of photon absorption would approach zero, which for finite light intensity would drive print times to infinity. Long print times are problematic for a variety of

reasons. Thus, it is desirable to operate at the upper limit of absorption, within the optically thin regime. We have found absorption on the order of 10% to be about right, enabling high-quality tomographic reconstructions while allowing for sub-minute print times on our printer. The question then becomes: for a fixed vial diameter, ℓ , how should we choose ϵ and C to yield an absorption of about 10%?

3.2 Photoinitiator Concentration and Molar Absorptivity

The concentration C of photoinitiator molecules should not be chosen to be so low that a print locally depletes a significant fraction of the photoinitiator. If this were to happen, the rate at which applied optical intensity creates radicals would significantly decrease during printing. This would cause a few problems. Firstly, it would violate the assumption of the image generation algorithm that sensitivity to applied dose is a constant. Secondly, significant local photoinitiator depletion would complicate calibrations needed for VAM process control, making it harder to predict the dose required to reach a desired degree of conversion. For these reasons, it is desirable to include a high enough concentration of photoinitiator that only a small fraction is cleaved during printing. As a rough guideline, let us state for now that we wish for only 10% of initiator to be locally depleted during printing.

Understanding how many initiator molecules are depleted during an optical exposure is straightforward under a simplifying assumption: that any photon absorbed cleaves an initiator molecule. This may not be correct, as other types of absorption occur, with the most significant being the absorption of photons by the initiator photoproducts - by the already cleaved molecules that will not create more radicals upon photon absorption. However, this is a good approximation since initiator byproducts are typically not significantly more absorbing than uncleaved initiator, and their concentration relative to initiator is small by the constraint of minimal initiator depletion. Consider an experiment where some constant intensity of monochromatic light is incident onto a thin layer of resin. Gelatin occurs after some amount of time, t_{gel} . The total energy per area applied during the exposure is the intensity I_0 multiplied by t_{gel} . Suppose the intensity exiting the

resin is $(1 - A)I_0$, where A is a number between zero and one. Each photon carries an energy $\hbar\omega$, where \hbar is the reduced Planck constant, and ω is the angular frequency of the light. Thus, the total number of photons absorbed in the volume, per incident area, is $A(I_0 t_{gel})/(\hbar\omega)$, which, under our assumption, is also the number of initiator molecules cleaved. The initiator concentration can then be chosen such that this number per volume is only about 10% of the initiator concentration, thus providing an approximate lower-limit to initiator concentration. Note that these calculations can be extended to polychromatic light by including ϵ as a function of wavelength, and integrating the equation with respect to wavelength. Given these rough constraints on C and on A , the equation $A = \epsilon C \ell$, with ℓ chosen to be 1 cm, tells us that ϵ should be chosen to be on the order of $1 L/(mol\ cm)$. This value is approximately four orders of magnitude smaller than the molar absorptivity of typical photoinitiators at their design-wavelength peak [36], thus requiring operation in the low-absorption tails of initiator spectra, away from their design wavelengths.

3.3 Viscosity

If the viscosity of resin is too low, a few problems can occur. The first is that a partially-gelled part can sink significantly before print completion, distorting the final part. One method to mitigate this is to use the latent cure technique discussed in a later chapter. Otherwise, however, viscosity must be chosen to be high enough that the settling velocity of gelled regions in the liquid resin is slow relative to print times. See [33] for an analysis of this.

The second problem that can arise from low resin viscosity is that molecular species such as photoinitiator, oligomers, or other potential components such as radical inhibitors could diffuse a significant distance relative to printed feature size during the course of a print. We have not measured diffusion times, and thus cannot place a lower bound on viscosity (given a print time) due to such effects [37]. However, they do appear to occur. As an extreme example, we have attempted prints under conditions where gelation was expected to occur after 8 minutes, but gelatin never occurred even after running the print for 25 min. This was likely due to the out-diffusion of oligomers and radicals, and the in-diffusion of replacement monomer, thus preventing gelation.

Instead of performing experiments to place a lower bound on viscosity, the approach we have taken is to simply use as high of a viscosity as is practical to work with. Further experimental work is needed to understand diffusion effects in a given resin, and what that implies about viscosity requirements.

The resin with which I have printed the most is a 3:1 mixture of Bisphenol A glycerolate (1 glycerol/phenol) diacrylate (BPA-GDA) and Polyethylene Glycol Diacrylate with average molecular weight of 250 (PEGDA250). BPA-GDA is much more viscous than PEGDA250, so viscosity can be tuned by changing the ratio of these monomers. Darren Liu has measured BPA-GDA:PEGDA250 ratios of; 1:1, 1.5:1, 2:1, and 3:1 to have viscosities of; 0.62, 1.1, 4.3, and 13.7 *Pa s*, respectively. Increasing viscosity beyond 13.7 *Pa s* is certainly possible, however it would make resin preparation (mixing, pouring, etc.) quite cumbersome. Our experience has been that a 3:1 mixture is sufficiently viscous for sinking effects to be negligible over order minute prints for printed feature sizes on the order of 50 μm .

Lastly, viscosity can have an interesting effect on reaction rates within a resin. In particular, high viscosity can cause the rates of reactions to be limited by mobility, and not by their chemical reactivity. This can be especially important if multiple competing reactive species are present - e.g.: cleaved photoinitiator, oxygen, monomer, or a radical inhibitor such as the radical inhibitor (2,2,6,6-Tetramethylpiperidin-1-yl)oxyl (TEMPO). In this case, reaction rates may differ from what their chemical reactivities would indicate, as they become limited by their ability to diffuse towards one another [38].

3.4 Measuring Molar Absorptivity

We wish to choose a photoinitiator with a molar absorptivity on the order of 1 $L/(mol\ cm)$ at our optical wavelength as discussed above. When designing our VAM printer, we considered wavelengths at which inexpensive and high power lasers were available. Ultimately, given the choices of lasers and of photoinitiators, we chose to use 405 nm light and Irgacure 907 (I907). However, this selection process required measuring initiator absorptivities because of a lack of

published data at their absorption tails - away from their design wavelengths where they were all too absorptive by about four orders of magnitude [36]. In this section we discuss the problem of measuring molar absorptivity, we present rough data on initiators that we discounted from our search, and we present data on the photoinitiator that we ultimately selected, I907.

Molar absorptivity ϵ can be extracted from the Beer Lambert Law: $\epsilon = \frac{A}{C\ell}$, where C is the concentration of the absorbing species (*mols/L*), A is the total optical absorption, and ℓ is the path length through the resin (*cm*). Experimentally, a spectrometer can be used to measure solutions in a cuvette with a fixed path length ℓ . For a given solution, the spectrometer will measure A as a function of optical wavelength. One simple approach to extracting ϵ is to repeat such a measurement for a variety of initiator concentrations, including a solution with zero initiator. ϵ can then be computed as the slope between A and $C\ell$. This method is relatively accurate and yields a low noise floor because it inherently rejects fixed losses such as Fresnel reflections or other absorbing species, even when losses dominate individual measurements.

Fig. 3.1 illustrates such results from four photoinitiators. This data was taken in search of an initiator with $\epsilon \approx 1$ at the candidate laser wavelengths of 405 nm and 473 nm. Each curve was generated by measuring $A(\lambda)$ with a Thermo Scientific Evolution 300 UV-VIS spectrometer for five initiator concentrations in either ethanol or acetone. In each case, the highest concentration was chosen to be near the solubility limit of in initiator in the solvent - typically just over 300 mM. The stock solution was then diluted by 10% for each subsequent data point. Data for both unexposed and photobleached photoinitiator is plotted. Unexposed photoinitiator was measured first. The samples were then flood cured for about 5 minutes with a few Watt 405 nm LED to produce photobleached samples. Regions of low absorptivity where the plotted line is jagged indicate a poor signal to noise ratio. While this data is rough, it was sufficient to eliminate these initiators from our search.

This experimental approach to measuring molar absorptivity has some limitations. First, to avoid polymerizing resin in the cuvette, the initiator was dissolved into a solvent - acetone or ethanol. This could affect the resultant measurements. However, solvatochromic shifts are typically

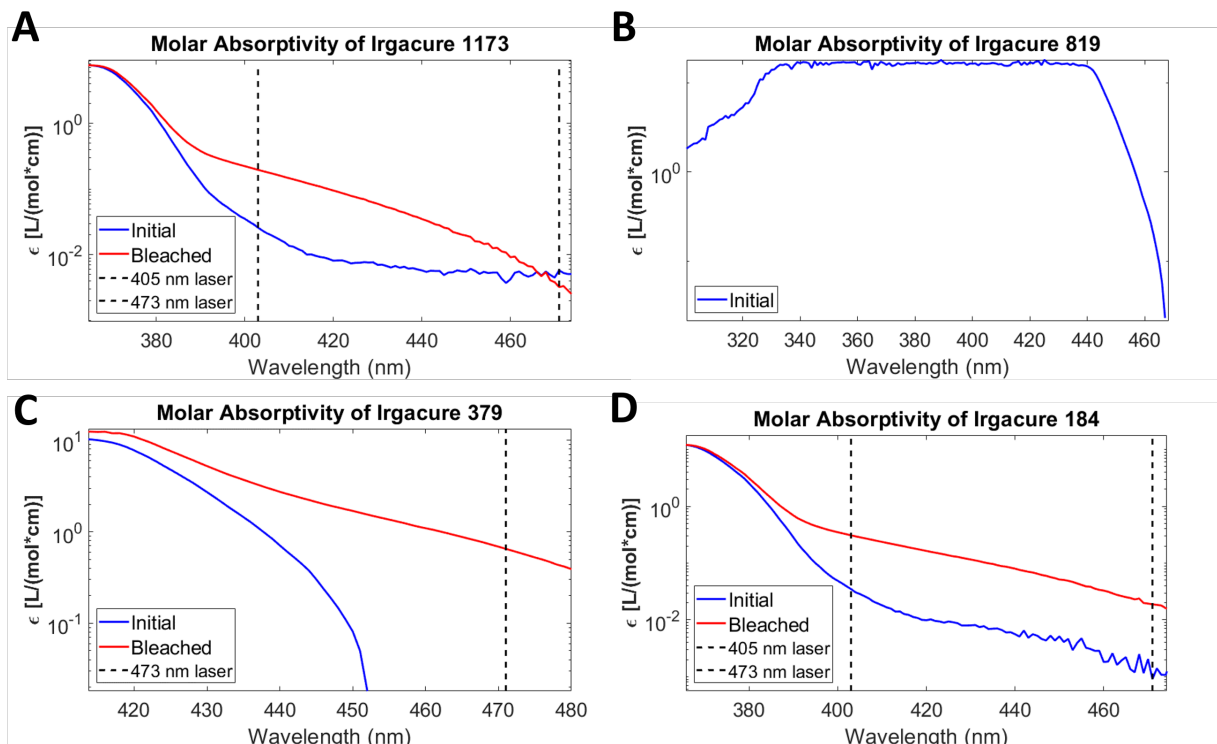


Figure 3.1: Initial rough measurements of molar absorptivity for four photoinitiators that were eventually discounted from our search for an initiator with $\epsilon \approx 1$ at the candidate laser wavelengths of 405 nm and 473 nm. **A:** Irgacure 1173 was not absorptive enough. Note that the isosbestic point - where initial and bleached absorptivities are equal - is near the 473 nm laser line. **B:** Irgacure 819 was appropriately sensitive only near 450 nm - not a candidate wavelength. No bleached data was collected for this sample. Note that the flat region of the plot between about 330 and 440 nm is saturated data. **C:** Irgacure 379 was too sensitive at 405 nm, and not sensitive enough at 473 nm. **D:** Irgacure 184 was too low in sensitivity at visible wavelengths.

small, and are likely negligible in our screening for initiators with $\epsilon \approx 1$. The second difficulty is that spectrometers have a limited signal to noise ratio. Since data at very low absorptivity values is desired, and solubility limits of initiator in solvent limit compensation for this via increased concentration, we must operate in a low absorbance and thus high noise measurement regimen. Finally, the limited dynamic range of a spectrometer means that a given solution will yield a saturated signal at some wavelengths, while not providing a large enough signal at others. This is important if trying to collect data over many wavelengths, which is particularly useful before a light source has been chosen.

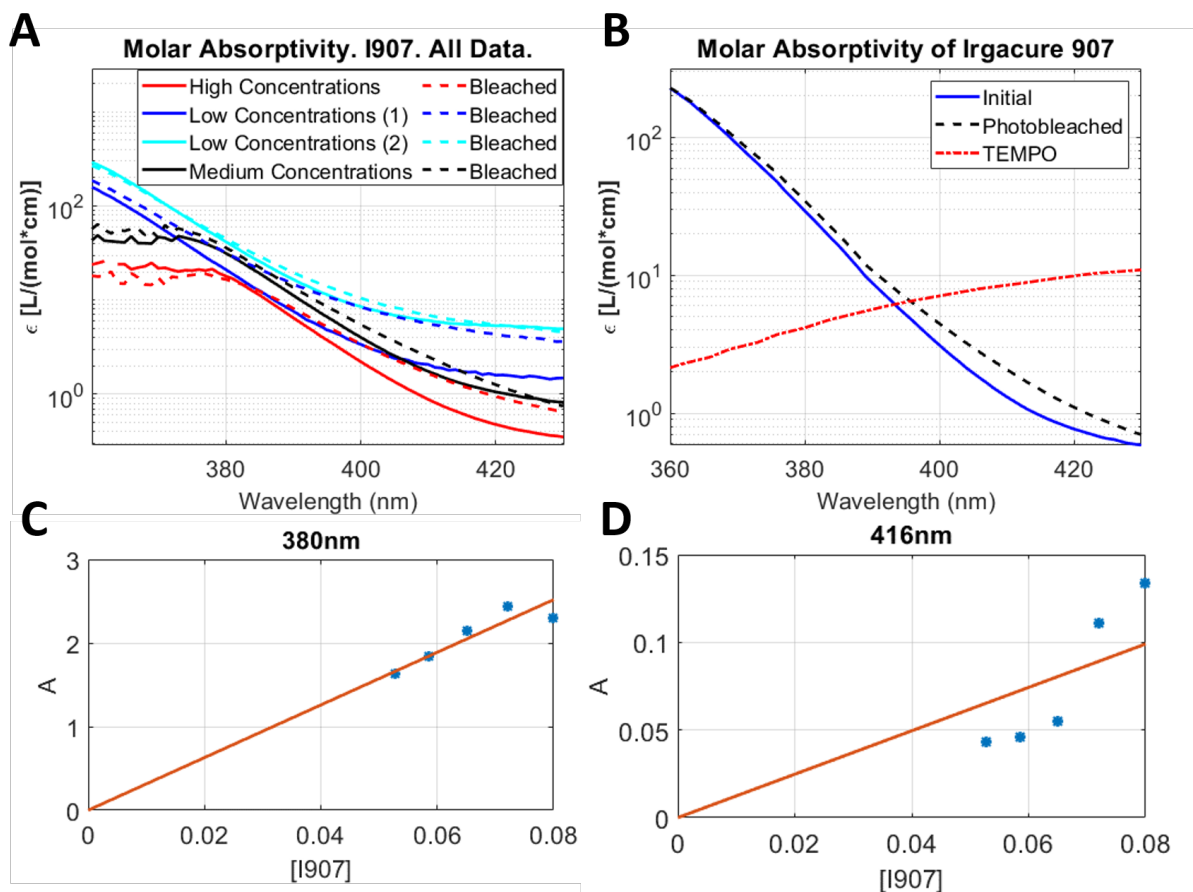


Figure 3.2: Molar absorptivity data for Irgacure 907 collected with assistance from Darren Liu. **A:** Molar absorptivity computed from different concentration ranges of I907 in ethanol. The solid and dashed lines show data for unbleached and photobleached initiator, respectively. **B:** The data from (A) was narrowed to exclude wavelengths where measurements were saturated or very noisy, then the data was averaged to plot molar absorptivity of both photobleached and unbleached I907. The molar absorptivity of the radical photoinitiator, TEMPO, is also plotted in red (data from [1]). **C and D:** raw data showing absorbance and I907 concentration for two optical wavelengths. For these, the line fit was constrained to pass through the origin. Both data sets were from the "High Concentrations" set of dilutions. All points should ideally perfectly fit a line. Especially for longer wavelengths, spectrometer noise significantly degrades the data accuracy. However, we found the final level of accuracy of this data (B) to be sufficient for an initial photoinitiator selection.

Fig. 3.2 illustrates some of these challenges. Under perfect conditions, the solid curves (unbleached) in Fig. 3.2A should overlap perfectly, as should the dashed curves (photobleached). The solid red curve, "High Concentrations", had its initial concentration chosen to be near the solubility limit in ethanol. A series of four dilutions, each 10% lower in concentration, was made

from the initial concentration. This data, computed from high concentration solutions, is more precise at low values of molar absorptivity than the rest of the data which was computed from lower concentration solutions. The solid black curve, "Medium Concentrations", was computed from a similar series of concentrations, with its highest initial concentration lower than the concentrations from "High Concentrations". Similarly, "Low Concentrations" (1) and (2) were computed from even lower ranges of concentration. This data performs better at shorter wavelengths, as it avoids spectrometer saturation. The dashed plots were computed with the same sample preparation, except that the samples were first photobleached. 3.2B shows an average of this data, with (C) and (D) showing example raw data with line-fit quality varying with wavelength.

The resin with which we have made the most prints is the aforementioned 3:1 (by weight) mixture of BPA-GDA:PEGA250 with a concentration of I907 in the 20-80 mM range. This yields an absorbance through a 1 cm print in the 4-18% range. The best compromise between print speed and tomographic reconstruction quality is dependent upon the particular structure to be printed. Thus, there is no best concentration value, and simulations must be run for each new structure to determine what degree of intensity attenuation will still allow for sufficient reconstruction quality. In the absence of such careful optimization, we have found a 40 mM concentration to work well, yielding a 5-10% process window (discussed in the image computation chapter), with 50-90 second print times.

3.5 Resin Preparation & Reuse

Here we briefly share some experimental methods of resin preparation: measuring components, mixing resin, and avoiding bubbles.

BPA-GDA, PEGDA250, and I907 in powdered form were mixed with the following steps. First, the desired final volume of resin is approximated. Then the densities of BPA-GDA and PEGDA250 are used to calculate the correct masses of these monomers to fill the desired volume in a 3:1 ratio. In order to facilitate pouring the viscous BPA-GDA, the stock bottle was heated in a 60°C water bath. Next, the BPA-GDA is poured into a jar in approximately the correct volume,

and the added mass is weighed. This mass is used to compute the correct mass of PEGDA250 to add - a much easier task because the low viscosity of PEGDA250 allows it to be pipetted. The total volume of the BPA-GDA and the PEGDA250 is then used to calculate the number of moles of I907 that must be added to achieve the desired I907 concentration. Then the molecular weight of I907 is used to calculate the mass of I907 that must be added. Fully mixing I907 into highly viscous material is difficult. Instead, it is easier to first mix the PEGDA250 and the I907. Then, this mixture can be added to the BPA-GDA. A speed mixer is an ideal tool for mixing everything together. However, a magnetic stir bar on a hot plate at 60°C works well enough if left mixing overnight.

When transferring resin into a print vial, or when transferring index-match fluid into the cuvette that surrounds the print vial, it is common to introduce bubbles. This is bad, of course, and the index mismatch from resin to air will cause optical scattering. Waiting for a matter of hours will allow the bubbles to leave, but this can be shortened to a few minutes by putting the material in a vacuum oven. A temperature of 40-60°C has been sufficient for all visible bubbles to leave under only the modest vacuum generated over the course of a few minutes by a roughing pump.

3.6 Post Processing

Upon completion of printing, the part must be removed, washed of any excess resin coating it, and post-cured. For binary structures, post-print processing falls into three categories: removal of the part from the bulk of the surrounding ungelled resin, washing, and post curing.

Removing the part after print completion requires some care as the parts are often quite soft and fragile. This is because the VAM print does not generally drive the part to a high degree of polymer conversion. Instead, the print is stopped just after gelation has been reached, and before any unwanted gelation occurs in regions surrounding the part. E.g., pulling the part out with tweezers tends to be too violent, resulting in breaking or severe deformation. A better strategy is to pour the entire print vial into a new print vial (perhaps to be used later for subsequent prints).

Just as the part is about to fall out of the first vial, it is instead poured into a vial of solvent, and the washing process begins.

The solvent that we have used the most for BPA-GDA and PEGDA resins is ethanol. Acetone and methanol are too aggressive and tend to destroy the part. Tetraethylene glycol dimethyl ether is another option. It is even less aggressive than ethanol, but also does a worse job at removing unreacted resin from the surface of the part. If using ethanol, it is important to gently tumble the part in a bath of ethanol. This mechanically expedites the removal of extra resin, important for minimizing the degree to which ethanol swells into the part - too much swelling and the part deforms. Starting with warm (60°C) ethanol can make the washing process easier. Once the part is cleaned of extra resin, it can be extracted from the solvent in preparation for post curing.

Post curing is done to increase the polymer conversion within the cleaned part, thus increasing properties such as mechanical modulus. For resin formulations including a thermally activated initiator, post curing may involve heating the part in an oven. For our resins, post curing is simply a matter of applying a high optical dose to the cleaned part. This is easily done with a 405 nm LED, or using the VAM laser light e.g. before it reaches the DMD.

Some parts tend to have a tacky surface, especially after a long solvent wash. It has been hypothesized that is due to the removal of photoinitiator during the wash - solvent swelling in, and initiator exiting the surface of the printed part. One remedy is to wash the part in solvent that has a high concentration of photoinitiator dissolved into it. Thus, lost initiator is replaced with new initiator, enabling conversion at the surface of the part during post cure. Early implementations of this approach qualitatively appear to be effective, but further work is needed to develop a reliable post-processing protocol.

3.7 Photo Product Fluorescence

One challenge in VAM is measuring and understanding the quantity and location of radicals produced by the cleavage of photoinitiator. The photoinitiator fragments that remain after cleavage have a property that may be useful for tackling this problem: the fragments fluoresce.

Fig. 3.3 shows an example of this effect. Here, a sample of around 600 mM I907 in PEGDA250 was placed between glass slides. A Nikon AR1 microscope was used to gel regions of the resin in 2D patterns with its 405 nm stimulation laser. The increased refractive index of these gelled structures were visible (Fig. 3.3, bottom). Upon stimulation with 561 nm light, these regions fluoresced, emitting light in the 575-625 nm wavelength range. The greater the initial 405 nm writing optical dose, the brighter the emitted fluorescence. This effect is likely due to the cleaved photoinitiator, as the photoinitiator in the regions not exposed with 405 nm light did not fluoresce. This could be used as a method of measuring the location and quantity of cleaved photoinitiator without the addition of fluorescent tags. For a 3D volume this could be done by a scanning focused stimulation source, or the entire volume could be stimulated, and tomographic measurement performed at the emission wavelength.




Figure 3.3: Microscope images of I907 in PEGDA250 placed between glass slides, then gelled into 2D shapes with a 405 nm laser. **Top:** 575-625 nm emissions via a 561 nm stimulation source. Local stimulation intensity was approximately proportional to the 405 nm writing dose. **Bottom:** The same microscope, a Nikon AR1, was used in differential interference contrast mode to examine the refractive index change in the gelled structures. Thanks to Joseph Dragavon for his assistance with collecting this data.

3.8 Conclusions

Many resin formulations are possible for VAM, and this work is certainly not comprehensive. However, many of the problems that we encountered are applicable to the design of any VAM resin. We discussed an upper bound on total resin absorption, driven by tomographic reconstruction constraints. We examined how the low bound on resin absorption is driven by desired print times and printer optical intensity. We then discussed the factors that determine this target absorption range: optical wavelength, the molar absorptivity of a given photoinitiator, and initiator concentration. We discussed how the desire to avoid significant local depletion of initiator concentration yields a lower bound on photoinitiator. Taken together, these factors suggested the need for a photoinitiator operating at the tail of its absorption spectrum, with respect to optical wavelength. The lack of data detailing the absorption tails of initiators required experiments to measure molar absorptivity for a variety of initiators. We examined the challenges involved in such measurements, and presented molar absorptivity spectra for a few photoinitiators. These data, together with the readily available high power laser sources, informed a choice of optical wavelength and photoinitiator: Irgacure 907 with 405 nm writing light. We discussed the re-use of resin, a basic post processing method, and a fluorescent property of photoinitiator products that could aid in future print metrology efforts.

Chapter 4

Image Computation

A fundamental challenge in VAM is choosing the images to project into a volume of resin to produce a desired print. Unlike most other lithographic 3D printing techniques in which image computation is relatively trivial, computation is a crux in VAM involving multiple compromises, with a profound effect on the quality and feasibility of prints. Most basically, the challenge of VAM computation lies in the fact that a beam delivering desired dose to a given voxel also inadvertently delivers dose to all the other voxels that lie along the its path. The problem is to choose the intensities of all beams such that when they are integrated together by the recording material, enough optical dose is delivered where required without delivering excess dose to surrounding regions.

The earliest methods of VAM computation attempted to construct desired optical dose distributions by clipping to zero the negative intensities resulting from the Filtered Back Projection (FBP) algorithm. This, however, caused print-shape errors, even under the assumption of perfect optical and material precision. These errors were improved by gradient descent optimization [26], which attempted to achieve dose above and below a gelation threshold to in-part and out-of-part voxels, respectively. Even in this modest goal, however, such methods did not always succeed for complex parts. This is not surprising; a typical 200^3 voxel reconstruction space with 360 azimuthal projection directions yields 14.4 megapixels over which to optimize - a difficult task given the complexity of the optimization problem. Furthermore, as soon as system errors are introduced, such as errors in optical uniformity, timing, material sensitivity, etc., these reconstructions, even when

theoretically perfect, produce print-shape errors. Accurate, practical VAM requires a reconstruction problem statement that includes material response, system error tolerances, etc., as well as methods to achieve the dose-control meeting these requirements. That is, we must understand and define the right problems to solve for various applications, then develop computation methods capable of the dose-control required to meet these challenges.

In this chapter we examine the challenge of image computation for VAM, we introduce quantitative methods of understanding reconstruction quality, and we present simple and effective methods of image computation that significantly improve on prior methods. We begin by defining the simplest VAM problem - that of printing binary structures in a gel-threshold resin - and we propose reconstruction quality metrics relevant to the problem. We present a novel approach to VAM computation, and we use the proposed metrics to demonstrate its improved control and performance. We demonstrate the flexibility of the method with novel print geometries and with the inclusion of inhibited material response.

We then extend the scope of VAM beyond binary printing, towards further increased dose control to enable prints with functionally graded material properties. We present a method for approximating arbitrary grayscale distributions of optical dose, and we examine the additional layers of inverse problems necessary for precise material property control using these grayscale reconstructions. We then consider a combined inverse problem to generalize the scope of tomographic reconstructions beyond *optical dose*, to the more flexible problem of tomographically constructing distributions of *monomer conversion*. We discuss how an expanded problem definition, along with the inclusion of sufficient resin inhibition, enables the application of FBP via the linearity of the back projection operator with respect to material conversion. We discuss how this theoretically makes arbitrary distributions of conversion printable without approximation, but how real world printing precision limits what is practically achievable. We compare such an approach to our iterative optimization method for the simple case of binary printing.

We discuss in depth a novel VAM printing geometry - tomosynthesis - which enables printing into flat sample surfaces. This is exciting for applications such as manufacturing microfluidic

chips, for printing into well plates, or for any application where packaging does not allow for projection access from all equatorial directions. Unlike with the standard tomographic geometry, tomosynthesis allows for an arbitrarily large number of addressable voxels in two dimensions. This enables highly parallelized simulations printing, and the the ability to print arbitrarily large patterned sheets, the thickness of which is equal to the size limitation of a standard tomographic VAM print. We discuss the theoretical resolution limitations of tomosynthesis, and how the highly coupled nature of the print geometry affects binary VAM reconstructions. We present the first proof-of-concept tomosynthetic VAM prints, and we discuss a first attempt at printing into a flat sample package. We examine the problem of coupling light into flat samples, and of archiving a tilted focus plane. We conclude by discussing how tomosynthetic VAM could be extended to the problem of grayscale dose control for the printing of functionally graded materials.

4.1 Binary-VAM Image Computation

The most common objective of 3D polymeric printing generally, and VAM specifically, is to fabricate a 3D part in which voxels receive local optical dose to be either above or below the gelation threshold of the resin. We refer to this here as **Binary VAM** printing. Note that while optical dose is defined as the product of intensity and time, the dose to gelation in real resins changes with intensity (i.e. reciprocity failure [39, 35]), requiring a calibration step to determine dose to gelation for a given intensity. This section discusses how to compute VAM projection images that will achieve this binary preferential gelation.

An ideal VAM image set would deliver gel-inducing optical dose to in-part print voxels while delivering zero dose to out-of-part voxels in the resin. Such an image set is impossible for general print-shapes because light must typically travel through out-of-part voxels to reach in-part voxels. We begin this section by analyzing this limitation more formally via the **Filtered Back Projection** (FBP) Algorithm. We then formulate a problem statement for binary VAM image computation that relies on a material nonlinearity (the gelation dose-threshold). We discuss prior approaches to the problem; we present a novel approach, a resultant algorithm, and we discuss parameter choice and

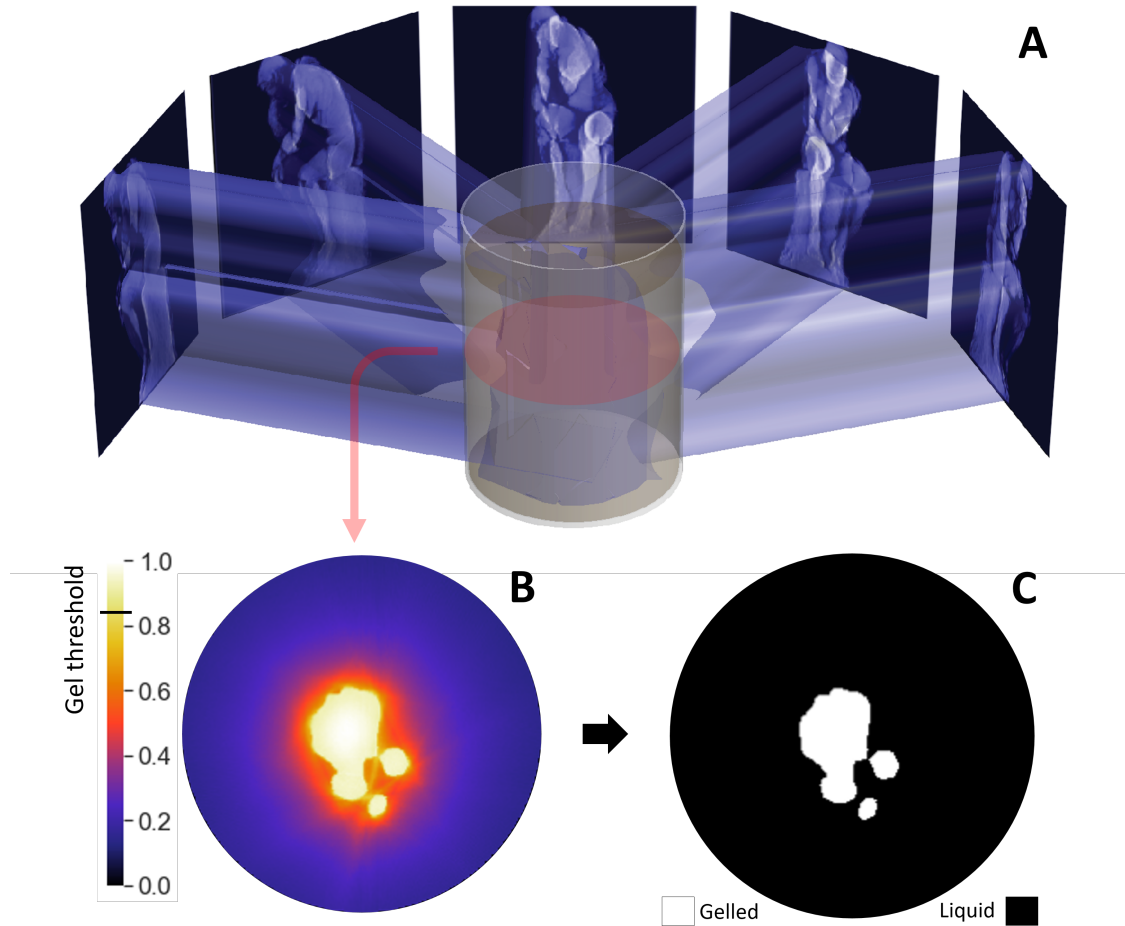


Figure 4.1: Overview of tomographic VAM concept. **A**: Computed images, optically projected into a vial of photosensitive resin. **B**: Cross-section of accumulated optical dose distribution (normalized units of dose). **C**: Perfectly-gelled cross section of a “thinker”; resin gels where the optical dose is above a threshold, forming a 3D printed part that is removed from remaining liquid resin. Image from [2].

reconstruction trade-offs; we define tomographic reconstruction quality metrics and we use them to evaluate the performance of our algorithm; and lastly, we discuss limitations and extensions to the method.

4.1.1 FBP and The Problem of Negative Intensity

Tomography is such an extensive field that it may come as a surprise to the reader that tomographic image computation for VAM is not a trivial, already solved problem. A common

example application is computed tomography (CT) x-ray scans for brain imaging. Here, x-rays are sent through an unknown object (a brain), the object attenuates the x-rays depending on its density distribution, and the intensities of the x-rays are measured. This is repeated for many azimuthal projection angles.

Mathematically, this step of going from an object - a function in *object-space* - to a set of external projections - functions in *image-space* - is known as a **forward projection**. For uniform, non-diverging, parallel projection beams, this can be modeled as a collection of parallel line integrals, with each line of integration representing the attenuation of a ray through the object.

In the medical imaging example, once data in image-space has been collected, a filtering operation is performed, then a simulation of **back projection** is performed on the data. Here, the spatial integration is run in reverse (in simulation), from image-space back into object-space. Each image is projected into object-space, and these projections add in superposition, creating a reconstruction of the original object (e.g. a brain).

The filtering step is necessary to compensate for the non-uniform sampling density inherent to a circular projection geometry. The VAM printing geometry introduced in [26], in which images are projected orthogonally to the resin's axis of rotation (Fig. 4.1A), can be considered as a collection of independent 2D reconstruction problems. This projection geometry is also common for tomographic metrology (e.g. brain scans). Fig. 4.2 shows an example illustrating the issue of non-uniform sampling density of circular projections. Fig. 4.2A shows an example 2D object comprised of four squares of constant value. Performing a forward projection (Fig. 4.2B), then back-projecting the resultant data results in a poor reconstruction (Fig. 4.2C) of the original object. This can be understood via the Projection Slice Theorem [40].

The Projection Slice Theorem says that the data collected by a projection of an object at a given azimuthal angle is the same as the data collected by sampling the spatial Fourier spectrum of the object along a line through its origin, normal to the projection direction. E.g. the spatial Fourier transform of the 0° projection in Fig. 4.2B, running from left to right, is equal to the data sampled by the vertical green line in Fig. 4.2D. Thus, the data collected by forward-projections

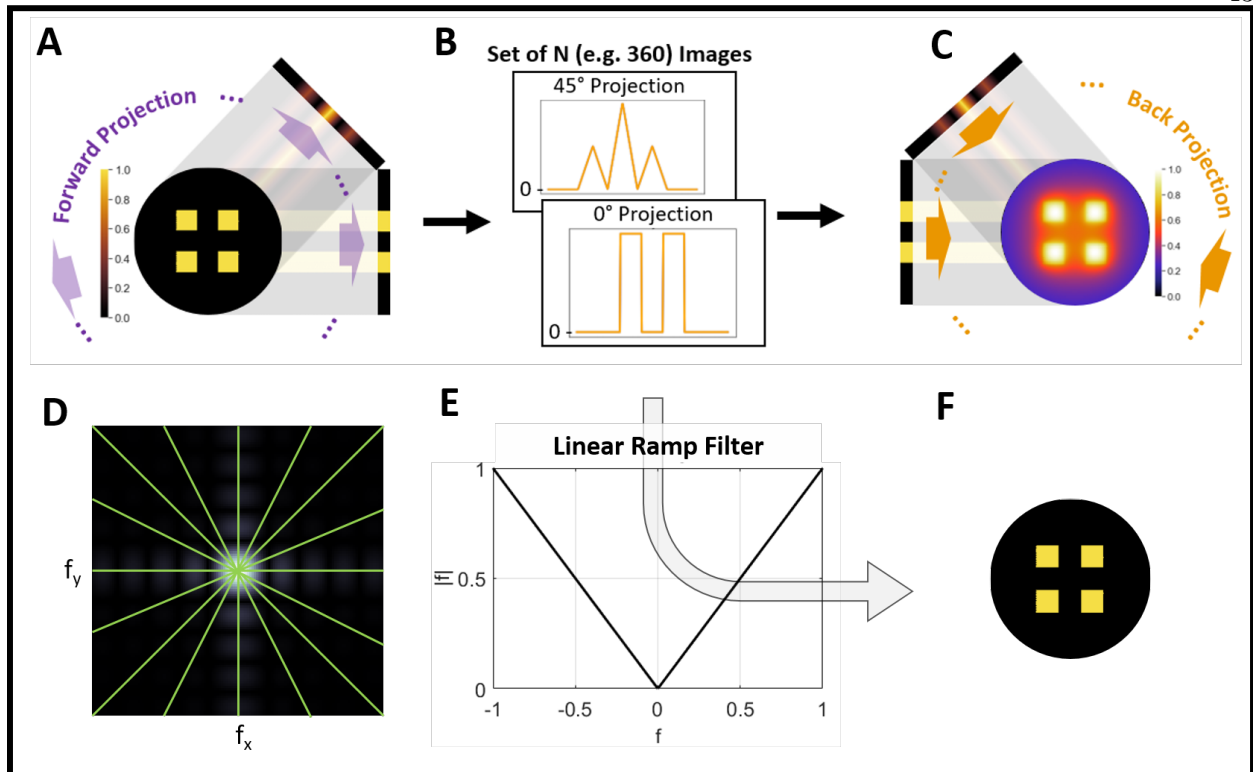


Figure 4.2: Illustration of FBP and how circular projection geometries non-uniformly sample the frequency-space of an object. **A:** A 2D object comprised of four squares of constant value, with forward projection illustrated. **B:** Different projection angle sample different data from the object. **C:** Simply back-projecting the projection data yields a poor reconstruction of the original object. **D:** A spatial Fourier transform of the object in (A). Via the Projection Slice Theorem, each projection image from (B) is equal to the inverse Fourier transform of a line passing through the origin of (D), perpendicular to the projection direction. **E:** The Fourier-space linear ramp filter that must be applied to the data in (B) to correct for non-uniform sampling. **F:** The object in (A) is correctly reconstructed after filtering. This filtering and back-projecting is the FBP algorithm.

of an object can be thought of as a sampling of the object's Fourier spectrum along lines running through its origin.

The sampling density shown in Fig. 4.2D is non-uniform; there is a higher density of green sampling lines near the origin (the center of the spectrum shown in Fig. 4.2D), than near the outside edges. This means that projection set shown in Fig. 4.2B has over-sampled low spatial frequency content with respect to high spatial frequency content. This is why the reconstruction Fig. 4.2C formed by Fig. 4.2B looks like it has some extra glow - that glow is an over-representation of the gradual changes (low frequency content) in Fig. 4.2A, compared to the fast changes (high

frequency content), such as the edges of the squares.

Correcting for this non-uniform sampling is straightforward. The sampling density in Fig. 4.2D falls off linearly with distance from the origin. That means that applying a frequency-space, linear ramp filter (Fig. 4.2E) to the image set (Fig. 4.2B) correctly re-emphasizes the higher spatial frequency content. Back-projecting this filtered data results in an improved reconstruction (Fig. 4.2F). This is the filtered back projection algorithm, and in the limit of continuous functions, infinite angular sampling, and well-behaved objects [40], it will yield perfect reconstructions.

While discrete approximations to FBP work well for tomographic metrology, FBP is not directly suitable to VAM. Because FBP's filtering step in frequency space (Fig. 4.2E) is analogous to performing a derivative on the projection data (Fig. 4.2B), it produces a filtered image set that contains negative values. This is an intuitively obvious result; if we wish to deliver positive values to in-part voxels, with beams that pass through out-of-part voxels, we need negative beams to apply to those out-of-part voxels to erase the unwanted values. In VAM, however, we must project images as intensity patterns. This leaves us with a serious problem: the lack of darkons in the lab means that we cannot project negative intensity values. It is for this reason that FBP cannot be directly applied to the problem of tomographically reconstructing arbitrary optical dose distributions via projections of patterned intensity. Even more alarming is that FBP essentially provides an invertible, one-to-one transformation between object-space and image-space (see [40] for details about why FBP is not technically invertible, one-to-one). Thus the fact that it cannot be used for VAM suggests that there exists no solution for tomographic reconstructions in VAM.

4.1.2 A Binary VAM Problem Statement and Prior Approaches

To overcome the fact that FBP cannot be directly applied to VAM, we must define the objective of computation for binary VAM to include a material nonlinearity. That is, since only FBP (or equivalent, such as projection back filtration) can reconstruct a target with perfect contrast, and VAM cannot directly use FBP to reconstruct optical dose, we must relax the constraints on our reconstructions. Instead of trying to compute image sets that create a dose distribution exactly

matching the geometry of the target-part, we instead seek a reconstruction with dose that is above a gelation threshold for in-part voxels, while being below the threshold for out-of-part voxels, thus achieving preferential gelation. That is, not only do we allow for imperfect contrast between in-part and out-of-part dose, but we also allow for non-uniformity in both in-part and out-of-part regions. Moreover, for practical reasons in the lab relating to precision limitations of timing, projector uniformity, resin formulation, etc., we wish the out-of-part voxel dose to be significantly lower than the in-part voxel dose. Without sufficient dose contrast between the desired part and the background, the VAM system must be excessively precise. We will discuss this more in the next section, but for now let us consider the objective of binary VAM computation to be: to achieve a dose reconstruction in which the lowest in-part doses are significantly higher than the highest out-of-part doses, such that the gelation dose-threshold can be placed between them with room for error to spare.

Perhaps the easiest method of computing such a VAM image set - and the one that was used in the initial demonstration of VAM [26] - is to use FBP, but to clip any resultant negative values to zero. This creates crude approximations to object reconstructions. For some objects, this approximation produces recognizable reconstructions. However, for many objects, the results are disastrous: voids are filled in, horn-like structures appear growing off of corners, and the reconstructions are generally quite inaccurate. This is not surprising: the FBP image set was precisely computed to have negative beams balanced against positive beams in just the correct spatial pattern to produce an accurate reconstruction. Arbitrarily setting the negative beams to zero without compensation destroys this balance and produces a poor reconstruction.

The first non-trivial VAM computation algorithms [26, 3] attempted to correct the effects of setting the negative values produced by FBP equal to zero. In general, the approach is to: **1.** Clip any negative values that result from FBP, or add a constant to the image set to avoid the negatives, or some combination of both. **2.** Back project the image set. **3.** Compare the resultant reconstruction to the original target function via some sort of error function. **4.** Slightly change the image set based on this error function via a gradient-descent optimization method. These steps are

then repeated until the image-set hopefully produces a reconstruction that causes the error function to be acceptably low. While these methods do, in fact, improve the outcome, VAM computation could significantly benefit from improvement in three areas: **ease of implementation and use**, computational **flexibility**, and **print accuracy** under material and optical imprecision.

In the following section, we show that algebraically optimizing directly on the structure, instead of the set of external projection images, enables notable improvement in these areas. We refer to this approach as *object-space model optimization* (OSMO), and we discuss multiple advancements: **1. Print Accuracy.** For complex parts, prior methods suffer from errors in the shape of a final part, especially when optical or material precision is imperfect. We show that by improving optical dose contrast between the target and surrounding regions, OSMO produces dose-reconstructions that can tolerate optical and material uncertainties, while still resulting in perfectly gelled prints. **2. Flexibility.** We show that this approach naturally allows for the application of object-space constraints such as dose-limits, as well as the application of object-space physical models, such as resin oxygen inhibition, Gaussian beam models, and occlusions for over-printing applications. **3. Ease of implementation and use.** The approach is intuitive and includes only two parameters, both of which directly set reconstruction dose constraints. While these parameters afford easy dose control, tuning them from their default values is unnecessary for fast and reliable convergence to desirable reconstructions. OSMO is easy to code in a handful of lines, and consists of only four algebraic steps per iteration. The ease of use and flexibility of this approach additionally facilitate VAM modalities that do not yet exist in the field, but that could enable new applications of VAM. This includes new printing geometries that would enable printing into non-cylindrical packages such as flat slabs without index-matching baths. It also includes the problem of printing arbitrary functionally graded materials by controlling the degree of polymerization – not just the state of gelation – throughout a printed part, a problem that we refer to as **Grayscale VAM**. We conclude by discussing how the availability of generalized, efficient, and simple image-computation algorithms could help enable new applications in the field of VAM, while avoiding the need for researchers to understand and tune complex computation algorithms.

4.1.3 Object Space Image Computation for Binary VAM Printing

In this section, we present an algebraic object space model optimization (OSMO) algorithm for computing image sets to 3D-print arbitrary structures in gel-threshold resins. The algorithm is easy to implement and run, it is flexible, and it yields improved accuracy compared to previously published alternatives, and has only two simple parameters that directly tune reconstruction properties in object space. We first describe the algorithm, we define and discuss reconstruction quality metrics, and we consider how to choose the two constraint parameters to achieve a desired balance of reconstruction quality metrics. We discuss an object-filtering step to reduce the number of OSMO iterations needed for convergence, and we compare the algorithm to previous methods. Note that the 2D reconstructions shown here can be trivially extended to 3D as a stack of independent 2D reconstructions for the printing geometry shown in Fig. 4.1. Computation for new printing geometries that cannot be decoupled into independent 2D regions is more challenging and is discussed in Chapter 5, but remains algorithmically unchanged from the basic 2D case.

4.1.3.1 OSMO Algorithm for Binary VAM

For a simple VAM projection model where each projection beam is assumed to be non-diverging, and attenuation through the resin obeys Beer-Lambert exponential intensity decay [41], transformation between volumetric dose and image set is well-described by the Attenuated Radon Transform (ART) [26, 3]. **Back projection** (projection from image-set into volume), as is physically implemented while printing, is similarly modeled with exponentially attenuated rays. Formally, back projection is the adjoint of the **forward projection** operator (from volume to image set) or, in the discrete case, its matrix transpose. The simplest method of generating an image set from a desired object is to mathematically perform a forward projection from object to image set, using the ART. However, back-projecting the resultant image-set results in a blurred reconstruction, a consequence of the suppressed high spatial frequencies inherent to forward then back projecting in a radial geometry [40]. Fig. 4.3, shows a target object, a binary model of the object,

and the reconstructed dose that results from forward projecting the model, then back projecting the resultant set of images. Here, instead of optimizing the intermediate image-set to avoid this blurring, we instead iteratively optimize a **model** of the desired structure, which, when then used to compute an image set, results in a desirable dose reconstruction.

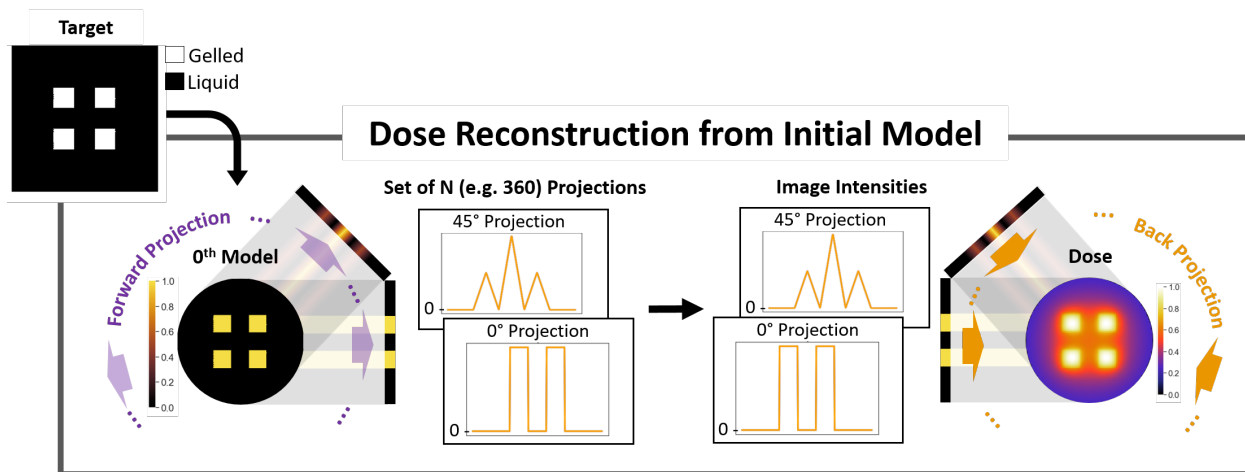


Figure 4.3: Dose reconstruction via forward projecting a binary model of a target geometry, then back projecting the resultant set of images. The reconstruction is blurred due to suppressed high spatial frequencies inherent to forward then back projecting. Instead of optimizing image intensities in projection-space to improve the reconstructed dose, we will instead optimize the model to improve dose. This first reconstruction serves as an initial step in the model-optimization algorithm shown in the next figure. Note that since all projection values are non-negative in this case (as no filtering nor thresholding was applied), they can be represented by intensities with no modification. Image from [2].

The model can contain negative values, however it yields image sets and dose reconstructions that are non-negative, and thus physically realizable. The process of optimizing a model is shown in Fig. 4.4A. The process of computing a non-negative image set (and its dose reconstruction) from a given model is shown in Fig. 3B – this process is used during the optimization procedure (Fig. 4.4A), as well as to compute a final image set and reconstruction. In this process, the model is forward-projected, and resultant negative image-values are set to zero. This image set is then back projected to form a reconstruction. The reconstruction is then normalized in preparation for optimization calculations (shown in Fig. 4.4A), or to produce a final reconstruction. Once optimization is complete, with a desired image set computed, the image set is scaled to match the

intensity of a real optical projection

The OS MO algorithm (Fig. 4.4A) tries to push in-part dose to be above an upper threshold D_h , and out-of-part dose to be below a lower threshold D_l by adjusting the model. Increasing the value of an in-part voxel in a model will increase the corresponding voxel in the resultant dose reconstruction. However, it will also increase dose in other reconstruction voxels, some of which may be out-of-part. These unwanted out-of-part doses can then be reduced by decreasing their corresponding model voxels. The OS MO algorithm alternates between increasing in-part model voxels and decreasing out-of-part model voxels to achieve a model that yields a desirable reconstruction. The algorithm can be mathematically defined as follows: Let f_T be a target function (the desired geometry to print). Let P be any forward-projection operator (e.g. the ART), and P^* the corresponding back-projection operator. That is, Pf_T is the forward projection of f_T , from object to image set, and the dose reconstruction P^*Pf_T is the back projection of that image set, back into object space. Let N be a normalizing operator that divides an input by its maximum. Let D_l and D_h be fixed high and low dose-threshold values, where $0 < D_l < D_h < 1$. Let $M_{j,j}$ be the object model after j iterations, where the initial model is defined to be $M_{0,0} = f_T$. Lastly, let $f_{j,j} = NP^*max(0, PM_{j,j})$. That is, the model is forward projected, any resultant negative values are set to zero, the result is back projected, and the reconstruction is normalized. Then, the steps to update a model $M_{j,j}$ by one iteration to $M_{j+1,j+1}$ are as follows:

STEP 1: Update the model by subtracting from out-of-part voxels with an excess of dose. For only the out-of-part voxels (in-part voxels remain unchanged from previous model, $M_{j,j}$):

$$M_{j,j+1} = M_{j,j} - max(0, f_{j,j} - D_l), \quad (4.1)$$

producing the intermediate reconstruction:

$$f_{j+1,j} = NP^*max(0, PM_{j+1,j}). \quad (4.2)$$

STEP 2: Update the model by adding missing dose to the in-part voxels. For only

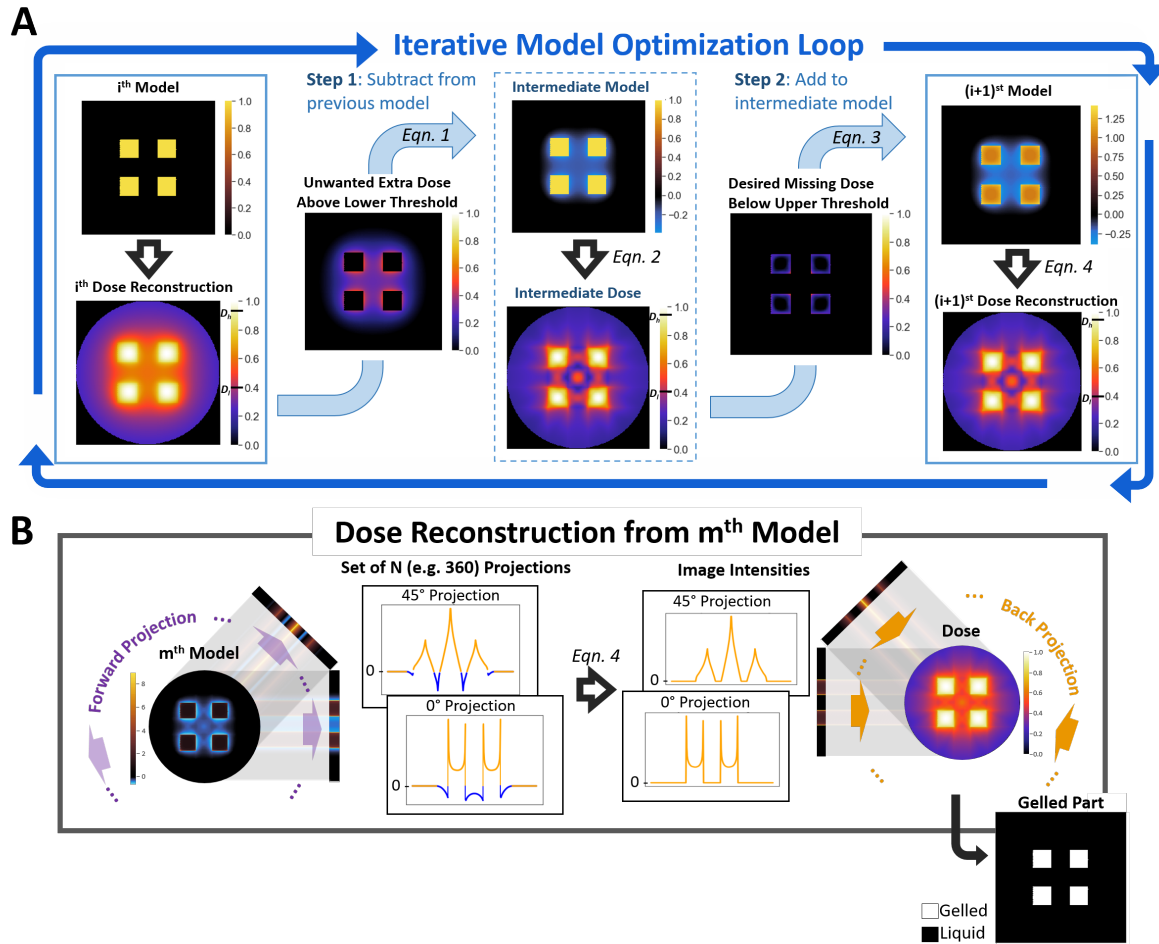


Figure 4.4: Object space model optimization (OSMO) algorithm for binary-gelation VAM printing. **A**: One model-update iteration of the OSMO loop. Each model in the loop produces a dose-reconstruction by the process shown in (B) forward projection, clipping any resultant negative values to zero, then back-projecting and normalizing to form a dose-reconstruction. Negative values must be clipped to convert projections to physically realizable intensities. Note that, unlike all subsequent models, the initial model (Fig. 4.3) has no negative values in its projections, and thus does not require a clipping step. **A, Left Column**: A given i^{th} model (in this case, the initial model is pictured) produces a reconstruction. **A, Step 1**: Unwanted extra dose in the out-of-part regions of the reconstruction are computed and subtracted from the i^{th} model to form an intermediate model. Only the extra dose that lies above a lower threshold, D_l , is included. **A, Middle Column**: The intermediate model is then used to compute a dose-reconstruction. **A, Step 2**: The prior step of subtracting out-of-part values from the model, inadvertently reduces some desired dose in the in-part regions. This missing in-part dose is calculated by comparing the in-part reconstruction voxels to an upper threshold, D_h . These values are then added to the intermediate model to form a $(i+1)^{\text{st}}$ model, completing one model-update iteration. **B**: The model produced by m iterations, the dose reconstruction it produces, and the resultant perfectly gelled part. Computationally, the gel threshold is modeled by setting all voxels above a gelation-dose to unity, and setting all voxels below that dose to zero. Image from [2].

in-part voxels (out-of-part voxels remain unchanged from $M_{j+1,j}$):

$$M_{j+1,j+1} = M_{j+1,j} + \mathbf{max}(0, D_h - f_{j+1,j}), \quad (4.3)$$

producing the reconstruction:

$$f_{j+1,j+1} = NP^* \mathbf{max}(0, PM_{j+1,j+1}). \quad (4.4)$$

The $\mathbf{max}()$ operation in STEP 1 means that for each out-of-part voxel, some extra dose is subtracted from the previous model, but not all way to zero; instead, the value is subtracted only down to D_l . Similarly, the $\mathbf{max}()$ operation in STEP 2 means that for each in-part voxel, some extra dose is added to the previous model, but only to reach the value of D_h , not to reach the highest-possible value of 1. Instead of trying to achieve a perfect reconstruction (a binary dose distribution), the algorithm is essentially pushing the reconstruction towards a reduced-contrast solution, allowing the in-part and out-of-part regions to be non-uniform in dose, as long as they lie above and below D_h and D_l , respectively. Thus, the two parameters have physical meaning in the object-space; they set the lower and upper goal-bounds on in-part and out-of-part dose, respectively. This algorithm is succinct and easy to code, especially in any language that includes projection functions (see Appendix A for a sample implementation in Matlab). It is also easy to understand, as computations are made in object space, not in the more abstract set of writing-images, and it has only two parameters, both which directly correspond to object space dose constraints. The two parameters, D_l and D_h , push the solution in contradictory directions; one enforces a decrease of out-of-part dose, and one an increase of in-part dose. The choice of these values affects competing aspects of reconstruction quality. Before discussing how to best choose these values, we must first discuss what our objectives should be; that is, what makes for an accurate reconstruction.

4.1.3.2 Reconstruction quality metrics

We define three metrics for evaluating the quality of a reconstruction: Voxel Error Rate (VER), Process Window (PW) size, and In-Part Dose-Range (IPDR). All three are simplest to understand while viewing in-part and out-of-part histograms of reconstruction dose, e.g. as shown

in Fig. 4.5. In normalized units of dose, a perfect reconstruction would yield an out-of-part histogram with values only at 0, and an in-part histogram with values only at 1. In reality, accurate reconstructions yield two non-overlapping histogram peaks, such that when the gelation threshold lies between them, all in-part voxels gel, while all out-of-part voxels do not. An example of a reconstruction with non-overlapping in-part and out-of-part dose histograms is shown in Fig. 4.5.

We define *Voxel Error Rate*, to quantify histogram overlap, as $VER = \frac{W}{N}$, where W is the number of out-of-part voxels that receive more dose than the lowest-dosed in-part voxel, and N is the total number of voxels in the reconstruction. Graphically, VER is the area of the out-of-part histogram that overlaps with the in-part histogram, normalized to the sum of the histogram areas. Any overlap between histograms indicates errors in the final gelled part, as any gel threshold will result in some error voxels (voxels that gelled when they should not have, and vice versa). VER is roughly proportional to Jaccard Index [26] for gelled reconstructions. However, VER does not require that a gel threshold be chosen, and thus is a more fundamental to the evaluation of tomographic dose reconstructions than Jaccard Index. For any reconstruction with $VER > 0$, the final part will not match the target function, even when the gelation threshold is an ideal step function positioned optimally between the dose histograms. We define *Process Window* to be the difference, in normalized units of dose, between the highest-dose out-of-part voxel, and the lowest-dose in-part voxel. Graphically, this is simply the distance between the inner edges of the histograms; it is negative when $VER > 0$. A large positive process window can allow for, without penalty to VER , a non-ideal gel threshold that occurs over a range of dose, for non-uniformities in optics, for uncertainty in the gel threshold dose location (e.g. due to imprecisions in measuring photo-initiator concentration when mixing resins), etc. Therefore, a large process window is desirable, as it makes it practically easier to gel the in-part voxels, without inadvertently delivering too much dose and also gelling any out-of-part voxels. Lastly, we define *In-Part Dose Range* (IPDR) to be 1 minus the lowest dose in-part voxel, in normalized units of dose. It is an indication of dose-variation of in-part-voxels. This metric is important in VAM printing because

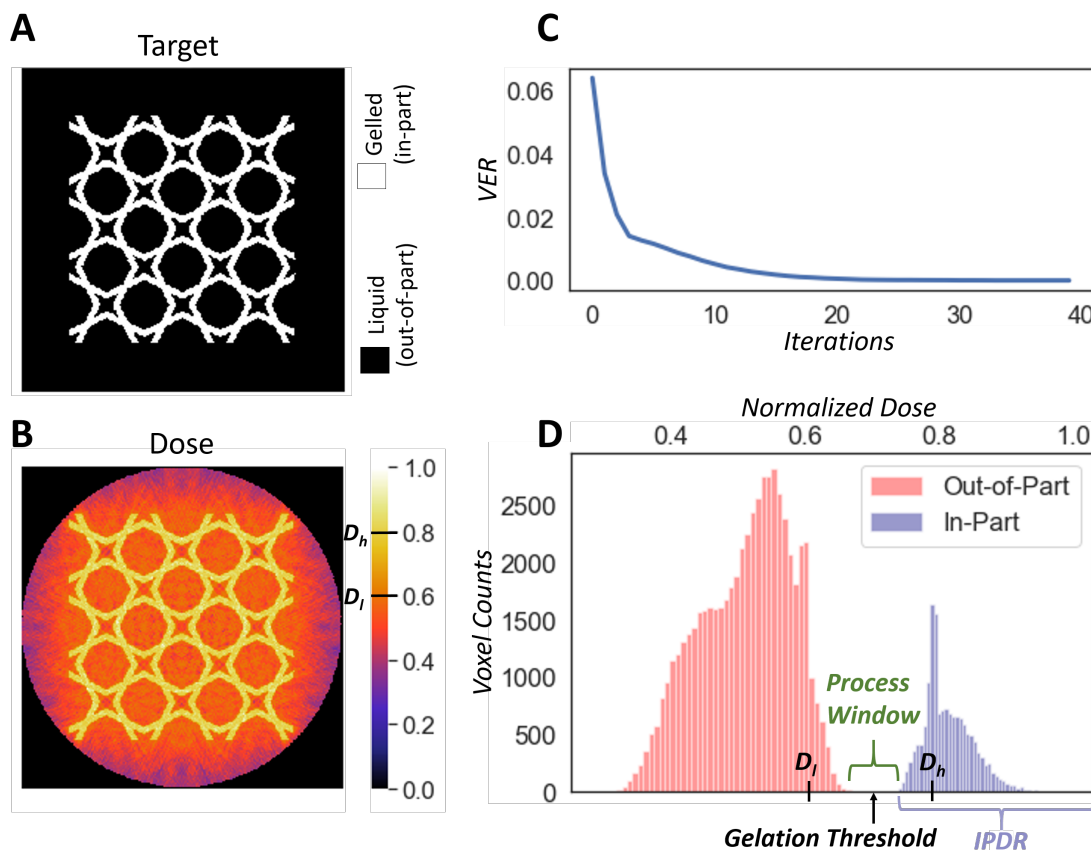


Figure 4.5: Reconstruction Quality Metrics. **A**: A target geometry to be printed. **B**: An optimized dose-reconstruction, shown in normalized units of dose. The parameter values were $D_l = 0.6$ and $D_h = 0.8$. **C**: **Voxel Error Rate (VER)** converging to zero. **D**: The distribution of doses in the reconstruction for the in-part and out-of-part regions. Since the histograms do not overlap, VER is 0, and thus the gel threshold is above all out-of-part voxels, and below all in-part voxels. Therefore, this dose-distribution will print a part that theoretically exactly matches the target geometry. The **Process Window** is the distance between histograms, and represents room for error with optics or materials, without a change in the final gelled shape. As long as dose is applied such that the gel threshold lies within the process window, the part will gel in the correct shape. The **In-Part Dose Range (IPDR)** is the range of doses, as a fraction of maximum dose, that in-part voxels receive. Image from [2].

we cannot deliver all optical dose instantly. Instead, it is delivered over time, by a projection system with limited dynamic range. A consequence of a large IPDR is that high-dose in-part voxels can surpass the gel threshold before the rest of the part. Once this happens, these gelled regions (with their increased index of refraction) will scatter subsequent writing beams, interfering with the completion of dose-delivery to the rest of the part. We will discuss this effect, and how to

computationally mitigate it, in greater detail in Section 4.1.7. All else equal, reconstructions with a small IPDR are desirable.

4.1.3.3 Choosing the two parameter values

The choice of the two parameter values must reflect the fact that a perfect dose reconstruction (where dose is delivered only to in-part regions) is impossible without negative image values. If a perfect dose-reconstruction were possible, the parameter values would be chosen to be $D_l = 0$, and $D_h = 1$. Then the algorithm would push out-of-part dose to zero, and in-part dose to 1 (in normalized units of dose). However, since such a reconstruction is not possible (due to the lack of negative-intensity images), choosing these extreme values results in a poor reconstruction with non-zero VER (thus leading to poor print fidelity upon gelation). Relaxing these constraints reduces the contrast of the dose-reconstruction but can drive VER to zero. This is because choosing $0 < D_l$, and $D_h < 1$ increases the freedom of the reconstruction, since voxels are constrained only to a range, not to particular values; in-part doses are above D_h , while out-of-part doses are below D_l . This produces reconstructions with reduced contrast, but with improved VER. An example of this is shown in Fig. 4.6.

Tuning the parameter values for each new target is unnecessary for fast and reliable convergence to desirable reconstructions. While some parameter optimization may slightly improve particular target-functions, the default values $D_l = 0.85$, and $D_h = 0.9$ resulted in reliable convergence to a VER of zero for all targets we tested, with the number of iterations required for convergence depending on target-complexity, and ranging from tens to hundreds of iterations. Choosing lower values, such as $D_l = 0.45$, $D_h = 0.5$, results in a slightly larger process window, which is unsurprising, as the increase in IPDR relaxes some reconstruction constraints for in-part voxels. Similarly, choosing values such as $D_l = 0.99$, and $D_h = 1$ results in a very small IPDR, at the cost process window size, and, for complex target structures, even reduced VER. Similarly, it is possible to request a larger than possible process window size from the algorithm, by choosing excessively separated parameter values. E.g. for all but simple targets, a choice like $D_l = 0.2$,

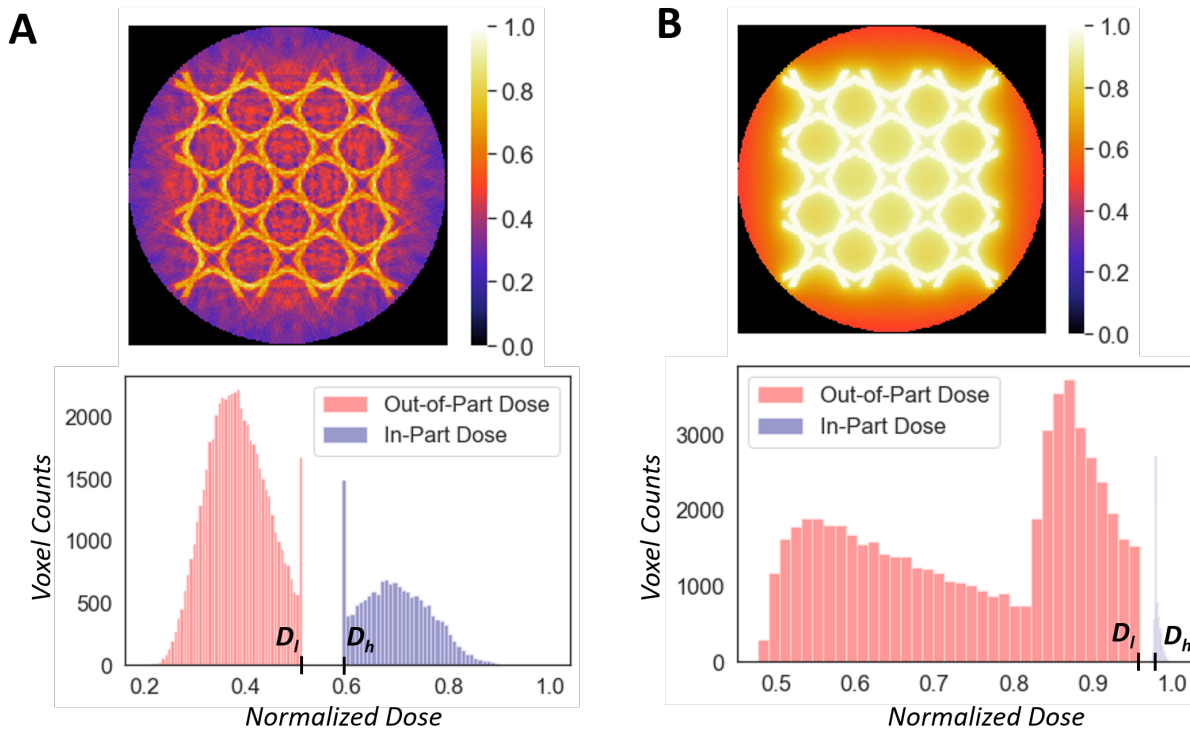


Figure 4.6: Trade-off between process window size, and IPDR. **A, Top:** Binary VAM reconstruction with $D_l = 0.51$, and $D_h = 0.59$. **A, Bottom:** Histograms of in-part and out-of-part dose. **B, Top:** Binary VAM reconstruction with $D_l = 0.95$, and $D_h = 0.98$. **B, Bottom:** Histograms of in-part and out-of-part dose. Reconstruction-A has a process window over 3 times larger than that of Reconstruction-B, but an IPDR that is over 10 times larger. Note that choosing parameter values that try to simultaneously accomplish a small IPDR and a large process window (e.g. $D_l = 0.7$, and $D_h = 0.98$) does not work – trying to achieve such a reconstruction results in histogram overlap, and a non-zero VER. Image from [2].

and $D_h = 0.9$ will not only fail to increase process window size, but it may also increase VER from zero. Lastly, as the choice of parameter values changes the magnitude by which object-voxel doses change upon each iteration, they can slightly affect the rate of convergence of the algorithm. Regardless of optimization possibilities, the default values yield reconstructions that improve upon previous methods, as discussed in section 4.1.3.5.

4.1.3.4 Object-Space Frequency Filtering – An Improved Initial Model

The OSMO algorithm reliably converges with the target geometry as the initial model. However, an improved initial model can significantly reduce the number of iterations needed to converge to the same result, saving computational run time. The initial model does not change the target nor any reconstruction constraints, but rather serves as initial guess for the OSMO algorithm. Using the target as an initial model results in an initially blurred reconstruction (Fig. 4.3), a consequence of the suppressed high spatial frequencies inherent to forward then back projecting in a radial geometry. FBP compensates for this by emphasizing high spatial frequencies with a linear frequency filter applied to the forward projected images [40]. While this results in negative-intensity projector image values, clipping these image-space values to zero yields an improved initial guess for gradient descent image-optimization methods [26]. We can use the same concept with OSMO by applying an object-space frequency filter to the target to generate an initial model. An object-space filter for the standard tomographic geometry (Fig. 4.1) is a 2D linear frequency-ramp of the form f_r , where $f_r = \sqrt{(f_x^2 + f_y^2)}$ is a radial coordinate in the Fourier-spectrum of a 2D object-slice, with Cartesian coordinates f_x and f_y . It is one dimension larger than FBP's projection-space filter since it applies to the object, and not to the image-set. Unlike FBP's filtering step which is computed for each 1D projection image, the object-space filtering step is computed for each object slice. This 2D filter is the same as is used in back projection filtration in which $f = R_{2D}P^*Pf$, where f is a 2D target slice, R_{2D} is an operator applying f_r , and P and P^* are forward and back projection operators. That is, just as P^*R_{1D} is a left inverse of P in FBP (where R_{1D} is the operator applying FBP's 1D filter), $R_{2D}P^*$ is also a left inverse of P . The difference between the filters emerges when negative-clipping projector-image values, since the negative values produced by the 2D filter are first smoothed by forward projection, while FBP's 1D filter produces negative values after forward projection occurs. An example of the application of this 2D filter to a target function is shown in Fig. 4.7.

A filtered target does not provide a reconstruction with a VER of zero; iterative optimiza-

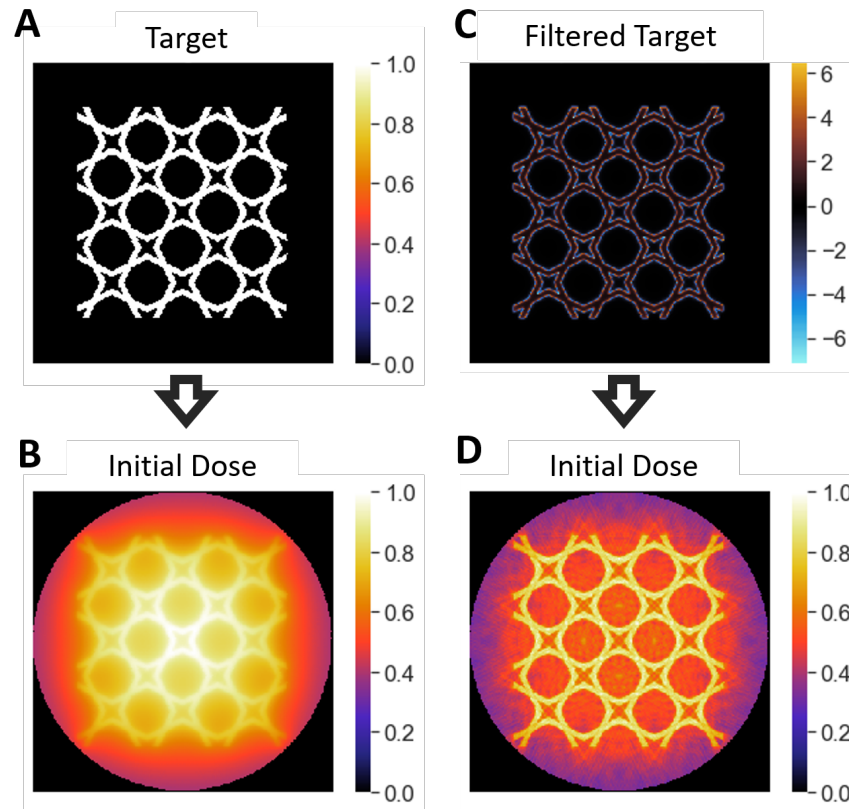


Figure 4.7: Frequency-filtering a target to generate an improved initial model. This reduces the number of algorithm iterations necessary to converge to good solution, compared to simply using the target as the initial model. For this target, with $D_l = 0.85$ and $D_h = 0.9$, required iterations were reduced by 25%. **A**: A target geometry. **B**: An initial reconstruction that results from using the target geometry as the initial model, a process illustrated in Fig. 4.4. **C**: The target geometry with a 2D frequency ramp filter applied. **D**: The initial reconstruction that results from using the filtered target as a model. Both initial models converge to qualitatively identical final dose reconstructions, with identical quality-metric values. Image from [2].

tion is needed to achieve that because the forward projection of filtered target results in negative numbers. Nor does initial filtering change the quality metrics of a final optimized reconstruction. However, it does typically reduce the number of iterations needed to converge to a VER of zero. E.g., for the target in Fig. 4.7, filtering reduces the number of OSMO iterations to reach a VER of zero by 25% (with $D_l = 0.85$, $D_h = 0.9$ and 360 projection angles). The exact iteration savings depends on the particular target and on the choice of parameter values, and can even be negative. Fig. 4.8 shows one of these unusual cases, along with convergence from initial models of unity or

of random numbers from zero to one. However, we observe significant iteration savings in most

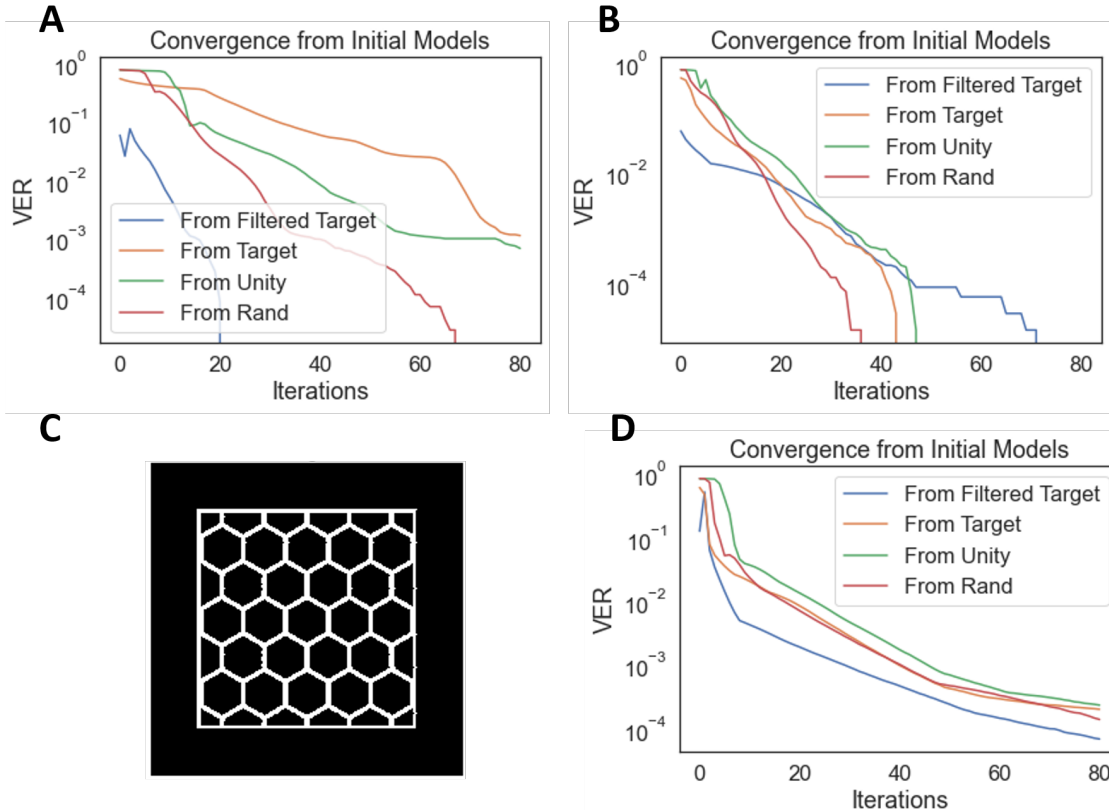


Figure 4.8: Examples of OSMO convergence from different initial guesses: the original target geometry, the filtered target as described above, an array of ones, and an array of random numbers between zero and one. **A and B**: Convergence for the mesh structure shown in Fig. 4.7A with 720 projection angles. For the plot shown in (A), D_l was 0.85 and D_h was 0.9. For (B), D_l was 0.65 and D_h was 0.75. This is an unusual case in which filtering slows convergence to $VER = 0$. **C**: A different target geometry, yielding the convergence shown in **D** with $D_l = 0.65$ and $D_h = 0.75$.

cases, at only the cost of the filtering step, which takes about 1% the run time of a single OSMO iteration. In summary, we find that applying a 2D ramp filter to the initial model significantly reduces the average run time of the OSMO algorithm.

4.1.3.5 Performance Comparison Example to Prior Methods

Here, we consider a binary VAM reconstruction example to compare OSMO to the state of the art in the literature: the CAL algorithm [26]. The original computed axial lithography (CAL) method [26] (referred to as “CAL2019” in what follows) uses gradient-descent optimization

on image-sets. It converges to a VER of zero for simple geometries, and is available on GitHub [42]. Another recent method [43] based on thresholded FBP, reports a relatively large number of voxel errors even for a simple geometry, so we do not show comparison calculations to this method. We consider an example reconstruction to compare the performance of OSMO and CAL2019. To ensure that CAL2019 parameters were chosen with expertise for this comparison, the algorithm was run by a member (authored here) of the research group that maintains it.

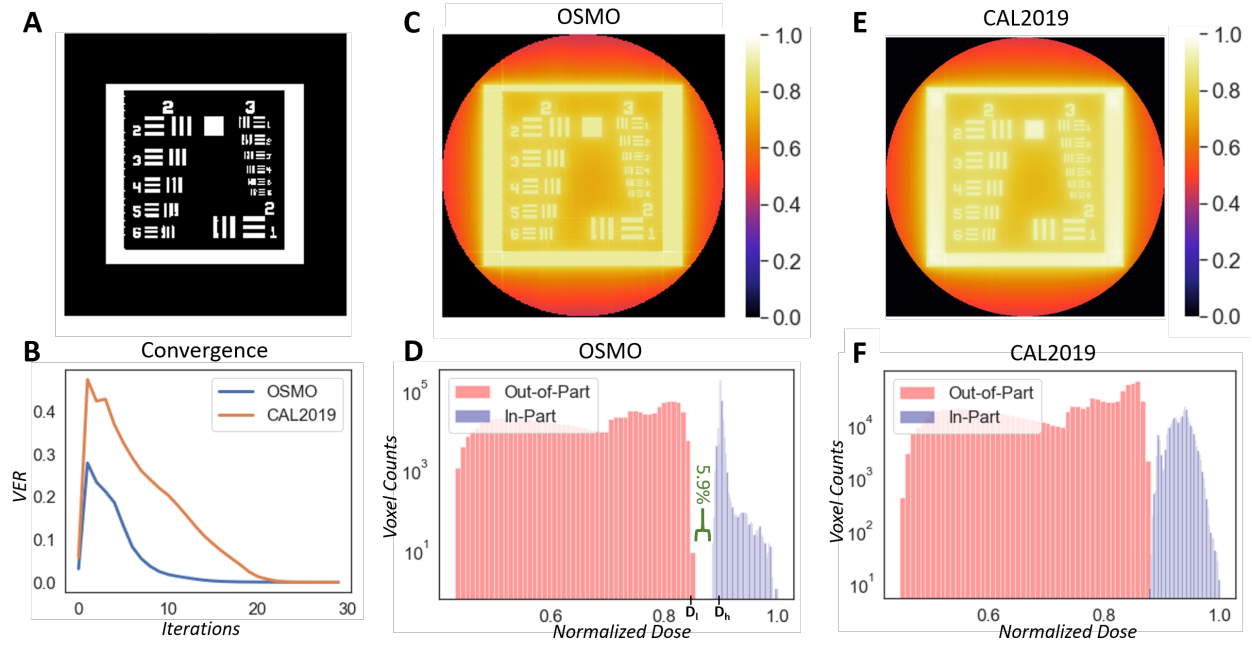


Figure 4.9: Convergence comparison example with the previously-published CAL2019 method. **A**: Binary target geometry. **B**: VER convergence for the first 30 iterations of the OSMO and CAL2019 methods. Note that the initial VER value for OSMO is slightly smaller than for CAL2019, due to different initial guesses: 2D model filtering versus clipped FBP, respectively. While CAL2019 plateaus to a small VER near zero (10^{-5}), further iterations of OSMO drive the reconstruction beyond a VER of zero, to a positive process window size. The CAL2019 reconstruction had an IPDR of 0.12 and a process window size of -0.23%. **C and D**: OSMO dose reconstruction and histogram. The OSMO parameter values were $D_l = 0.83$, and $D_h = 0.9$, $VER = 0$, and the process window size was 5.9% of the maximum dose with an IPDR of 0.11. **E and F**: CAL2019 dose reconstruction and histogram. The histogram overlap yields a non-zero VER value with no positive process window size. Since $VER > 0$, a print would have voxels with gelation error. While these voxels would be few, the lack of a process window would result in further error in a physical implementation, due to imperfect precision in optics and materials. Image from [2].

We find that OSMO improves upon the reconstruction quality of CAL2019. Fig. 4.9A shows

the object used for this comparison: a resolution target. Fig 4.9B shows the convergence of VER towards zero for the first 30 iterations of CAL2019 and OSMO. The initial VER for OSMO is slightly lower than that for CAL2019 because OSMO uses an object-space filtered initial guess while CAL2019 uses a clipped FBP initial guess. Both algorithms show an initial increase in VER as they work towards a reconstruction with the requested IPDR. CAL2019 plateaus to a VER near zero (10^{-5}), while further iterations of OSMO drive the reconstruction past a VER of zero to a process window size of 5.9% of the maximum dose. While the number CAL2019 error voxels are few, its lack of process window means that a physical implementation would require perfect precision in optics and materials to achieve that low level of gelation errors. The OSMO 5.9% process window allows for some errors in material and dose-delivery, while still resulting in perfect gelation. See Figs. 4.10, 4.11, and 4.12 for exposure timing and optical non-uniformity error simulations for both CAL2019 and OSMO.

Both methods take on the order of a few seconds per iteration on a typical desktop computer. CAL2019 has six hyperparameters which control: (1) learning rate, (2) descent momentum, (3) prescribed or floating threshold (note: prescribed threshold, not originally in CAL2019, was added to make a fair comparison between OSMO and CAL2019), (4) non-negativity constraint relaxation, (5) width of sigmoid thresholding function, and (6) weight of robust thresholding heuristic. For more details, the reader is directed to the Supplementary Information of Kelly Et al.[26]. Although, typically only the first two or three are manually tuned to arrive at a sufficient reconstruction given a new structure. OSMO has 2 parameters whose default values do not require tuning to produce a positive process window size for a new structure. OSMO uses four projection operations per iteration – twice the number that CAL2019 uses. In Matlab, with its relatively slow projections (compared to e.g. [44]), OSMO takes about 3 times longer per iteration than CAL2019, however OSMO iterations produce an improved reconstruction, and some of that additional time may be offset by a lack of parameter-tuning.

While this comparison (Fig. 4.9) is illustrated using a 2D object of relatively high complexity, it supports the observation that OSMO typically results in quality metric improvements - especially

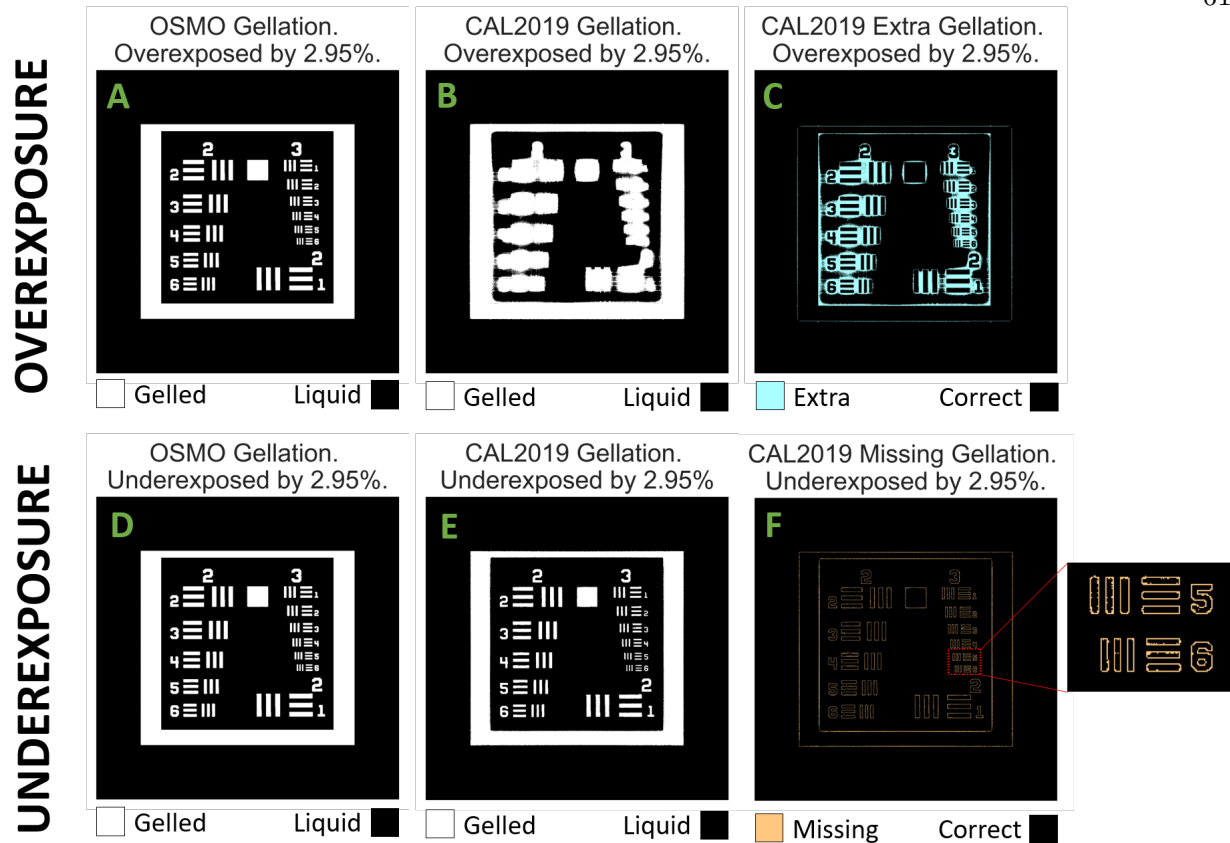


Figure 4.10: Modeled effects of total dose errors upon gelation for OSMO and CAL2019. Reconstructions shown in Fig. 4.9. **A**: Modeled gelation when the OSMO reconstruction (shown in Fig. 4.9C) is overexposed by 2.95%. Since this error falls within the reconstruction process window, no gelation errors occur, and the final simulated part matched the target geometry. **B**: CAL2019 gelation upon the same 2.95% overexposure. **C**: The extra gelation that occurs in (C) due to this error. **D**: Modeled gelation when the OSMO reconstruction (shown in Fig. 4.9C) is underexposed by 2.95%. Since this error falls within the reconstruction process window, no gelation errors occur. **E**: CAL2019 gelation upon the same 2.95% underexposure. **F**: the missing gelation in (E) that occurred due to underexposing the CAL2019 reconstruction by 2.95%. Inset shows the missing gelation size relative to small print features. Image from [2].

in increased process window size. OSMO's improvement in final output, coupled with its simplicity - both to run and to code - provides a significant advancement in computation for binary VAM.

4.1.4 Scaling Between Independent Reconstruction Slices

For the circular VAM geometry in which images are projected orthogonally to the resin's axis of rotation (Fig. 4.1A), a 3D computation problem can be split into a stack of independent

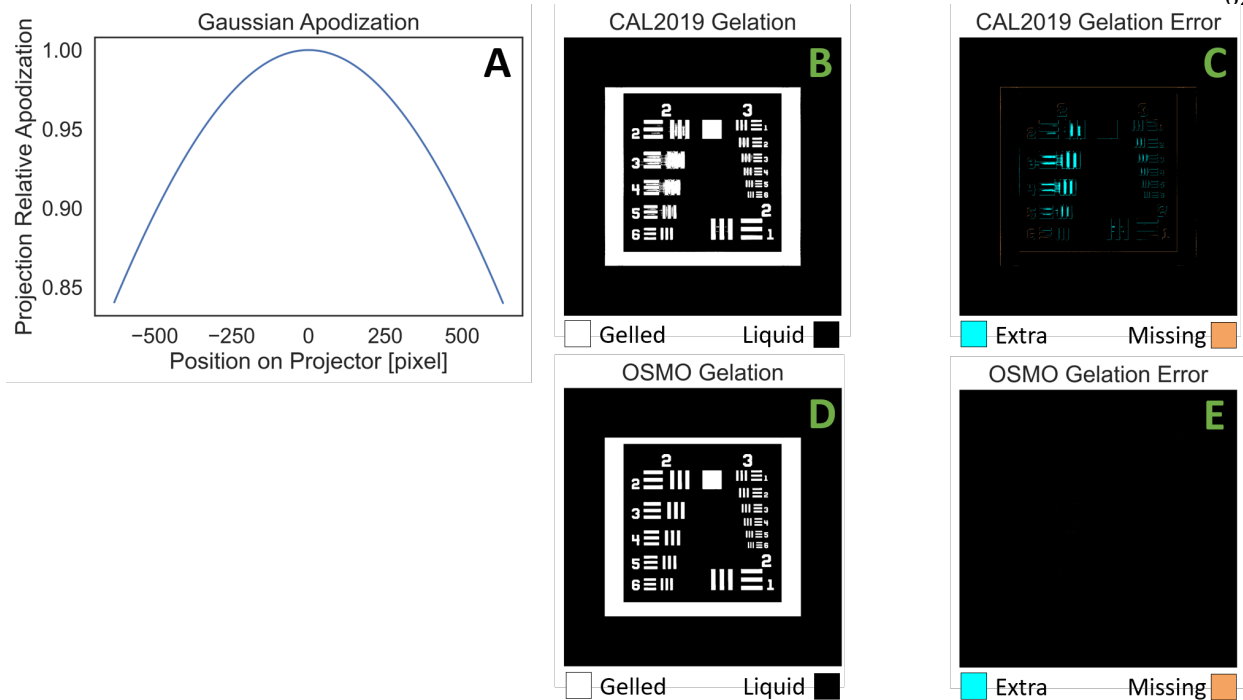


Figure 4.11: Modeled effects of total gelation errors due to projector non-uniformity, with peak intensity at the center of the projections. Apodization functions were applied to the image sets computed in Fig. 4.9. before back projecting them to form reconstructions. **A**: Gaussian apodization (0.84 at edges, and unity at center) applied to each projector image computed in Fig. 4.9. **B**: Resultant gelation for CAL2019 - the unwanted and missing gelation voxels are detailed in **C**. For (B/C), VER was 0.16, and the threshold that minimized the number of gelation voxel errors was 0.87. **D**: Resultant gelation for OSMO with the apodization in (A) applied to all projections. **E**: Extra and missing gelation from (D). Here, VER was $1.8E-5$, and the threshold that minimized the number of gelation voxel errors was 0.89. Image from [2].

2D problems. When OSMO optimizes a given slice, a normalization is performed on the resultant reconstruction. However, in the lab the reconstruction slices are not independent - their relative scaling is coupled by the optical intensity incident on the DMD. Consider an optimized OSMO model, producing the image set $g(\theta, x, y)$, where θ is the azimuthal angular projection coordinate, and (x, y) are DMD pixel coordinates, with y running parallel to the resin's axis of rotation. Back projecting g produces the 3D dose reconstruction $F(Z, X, Y)$, where Z is coordinate indexing 2D slice regions orthogonal to the axis of rotation, and (X, Y) are Cartesian coordinates within each slice. Each row i of DMD pixels $g(\theta, x, i)$, produces a 2D dose-reconstruction $F(i, X, Y)$. However, without performing some sort of scaling between independently computed slices, the dose-ranges

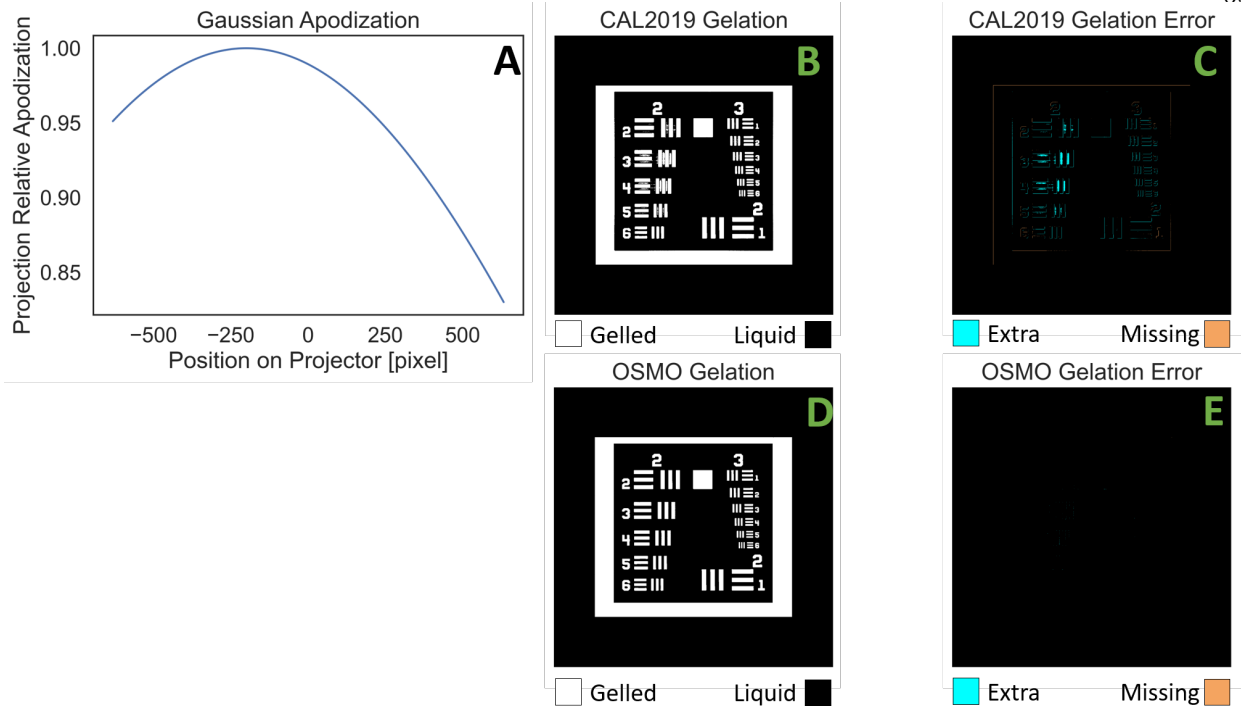


Figure 4.12: Modeled effects of total gelation errors due to projector non-uniformity, with peak intensity offset from the center of the projections. Apodization functions were applied to the image sets computed in Fig. 4.9. before back projecting them to form reconstructions. **A**: An apodization function, with a maximum value of 1 offset by 20% the total width of the function, with a value of 0.95 at the left edge, and a value of 0.83 at the right edge. **B**: Resultant gelation of applying (A) to the CAL2019 image set. **C**: The resultant extra and missing gelation voxels in (B). Here, VER was 0.15, and the threshold that minimized the number of gelation voxel errors was 0.87. **D**: Resultant gelation from applying (A) to the OSMO image set computed in Fig. 4.9. **E**: The resultant extra and missing gelation voxels in (D). Here, VER was $7.8E-5$, and the threshold that minimized the number of gelation voxel errors was 0.89. Image from [2].

in the different slices of F could be very different. Thus, we must perform some sort of scaling on the rows of g so that the resultant slices of F are scaled to one another in a desirable way.

One option is to scale the image set such that each reconstruction slice will have the same maximum dose. That is, we **scale to the maximum in-part dose** of each slice. To do this, for each slice, we back-project $g(\theta, x, i)$ to compute the 2D reconstruction $F(i, X, Y)$. Then the new image set for that slices becomes $g_{scaled}(\theta, x, i) = g(\theta, x, i)/max(F(i, X, Y))$. Then, each slice-region of the reconstruction will have the same maximum dose.

Another scaling option is to scale the image set such that each reconstruction slice will have

the same minimum in-part dose. That is, we **scale to the minimum in-part dose** of each slice. As before, we back-project $g(\theta, x, i)$ to compute the 2D reconstruction $F(i, X, Y)$. Then new image set for that slice becomes $g_{scaled}(\theta, x, i) = g(\theta, x, i)/\min(F(i, X_{in\ part}, Y_{in\ part}))$, where $X_{in\ part}, Y_{in\ part}$ index the set of in-part reconstruction voxels. Then, each slice-region of the reconstruction will finish gelling simultaneously. This simplifies detecting the completion of a print via a shadowgram, as the entire print will simultaneously stop changing upon print completion.

One advantage of not scaling to the maximum in-part dose has to do with the variability of reconstruction metrics. Consider a reconstruction for which each slice has been scaled to the maximum in-part dose. If reconstruction metrics for each slice turn out to be identical, then the entire volume will finish gelling at the same time, there will be a single global process window in which to fit timing and other errors, and the 3D part will print with the same behavior that a single reconstruction slice would predict. However, it is often the case that different 2D slice-regions will exhibit varying degrees of optimization success. For example, suppose that the OSMO algorithm is used with $D_l = 0.8$ and $D_h = 0.9$. For 2D slices with a simple geometry, the largest out-of-part dose will likely be 0.8, and the smallest in-part dose will likely be 0.9. However, for a particularly complex 2D structure, this choice of D_l and D_h will likely demand a larger process window than is possible, and the resultant reconstruction will not match these values. E.g., perhaps the largest out-of-part dose will be 0.84, with the smallest in-part dose being 0.86. Then, the 3D reconstruction would contain slices with varying minimum in-part dose, yielding a reduced global process window. Instead, for example, scaling slices to one end of the process window - like the minimum in-part dose - has the effect of increasing process window overlap between slices, thus improving the global process window size. An example of this is shown in Fig. 4.13.

Lastly, another option is to scale the image set such that each reconstruction slice will have the same mid-process-window dose. That is, we **scale to the process window mid-point** of each slice. To do this, we compute each scaled image as:

$$g_{scaled}(\theta, x, i) = g(\theta, x, i)/\text{mean}(\min(F(i, X_{in\ part}, Y_{in\ part})), \max(F(i, X_{out\ of\ part}, Y_{out\ of\ part}))).$$

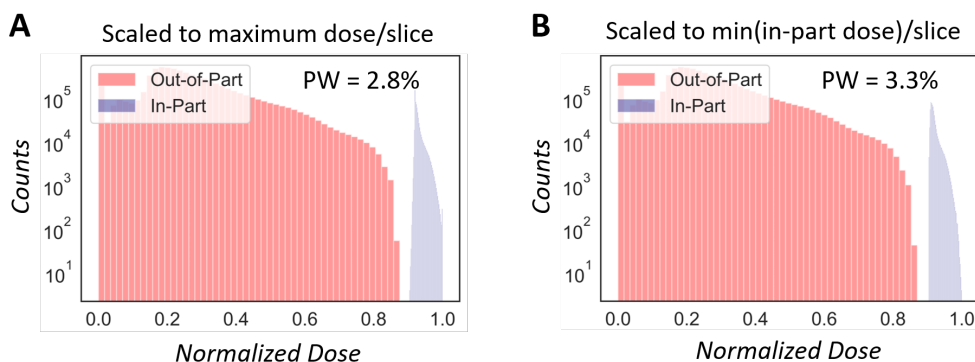


Figure 4.13: Global dose histograms from different slice-scalings for reconstruction of a Thinker geometry. **A**: Scaling the independent 2D slices to have the same maximum reconstruction dose results in a process window of 2.8%. **B**: Scaling the slices to the minimum in-part dose increases the process window size to 3.3%. This is due to the increased process-window overlap caused by scaling the slices to have the high end of each process window equal in dose, thus increasing the global process window.

With this scaling, a gel point can be chosen to fall exactly in the center of each slice's process window. This minimizes the chance that any given slice will suffer gelatin errors in the event of e.g. a timing error during printing.

A different method of plotting reconstruction details can help elucidate the effects of different slice scaling choices. Fig. 4.14 shows an example of this using a reconstruction of a Thinker geometry for which slices were scaled to the minimum in-part dose. Here, for each slice, index on the vertical axis, three pieces of data are plotted. A red dot indicates the largest out-of-part dose. The left end the blue line indicates the smallest in-part doses, and the right end indicates the maximum in-part dose for the slice. The process window of a given slice is the distance from the red dot to the left end of the line. The length of the line is the IPDR. This helps visualize the reconstruction properties of different slices. In this example, we can see which slices have the highest overall dose by seeing which lines extend furthest to the right; these are the slices that will start to gel first. Similarly, the slices that will start to gel last are the ones with the lowest maximum dose, with the tight line-end furthest to the left. This also lets us see which slice will first experience unwanted out-of-part gelation upon over exposure.

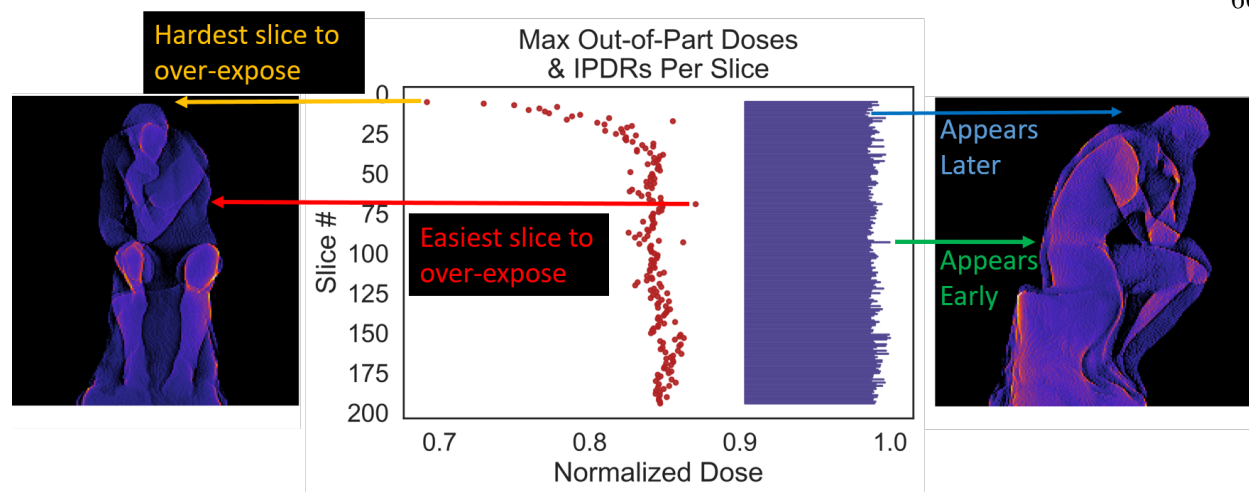


Figure 4.14: Method of plotting reconstruction dose details to better visualize slice-scaling choices. For each slice (vertical axis), a red dot indicates the maximum out-of-part dose. The length of the blue line is the IPDR. The left end of a given line is the minimum in-part dose for that slice, and the right end is the maximum dose.

We can apply this method of plotting to the three types of scaling discussed here. See Fig. 4.15 for an example comparing three methods of re-scaling for the Thinker geometry. For this particular structure, we can see that scaling to the mid point of the process window results in different slices both starting and ending gelation at significantly different times. We also see that, in this case, there is very little difference between scaling to the smallest in-part dose or to the maximum dose. Our default choice is to scale to the minimum in-part dose for simultaneous completion of all layers.

Re-scaling between reconstruction slices is an important step that adds an additional degree of freedom to a 3D reconstruction. As we showed in Fig. 4.13, re-scaling can significantly improve the global properties of a tomographic VAM reconstruction. Improvement will be greatest for collections of 2D reconstructions in which some slices are more difficult to optimize than the rest. A re-scaling step is not specific to the OSMO algorithm, but can be applied to any collection of reconstruction slices, regardless of how they were computed.

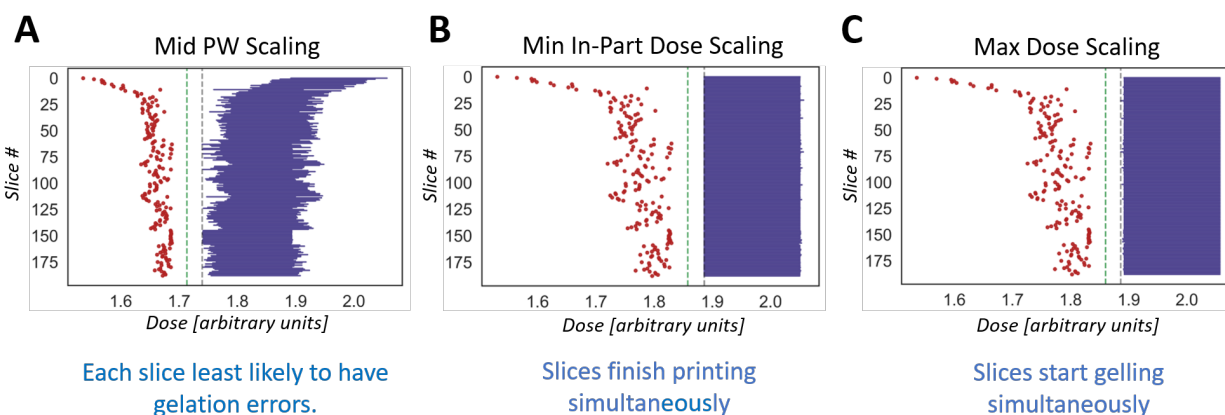


Figure 4.15: Per-slice reconstruction details from re-scaling slices of a Thinker geometry using the three criteria presented in this section. For each slice: the red dot is the largest out-of-part dose, the left end of the horizontal blue line is the smallest in-part dose, the right end of the blue line is the largest dose, and the green line represents the gelation threshold that is at the midpoint between the rightmost red dot and the leftmost line-end.

4.1.5 Scaling to Partial Rotations

Upon computing a VAM image set, we are left with a set of projection images in arbitrary units [a.u.]. Additional steps are required to apply this idealized tomographic reconstruction to the problem of predicting and controlling real VAM prints.

Although a single full angular range is typically computed during image optimization (e.g. 360 images at 1 image/degree), real printing does not usually occur over exactly one rotation of the resin. Printing usually occurs over 5 to 30 rotations, depending on projector intensity and resin sensitivity. The simplest approach to printing over many rotations is to repeat the same computed image set for each rotation. However, we often run into the problem that a print does not complete in exactly an integer number of rotations. This is a problem if the image-computation step assumes an integer number (e.g. 1) of rotations.

For example, suppose a print completes after 7.25 rotations using an image-set that was computed assuming a single rotation. The first quarter of the image set is projected 8 times, but the last 3/4 of the images are projected only 7 times. Thus, in summation, the first quarter of the images (e.g. the first 90 images in a 360 set) are topographically over-represented by about

14%. Fortunately, the mathematical correction for this is trivial - we must simply multiply the first quarter of the image set by $7/8$. This would result in a reconstruction that is proportional to the originally-computed reconstruction. Unfortunately, it can be difficult to predict exactly how many rotations will be required for a print to form. For this, we need good process control: small tolerances on optical intensity, and high precision on material sensitivity - including material nonreciprocity.

4.1.6 Scaling Tomographic Reconstructions to Real Optical Units

The first step towards process control in a VAM print is to translate a normalized computed reconstruction into real optical units specific to a printer. This can be done with a few simple steps.

First, we must determine the maximum intensity projected by the DMD pixels. A simple way to do this is to display a white (100%) image, then measure the projected intensity at the print plane. Let us refer to this value as I_0 . Then, we multiply the image set, $g(\theta, x, i)$, by I_0 . We also must multiply by the percentage of the time that the DMD projects an image, as there is a short dark-state when switching between images. This yields an image set that is scaled to intensity, g_{scaled} .

Back-projecting g_{scaled} in simulation yields a reconstruction in units of intensity. However, in simulation, we assume that all projections occur simultaneously, while in reality they are multiplexed over time. Thus, to compute the optical dose from a print, we must compute the duration of time spanning each image projection. For example, in an image set with 360 projections, each image is projected $(1/360)^{th}$ of the time. So the dose in a reconstruction could be computed as $g_{scaled}t_{print}/360$, where t_{print} is the total duration of the print. This leaves us with a simulated reconstruction in which each voxel has a unit of optical dose.

This method of computing dose is not unique; many methods are possible. More importantly, it is not optical dose that we ultimately care about - it is material conversion. Computing reconstruction dose loses the intensity information that is important for accounting for material non-reciprocity - a problem that we will discuss later. However, computing intensity in real units,

as described here, is a prerequisite step to predicting and controlling real VAM prints.

4.1.7 Controlling Print Dynamics

Above, we have calculated final dose distributions and ignored the fact that this dose is deposited by the exposure of each image at a particular intensity for a finite time. However, the computed dose delivered by each image does not have to be applied in a single exposure. By visiting each angle many times in a single print it is possible to use different images for each rotation, which may result in additional control over the print dynamics. Instead of considering an image set to contain e.g. 360 images (1 per degree) that are then repeated over N rotations, we can instead consider the image set to contain $360 \times N$ independent images. This increases control over the dynamics of dose delivery within the total exposure time. We hypothesize this can improve print quality e.g. by reducing the early gelation of localized volumes which can sink or refract subsequent writing beams.

We can avoid gelling a portion of the print long before the rest of the part forms by delivering the majority of the dose at the end of the of the printing process. While the OSMO algorithm does not inherently account for temporal variations in the resin, the approximate energy doses delivered over time can be estimated by considering a fraction of the total projections. By choosing to order the projections such that the high intensity projections occur at the end of the build, the in-part regions are expected to cross the gelation dose-threshold nearly simultaneously. Equivalent dose distributions can be generated by re-distributing the pixel intensities across multiple rotations as given by equation 4.5.

$$g_0 n_{rot} = \sum_{i=1}^{n_{rot}} g_i \quad (4.5)$$

where:

$g_0(\theta, y, x)$ = image set assuming a single rotation, where θ ranges from 0 to 360°

$g_i(\theta, y, x, i)$ = image set for i of n_{rot} rotations

n_{rot} = number of rotations over which to vary the light field.

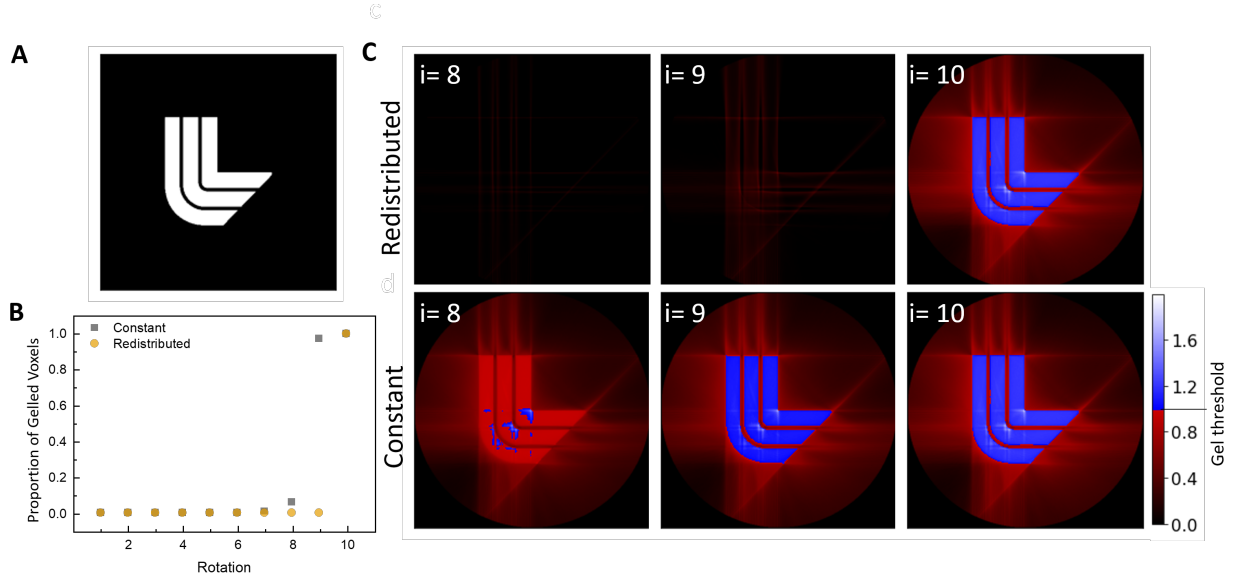


Figure 4.16: Re-binning pixels across multiple rotations enables simultaneous part formation. **A**: Binary target geometry. **B**: When a constant image set (gray squares) is projected across 10 rotations, the object starts forming (surpassing the gel threshold) after 9 rotations. When the intensity of pixels is re-binned such that the majority of the dose is delivered at the end of the build (orange circles), the object does not form until the final rotation. **C**: Rate of gelation for a constant image set versus the image-set with value redistributed as per Eqn. 4.6. **C, Top Row**: Dose reconstruction after the 8^{th} , 9^{th} and 10^{th} rotation for redistributed image set. **C, Bottom Row**: Dose reconstruction after the 8^{th} , 9^{th} and 10^{th} rotation for the constant image set. Image from [2].

While any combination of g_i can be used such that equation 5 is satisfied, to deliver the majority of the dose at the end of the build, we define g_i as:

$$g_i = \begin{cases} n_i \max(g_0), & 0 < n_i < 1 \\ \max(g_0), & n_i > 1 \\ 0, & n_i < 0, \end{cases} \quad (4.6)$$

where:

$$n_i = \frac{k_i - g_0}{k_{i+1} - k_i}$$

$$k_i = \frac{\max(g_0)}{n_{rot}} (n_{rot} - i)$$

where: i = the rotation number $1, \dots, n_{rot}$. Note that k_i is defined up to $i = n_{rot} + 1$ to calculate g_i up to $i = n_{rot}$. k_i is the proportion of the highest value projection pixels remaining in the image set after rotation number i , and n_i is a normalized difference between the single-rotation image set and this remaining proportion. Equation 4.6 maximizes projection pixel values in the latter rotations as much as possible, while not setting any pixel value above $\max(g_0)$, and while still obeying equation 4.5. E.g., suppose an image set g_0 has a maximum value of 1, a particular pixel has a value of 0.4, and there are a total of 3 rotations. This pixel contributes a total value of 1.2 over three rotations. Equation 4.6, would set this pixel value to 1 for the last rotation, to 0.2 for the second rotation, and to 0 for the first rotation. This does not change the total contribution of this pixel, but it allows the pixel to deliver its dose as late as possible.

For an example reconstruction, when the same image set is used over 10 rotations, voxels in the object start reaching the gelation threshold after 6 rotations (Fig. 4.16). However, when the pixel intensities are redistributed across 10 rotations as described by equation 6, object voxels do not surpass the gelation dose-threshold until the final rotation. Note that this requires good process control, as it assumes that the print will run until the full completion of the last rotation. Experimental results demonstrating such an avoidance of early gelation are shown in Fig. 4.17.

This process can be applied to any build to minimize the time between when the first and last in-part voxel crosses the gelation threshold. This method allows all regions of a given part to cross the gelation dose-threshold nearly simultaneously, with the largest benefit gained by reconstructions with a large IPDR. Simultaneous gelation minimizes the scattering due to increases in index of refraction as the polymer gels, and the problem of parts sinking since gelled regions are denser than their monomer counterparts. This technique could likely be combined, to good effect, with the latent VAM method described in Section 6.2, especially since the last rotation of re-binned exposures better matches a uniform flood exposure than does a single-rotation image set.

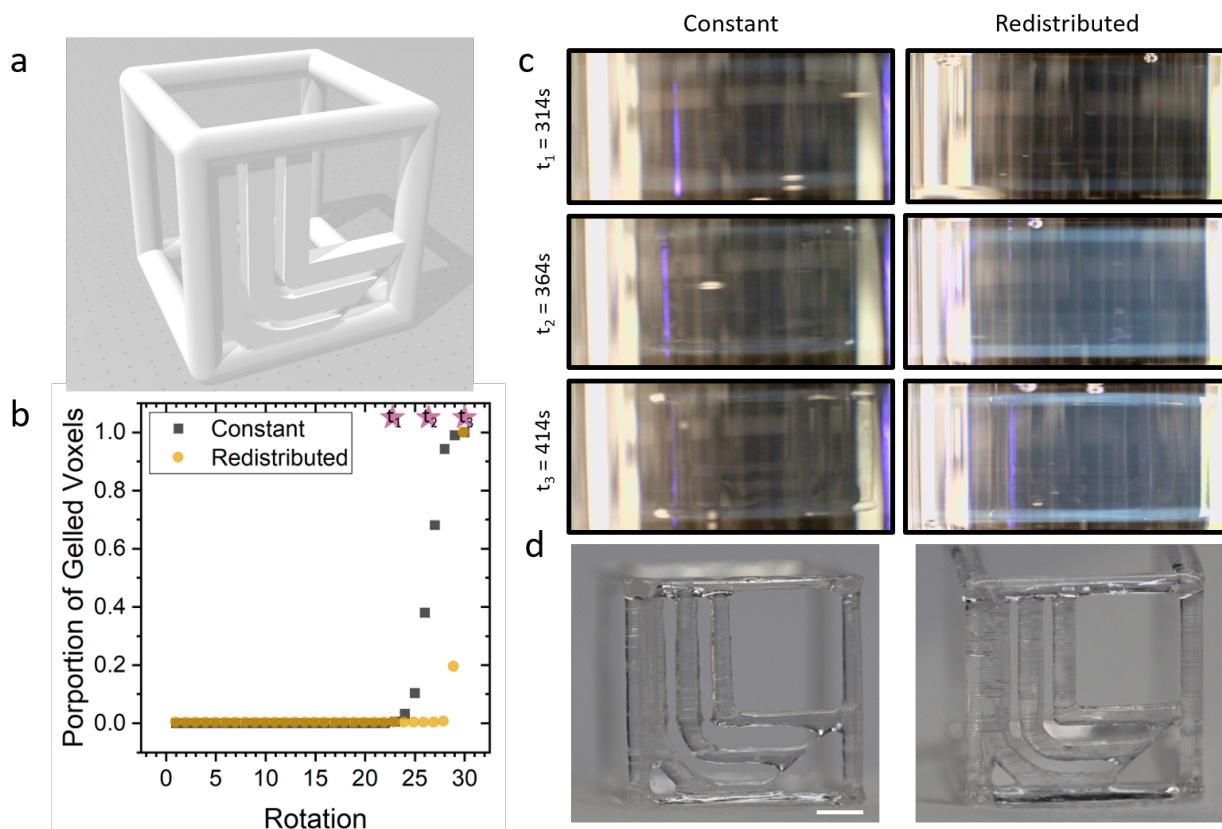


Figure 4.17: Experimental demonstration of redistributing light intensities. **A**: model 3D part. **B**: Calculated proportion of gelled voxels after each rotation when the intensity is either constant for each rotation (gray squares) or redistributed across 30 rotations such that the majority of the dose is delivered at the end of the build (orange circles). **C**: Snapshots from videos of parts forming without (constant) and with redistributing the light intensities at the time points denoted in (B) by pink stars. At t_1 no object is visible in either the constant nor the redistributed light intensities. At t_2 the edges of the cube are visible for the constant case, but no part is visible in the redistributed case. At t_3 the full part is visible in both cases. **D**: Extracted parts for the constant (left) and redistributed (right) light intensities. Scale bar: 2 mm. Experimental parts were printed using a custom VAM printer equipped with a 405 nm LED light engine (CEL5500, Digital Light Innovations), with a maximum intensity of $55\text{mW}/\text{cm}^2$ illuminating a resin vial fixtured to a rotation mount (HDR50, ThorLabs). [3] Bisphenol A glycerolate diacrylate (BPAGDA), poly(ethylene glycol) diacrylate (Mn 250) (PEGDA), and - Methyl-4'-(methylthio)-2-morpholinopropiophenone and Irgacure 907 photoinitiator were purchased from Sigma Aldrich and used without further purification. The photosensitive resin mixture is a 75:25 wt% mixture of BPAGDA and PEGDA, respectively, with 40 mM of Irgacure 907. [4]. Image from [2].

4.1.8 Including Material Response

Section 4.1.3 used an idealized polymer resin described only by a threshold dose of gelation. Algebraic optimization in object space facilitates more general material response models.

4.1.9 Binary Printing with Inhibition

Resins containing radical inhibitors exhibit a delay in the commencement of monomer conversion as optical dose accumulates. This delay, or inhibition period, is caused by radical inhibitors reacting with photo-generated radicals at a much higher rate than monomer-radical reactions. Only once inhibitor is locally depleted does significant conversion commence. Oxygen can provide a small inhibition period for acrylate and methacrylate resins using radical initiators, and is often present in photo-sensitive resins when they are used in ambient conditions. A larger inhibition period can be engineered by the inclusion of a radical inhibitor [28]. In VAM printing, inhibition relaxes the requirements of an optical tomographic reconstruction, as it increases the allowable dose without gelation in the out-of-part regions. We find that even a small inhibitory dose significantly improves the contrast between in-part and out-of-part regions in a VAM print. The OSMO algorithm can be easily changed to model this effect. With respect to conversion, inhibition effectively subtracts from the dose applied in object-space since some amount of dose causes no conversion. Including inhibition is mathematically equivalent to including negative values in image-sets, with larger inhibition periods corresponding to larger negative magnitudes. By the linearity of Radon Transform, this effect is equivalently described as a subtraction in image-space, resulting in negative values in simulated image-sets. Thus, for computational convenience, we can apply the subtracting effect of inhibition in image-space during the negative-clipping step when computing a reconstruction from a model (Fig. 4.4B). After forward projecting a model, instead of clipping all resultant negative values to zero, we clip them to the negative value g_{min} , where the ratio of $-g_{min}$ to 1 is the same as the ratio of the inhibition period to the conversion period of the resin. Then, after calculating a reconstruction with a back projection, any negative dose values in object-space are set to zero,

since they correspond to no polymeric conversion. That is, all of the reconstruction steps described in Section 4.1.3 such as $f_{j,j} = NP^* \max(0, PM_{j,j})$, become $f_{j,j} = \max(0, NP^* \max(g_{min}, PM_{j,j}))$. Once a final image-set has been computed, the images are shifted by a constant such that their smallest value is zero (again relying on the linearity of the Radon Transform). This process is illustrated in Fig. 4.18. The rest of the algorithm remains unchanged.

Fig. 4.19 shows reconstructions comparing the effects of including material inhibition versus including no inhibition. Here, including only a 10% inhibition period results in a 4x increase in process window size. While both reconstructions would result in perfect gelation under ideal assumptions, including inhibition makes the reconstruction four times more tolerant to optical or material imprecisions, while still resulting in perfect gelation.

4.1.10 Additional Printing Geometries and Optical Models for Binary VAM

4.1.11 Image Computation for More Complex Projection Models

Real VAM printing does not occur with ideal, perfectly sharp, perfectly non-diverging beams. Likewise, when over-printing around occluding structures, the ideal tomographic condition where each beam spatially overlaps with all other beams, is broken. Fortunately, the OSMO algorithm still works, without modification, for more complex projection models than the Attenuated Radon Transform (ART). Here, we demonstrate its outputs for two projection models: for Gaussian beams, and for printing around occlusions. Gaussian beam projections can be modeled by changing the line integrals used in the ART from straight paths, to curved paths represented by a Gaussian kernel. This tends to concentrate dose towards the center of a reconstruction. It also effectively reduces the degrees of freedom of a reconstruction, since each image-set pixel addresses a larger region within the reconstruction (the out-of-focus parts of the Gaussian beams), resulting in an increased overlap with other writing beams. Both effects increase the difficulty of achieving good tomographic reconstructions, with more iterations needed to achieve an inferior reconstruction (e.g. reduced process window size), as compared to the ideal non-diverging beam case. An example

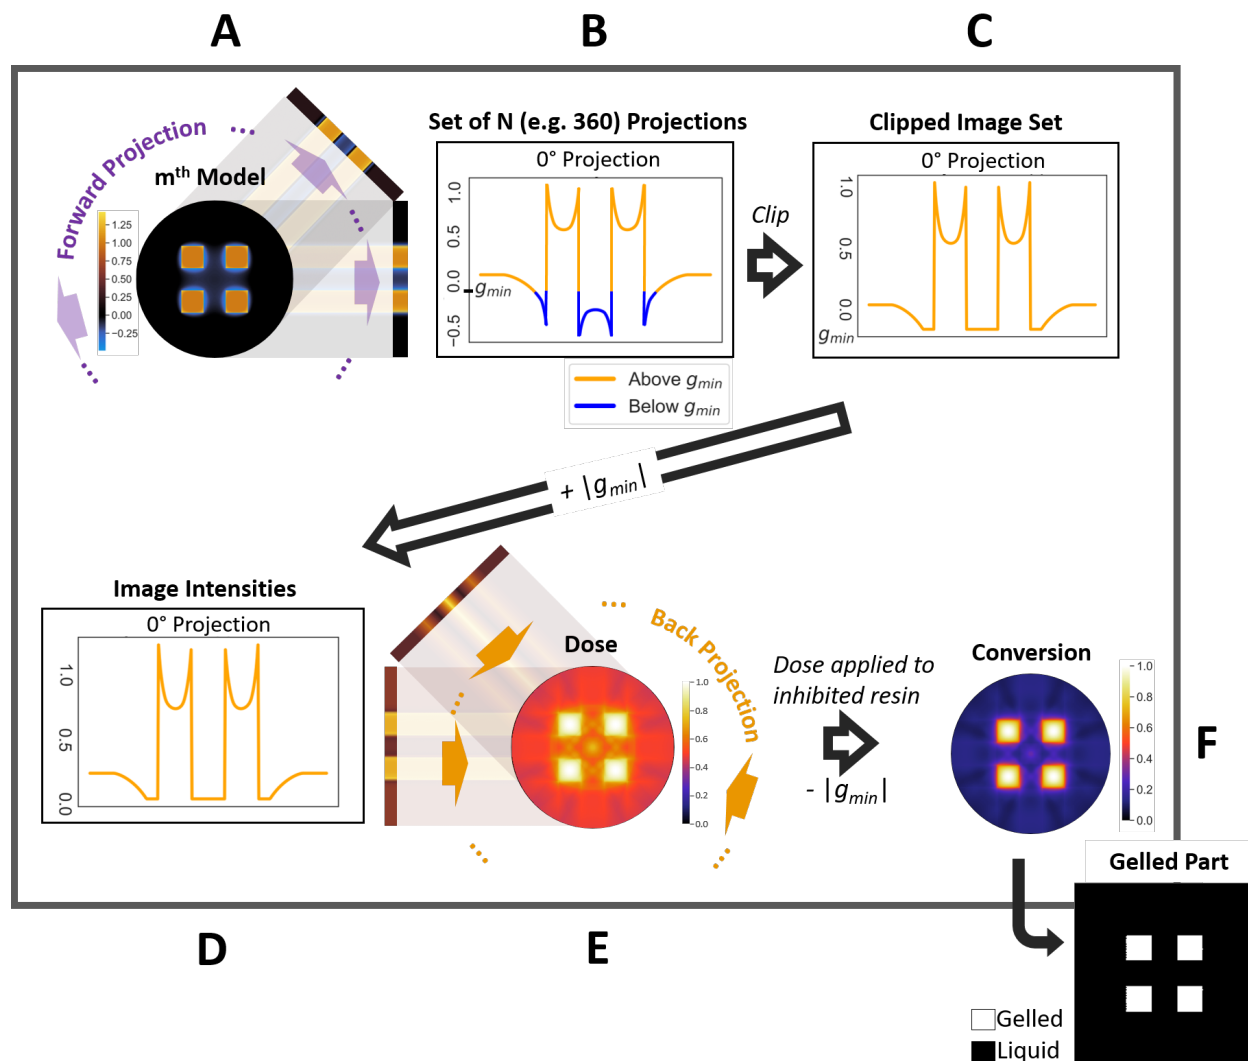


Figure 4.18: Dose reconstruction from m^{th} OSMO model, applied to inhibited resin. This process is used during model optimization, and to compute a final image set. **A:** m^{th} model. **B:** One of N (e.g. 360) forward projections. **C:** Image set is clipped not at zero, but at $g_{\min} < 0$. **D:** Image set is shifted by $+|g_{\min}|$, yielding non-negative intensities. **E:** Image set is back projected, forming a dose reconstruction. **F:** The dose is applied to an inhibited resin. The ratio of $-g_{\min}$ to 1 is the same as the ratio of the inhibition period to the conversion period of the resin. In terms of conversion, when all units are normalized, this has the effect of subtracting a factor of $|g_{\min}|$ from the resultant conversion distribution. For mathematical convenience when optimizing a model, adding $|g_{\min}|$ from step (C) to (D), and subtracting $|g_{\min}|$ at the final step, can be omitted. By the linearity of the Radon Transform, the addition and subtraction of $|g_{\min}|$ cancel each other out (under normalization), yielding the same final conversion distribution.

of this in 2D is shown in Fig. 4.20. Details of this example reconstruction are shown in Fig.

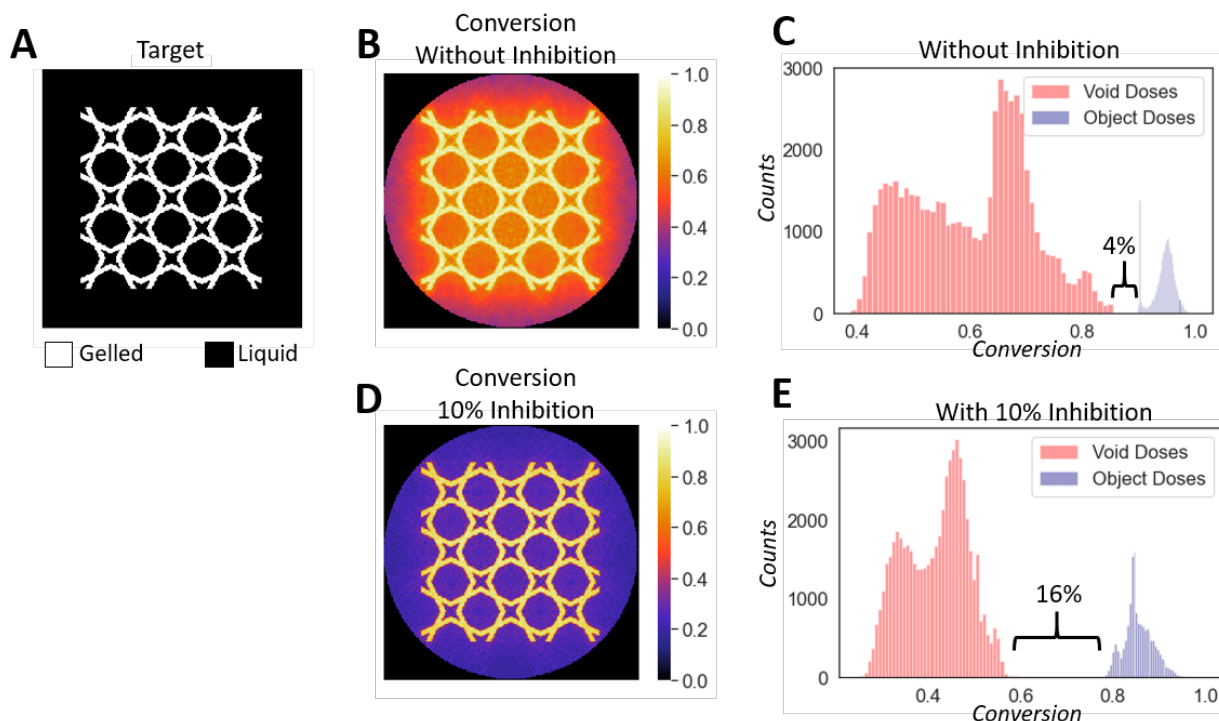


Figure 4.19: The effects of including material inhibition when computing VAM image-sets. **A**: A target geometry. **C**: A conversion profile computed under the assumption of no material inhibition. **C**: The voxel distribution from the reconstruction shown in (B). A 4% process window size allows for some material and optical imperfections, without affecting the final shape of the gelled part. **D**: A conversion profile computed assuming a 10% inhibition period. **E**: The voxel distribution from the reconstruction shown in (D). The process window size is four times larger than what was achieved in (B)/(C) without inhibition, resulting in a reconstruction with four times more tolerance to material and optical errors. A high-contrast reconstruction such as this would be desirable for applications with large optical or material uncertainties, or for applications where the initial-gelled part must be driven to higher degrees of polymeric conversion before removal from remaining resin, and post-cure. Image from [2].

4.21. Note that extending this to 3D would break the independence of slice-regions in a print, but the optimization process would remain algorithmically unchanged. The ability to model Gaussian beam projections becomes especially important when a long depth of focus relative to the maximum target dimension cannot be achieved experimentally, resulting in beams that are significantly non-collimated. This may occur when using LED illumination instead of lower etendue laser-sources [27], or when applying VAM at especially large or micro scales.

Similarly, image computation for printing around occlusions requires that ART line integrals

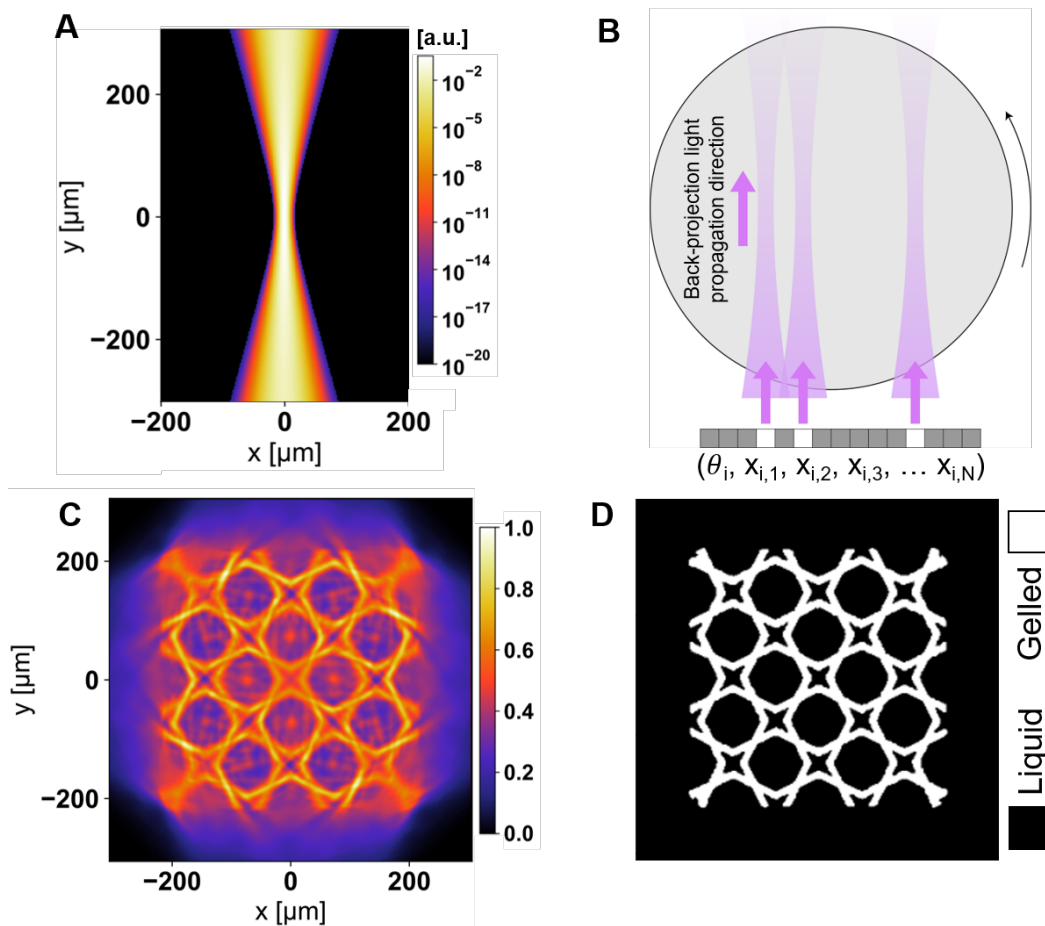


Figure 4.20: **A**: Gaussian point spread function (PSF). **B**: Back projection with Gaussian-pixel-PSF projector. **C**: Optimized reconstruction with $D_l = 0.5$, $D_h = 0.6$, and 150 iterations. Optical parameters: wavelength = 405 nm, numerical aperture = 0.1, digital micromirror device pixel pitch = $10.8 \mu\text{m}$, and magnification = 0.36. **(D)** Thresholded reconstruction, threshold = 0.545. Here, VER was 0.074, IPDR was 0.512, and process window size was -15.5%. Image from [2].

be truncated at the occlusions. This also reduces the freedom of a tomographic reconstruction, however the OSMO algorithm still works without modification. An example of such a reconstruction is shown in Fig. 4.22. Such computation would be necessary for applications involving over-printing around existing, occluding structures. Here, VER was 0.0694, IPDR was 0.605, and process window size was -24.1%. Note: IPDR and process window size are large, compared to a visual inspection of the dose histogram (Fig. 4.23), because a small number of in-part pixels are in error. Further details about this occlusions reconstruction are shown in Fig. 4.23.

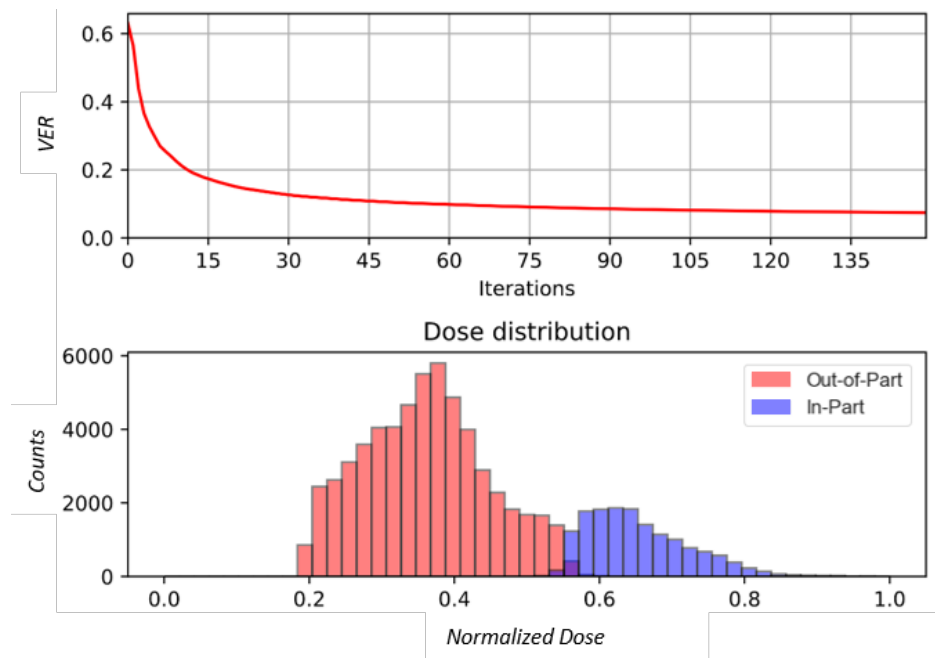


Figure 4.21: Details for the Gaussian beam reconstruction shown in Fig. 4.20. **Top:** Algorithm Convergence, measured in VER. **Bottom:** Histograms of dose for Gaussian beam reconstruction. Image from [2].

Printing around occluding structures is not limited to sparse target geometries. Dense overprinting around occlusions is also possible with the OSMO algorithm. An example of a maximally dense print around occlusions is shown in Fig. 4.24. For this challenging reconstruction, the number of error voxels is quite noticeable. However, this quality of reconstruction could still be useful for many applications, and it further illustrates the flexibility of the OSMO algorithm.

4.1.11.1 Image Computation for Alternate Printing Geometries

In prior VAM literature [26, 27, 28], only a basic tomographic projection geometry was considered; one in which writing beams are normal to the axis of rotation, and are evenly distributed about a rotation. Other printing geometries, however, are possible and may have application advantages. In particular, the standard geometry requires access to the sample from all directions normal to the rotation axis. One simple modification is to use only a limited angular range for printing. Fig. 4.23, A, B shows an example of a reconstruction slice of a thicker geometry. Here, only an




Figure 4.22: **A**: Target and occluding geometries. **C**: Back projection with custom projector; occluded regions are shown as black regions adjacent to the occluding geometries. **C**, **D**: Optimized raw and thresholded reconstructions with $D_l = 0.35$, $D_h = 0.60$, threshold = 0.515, and 120 iterations. Image from [2].

angular range of 130° was used, yet the OSMO algorithm was able to produce a reconstruction with a VER of zero. Another geometry, inspired by tomographic medical imaging and computed laminography, uses a ring of projections that are at a fixed, non-normal angle, to the axis of rotation (see Fig. 4.23. C, D). Although this geometry excludes a cone of available spatial frequencies, it is commonly used for imaging and provides sufficient resolution for many applications [45, 46]. This could allow for VAM printing into flat sample packages, perhaps with applications in microfluidics. In this case, the OSMO algorithm achieved a VER near zero ($1.7\text{E-}4$). However, deviations from the standard geometry come at a cost, particularly as a reduction in process window size. For both the limited angular range reconstruction, as well as the tomographic reconstruction, while a VER of zero or near zero was achieved, the resultant process window sizes were much smaller than what

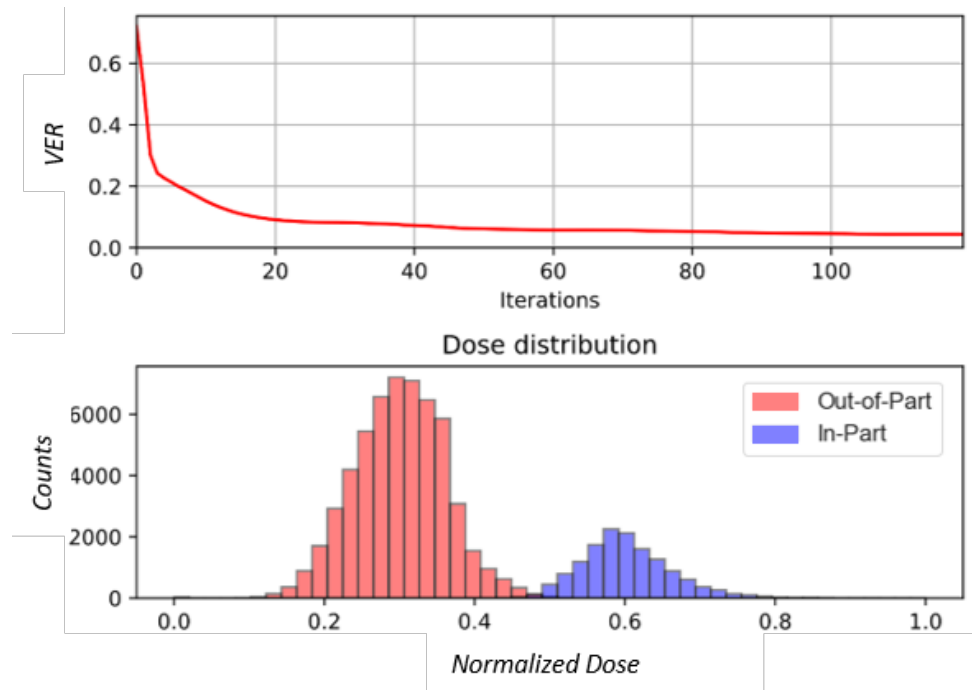


Figure 4.23: Details for the occlusions reconstruction shown in Fig. 4.20. **Top:** Algorithm Convergence, measured in VER. **Bottom:** Histograms of dose for the occlusions reconstruction. Image from [2].

a standard tomographic geometry could achieve. In both cases, this is due to a reduction in the degrees of freedom of the reconstruction process. For the limited angular range case, only 130° (36%) of a full rotation was used. Note that since attenuation is included, a 0° projection is different from a 180° projection, and 130° is 36% of a full rotation image set. If attenuation were not included, a 130° range would comprise 72% of the images in a full rotation. With tomosynthesis, the degrees of freedom of the reconstruction are reduced because the volume can no longer be decoupled into slices addressed by 1D projections, and there is an increased spatial overlap of writing beams. Small deviations from the standard geometry hardly suffer a decrease in reconstruction performance, with larger deviations (the limit being 90 degrees for tomosynthesis) exhibiting worse reconstructions. However, in spite of these fundamental limitations, image sets can be generated for complex, arbitrary target-objects by the OSMO algorithm.

Additional details of the limited angular range and tomosynthetic reconstructions are shown

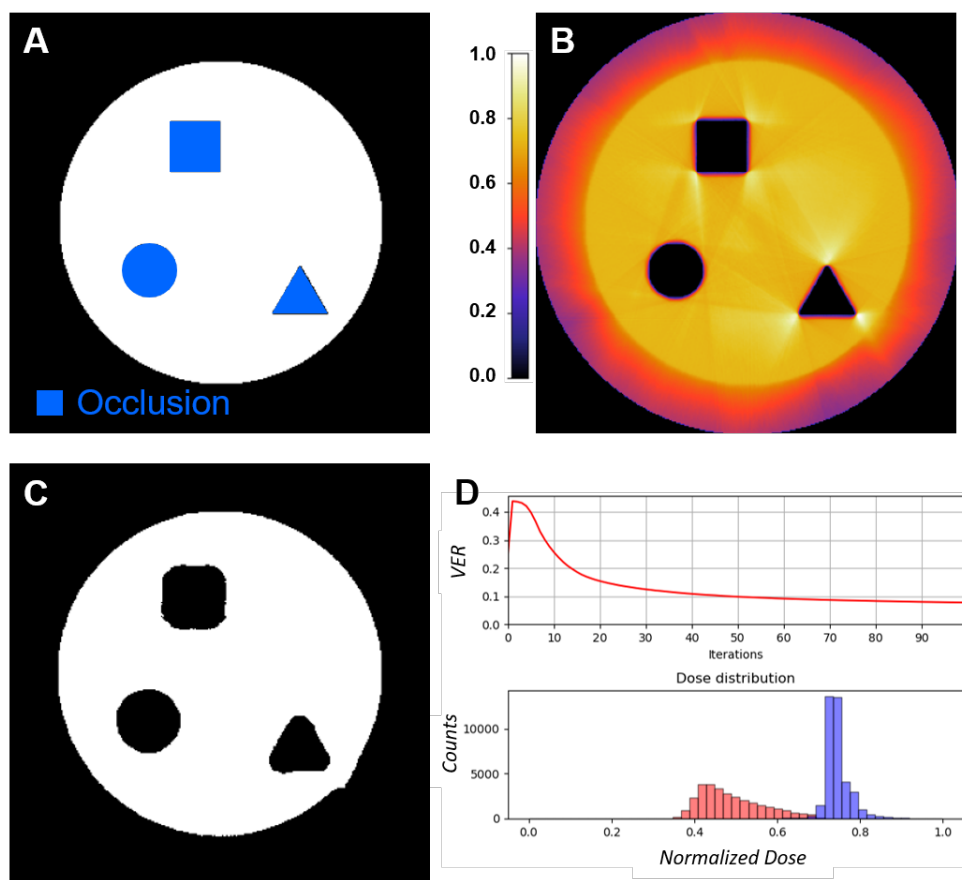


Figure 4.24: **A**: Dense target and occluding geometries. **B and C**: Optimized raw and thresholded reconstructions with $D_l = 0.73$, $D_h = 0.74$, threshold = 0.687. **D**: Algorithm convergence and histograms of dose in B. Here, VER was 0.0756, IPDR was 0.4341, and Process Window was -38.3%. Image from [2].

in Fig. 4.26. Tomosynthetic VAM printing, in particular, could enable new applications of VAM. The flat package geometry is more convenient than cylindrical in many cases and could enable hybrid printing methods such as printing into microfluidic devices. Further, the flat package enables a step-and-repeat scheme such that tomosynthetic VAM could allow for writing into very large areas, one small, high-resolution area at a time. See the next chapter for a more detailed discussion about tomosynthesis for VAM.

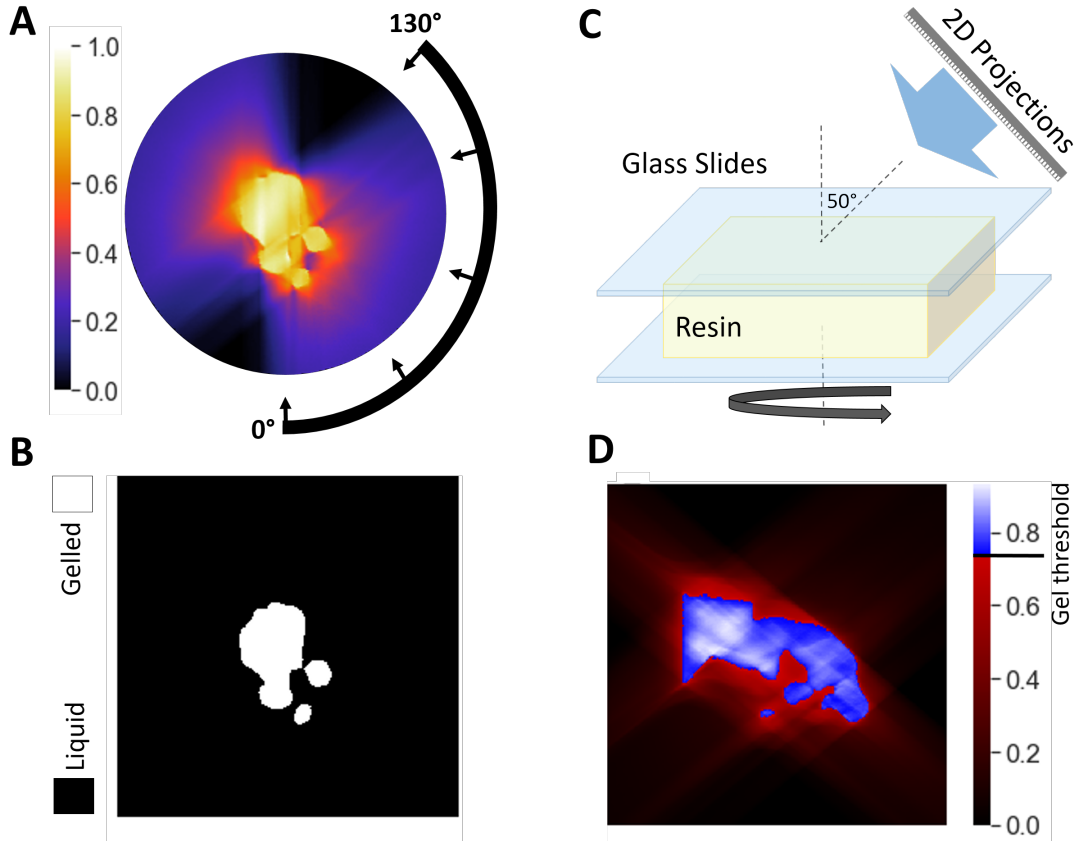


Figure 4.25: **A:** A 2D dose reconstruction yielding a VER of zero, made by a set of projections over only a 130° range, with $D_l = 0.78$ and $D_h = 0.79$. **B:** A gelation threshold applied to the dose shown in (A) yields a perfectly gelled slice of a thicker geometry. **C:** A tomosynthetic printing geometry, with projection direction vertically angled by 40° from the standard printing geometry. **D:** A slice of a 3D thicker-geometry dose-reconstruction with VER near zero ($1.7E-4$), made with the tomosynthesis writing geometry shown in (C), with $D_l = 0.74$ and $D_h = 0.75$. The process window size was -0.70%, and the IPDR was 0.26. Note that (D) shows a 2D slice of a 3D reconstruction, not a 2D reconstruction. In this case, 360 angles, all at a 50° to the axis of rotation, were used. Image from [2].

4.1.12 Experimental Demonstrations of OSMO

Here we present a few examples of VAM prints made with OSMO-computed projection sets. These prints are not used to directly compare OSMO to alternative computation methods - the complexity of washing, swelling, etc. make the shape of optical dose delivery a difficult variable to isolate. Instead, we present a few prints only as an example of the method working for a variety of target geometries.

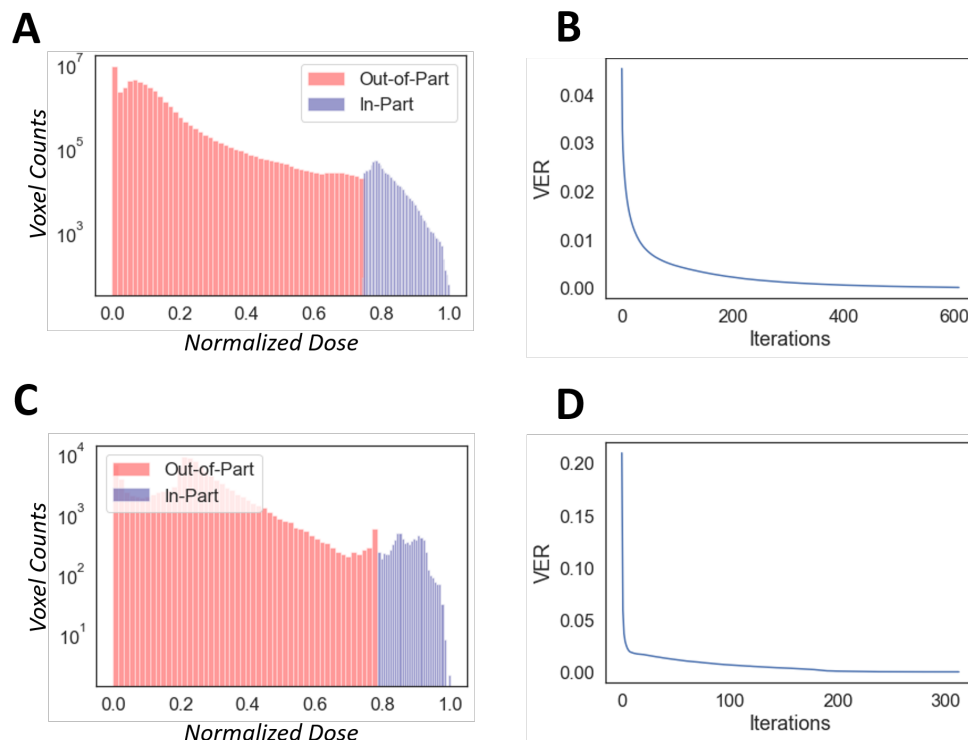


Figure 4.26: Algorithm Convergence and Histograms of Dose for reconstructions shown in Fig. 4.25. **A and B** Dose histograms and algorithm convergence for tomosynthetic reconstruction shown in Fig. 4.25D. Here, the algorithm converged to a VER of $1.7\text{E-}4$. **C and D** Dose histograms and algorithm convergence for limited angular range reconstruction shown in Fig. 4.25A. Here, the algorithm converged to a VER of zero. Image from [2].

One noticeable benefit to printing with OSMO image sets has to do with in situ print timing. We have found that when printing with image sets computed with gradient descent methods, we observe continuous gelation growth via a shadowgram at the end of a print. This makes it difficult to know when to stop the print, as it is unclear which additions to gelation are in-part and which are out-of-part as the print over-exposes. With OSMO, however, we observe an increase in the number of gelled voxels, then no change in the print-shape for a few seconds before more gelation then occurs. This pause in gelation growth is a manifestation of a positive process window, and makes it easy to know when to stop the print. Improved print methods would calibrate the printing process well enough to know the target print time before the start of a print, making manual print-timing unnecessary. However, the process window would still present a tolerance for timing and other

impressions.




Figure 4.27: VAM print using OSMO algorithm, yielding a PW of 5%. The resin was 3:1 BPA-GDA:PEDGA250 with 80 mM of I907. The print time was 48 s with a maximum optical intensity of 233 mW/cm^2 . **A**: SEM image of the print taken by Archish Muralidharan. **B-C**: Photographs of the part. **D-F**: Corresponding views of the target STL file.

Figs. 4.27 and 4.28 show examples of VAM prints of dental aligner geometries made with a OSMO-computed image sets. Note that the flat spot at the top of the image in Fig. 4.27B was from resin remaining after washing that drained down and collected onto the slide upon which the part was sitting; post cure then hardened this small puddle. This exemplifies the significant impact that washing procedure has on a final print, and the resultant difficulty of evaluating a reconstruction algorithm via the final product. Also note the significant striations - similar in appearance to traditional 3D print layering - that affect the optical clarity and surface uniformity of the printed part. These striations are ubiquitous in the field of VAM. In a later chapter we discuss what causes them, and we present a simple and effective method of eliminating them.

Further diversity of target geometries printed by OSMO-computed image sets is shown in Fig. 4.29. Note the large striations in the box structure causing surface ripples. The octet strut exhibits




Figure 4.28: VAM print using OSMO algorithm, yielding a PW of 4.8%. The resin was 3:1 BPA-GDA:PEDGA250 with 80 mM of I907. The print time was 72 s with a maximum optical intensity of 233 mW/cm^2 . **A-B**: Photographs of the part. **C-D**: SEM image of the print taken by Archish Muralidharan (scale bar: 5 mm). **E-H**: Corresponding views of the target STL file.

nonuniform strut widths. This could have been due to non-uniform optics, to diffusion effects, or other system errors. The simulated image set, of course, did not exhibit such nonuniformities. Despite the few percent process window, print shape errors still occurred in this case. Lastly, Fig. 4.30 demonstrates the significant effects of swelling during post-processing, as well as an index-structure print into a holographic photopolymer. Fig. 4.30A shows a projection view of the 3D target. Fig. 4.30D shows a dose reconstruction for the top of the box. Fig. 4.30B shows a part, still immersed in 3:1 BPA-GDA:PEGDA250 resin immediately after print completion; note the undistorted box shape. Fig. 4.30C shows the same part after a 10 min soak in ethanol. All sides of the box swelled and severely distorted. Fig. 4.30E-F show the same box printed as an index structure into a solid holographic photopolymer (see [47] for material details). While these prints indicate the need for additional system and process development and calibration, they do provide the most basic validation of OSMO's efficacy: VAM prints for a variety of target geometries. OSMO was also used for the VAM prints in [5] and in [6] - later work with improved systems, capable of demonstrating the print precision possible with the OSMO approach.




Figure 4.29: VAM prints using OSMO algorithm. The box was printed into a 3:1 BPA-GDA:PEDGA250 resin with 80 mM of I907. The print time was 37 s with a maximum optical intensity of 426 mW/cm^2 . The octet strut was printed into a 3:1 BPA-GDA:PEDGA250 resin with 20 mM of I907. The print time was 184 s with a maximum optical intensity of 647 mW/cm^2 . **A-B**: SEM images of the prints taken by Archish Muralidharan (scale bars: 5 mm). **C-D**: Corresponding views of the target STL file.

4.2 Image Computation for Grayscale-Dose VAM Printing

One of the appealing features inherent to VAM is that it is possible to cure resin with spatial control over the dose delivered, creating functionally graded properties through a printed part. Whether the objective is to control modulus, refractive index, or any other dose-tunable




Figure 4.30: Prints of a hollow box made up of mesh walls. **A**: Projected side view of the target. **B-C**: Print immersed in resin and after a 10 min ethanol wash, respectively - severe swelling and distortion occurred. **D**: A dose reconstruction profile of the top of the box. **E-F**: The same box target printed as in index structure into a 2-stage holographic photopolymer [5] prepared by Gabe Seymour. The printed boxes had 8 mm side lengths.

property, arbitrary 3D property control requires arbitrary 3D dose control. The particular dose values required depend on the property-dose response of a given material, requiring a material calibration step. What we call the **Grayscale VAM** printing problem is to compute the image-set that will reconstruct such an optical dose distribution. Here we present an object space approach to image computation for grayscale VAM. As with binary VAM, we seek to compute a model that will yield a desirable dose reconstruction by the process shown in Fig. 4.4B. This can be done with two steps: first by computing an initial model, and then by adjusting that model to improve the final reconstruction. An initial model will produce a grayscale dose reconstruction that suffers from reduced contrast and from non-uniformities. It can be computed by the same process as for binary VAM: by applying a frequency filter to the object, as detailed in Section 4.1.3.3. Fig. 4.31A-C shows a grayscale target, the initial model produced by filtering the target,

and the dose reconstruction produced by that model. This initial reconstruction (Fig. 4.31C) exhibits reduced contrast, as well as non-uniformities in each region that should be constant in dose. This is unsurprising, as grayscale reconstructions are highly constrained since they attempt to achieve a particular function value at every point in a reconstruction – a task that is generally impossible without negative image intensities, although there do exist functions that can be exactly reconstructed by non-negative image sets.

The initial model can be adjusted to balance reconstruction contrast and uniformity, and to choose which dose values are best reproduced. This is done by shifting the initial model (M) by a constant value (c) before using it to produce a reconstruction (f). That is, using the notation from Section 4.1.3, $f = NP^* \max(0, P(M+c))$. Fig. 4.31D-F shows examples of this. Adding a constant to the initial model (Fig. 4.31B) will shift where, in the reconstruction dynamic range, the object will be reproduced. Adding a negative constant to the model allows for adjustment in the opposite direction, better-matching the smaller values of the target function. This process lets us choose which values from our original target structure are most-accurately reproduced. Larger constants will better distinguish high dose regions, and will improve their uniformity, but at the cost of an overall reduction in contrast. This process is fast, making it computationally practical to check many shift values to find a desirable reconstruction for a given target. While this method of filtering and shifting does not control material properties perfectly, it allows for adjustable approximation, opening the doors to the application of VAM to the fabrication of functionally graded structures.

An application example of this method is the tomographic printing of a gradient index (GRIN) lens, hyperbolic secant radial profile (Fig. 4.32A). Suppose we wish to fabricate a flexible GRIN lens for an endoscopic application. Suppose it should be much longer than its diameter; e.g. we wish to fabricate a 9 mm long by 1 mm diameter lens. One approach is to pattern a holographic photopolymer whose refractive index increases with applied optical dose (while it is sensitive - a thermal cure step stops further change). We have found that projecting the desired GRIN profile through the length of the sample does not work well. The highly directional, long path-length exposure results in haze throughout the sample. An alternative approach is to tomographically

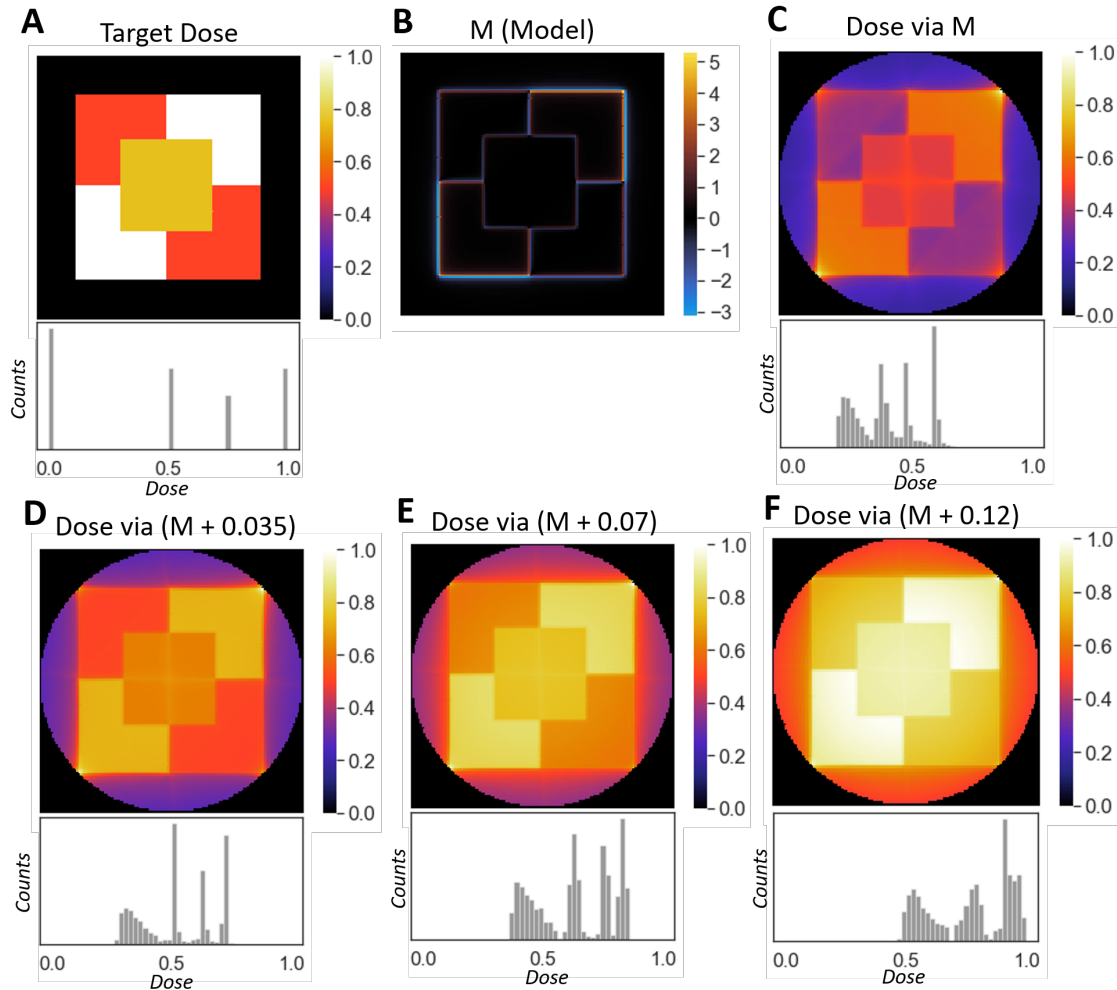


Figure 4.31: Grayscale tomographic reconstructions by non-negative projection sets. **A**: A target dose to reconstruct and its histogram. **B**: A model, M , formed by frequency-filtering the target. **C**: The reconstruction and histogram yielded from the model. **D-F**: Reconstructions and their histograms yielded by adding a constant to the model. The more the model is shifted, the higher the dose values that reconstruct the object in the necessarily reduced-contrast reconstructions. There is a trade-off between uniformity and contrast. E.g. the average spread of each dose-region in the reconstruction in (F) is smaller than those in (C-E). However, the reconstructions shown in (C-E) exhibit a larger dose contrast than the reconstruction shown in (F). Image from [2].

pattern the material from the side, though its diameter. Fig. 4.32 shows how the method presented in this section can be used to tomographically reconstruct an approximation to a GRIN lens profile.

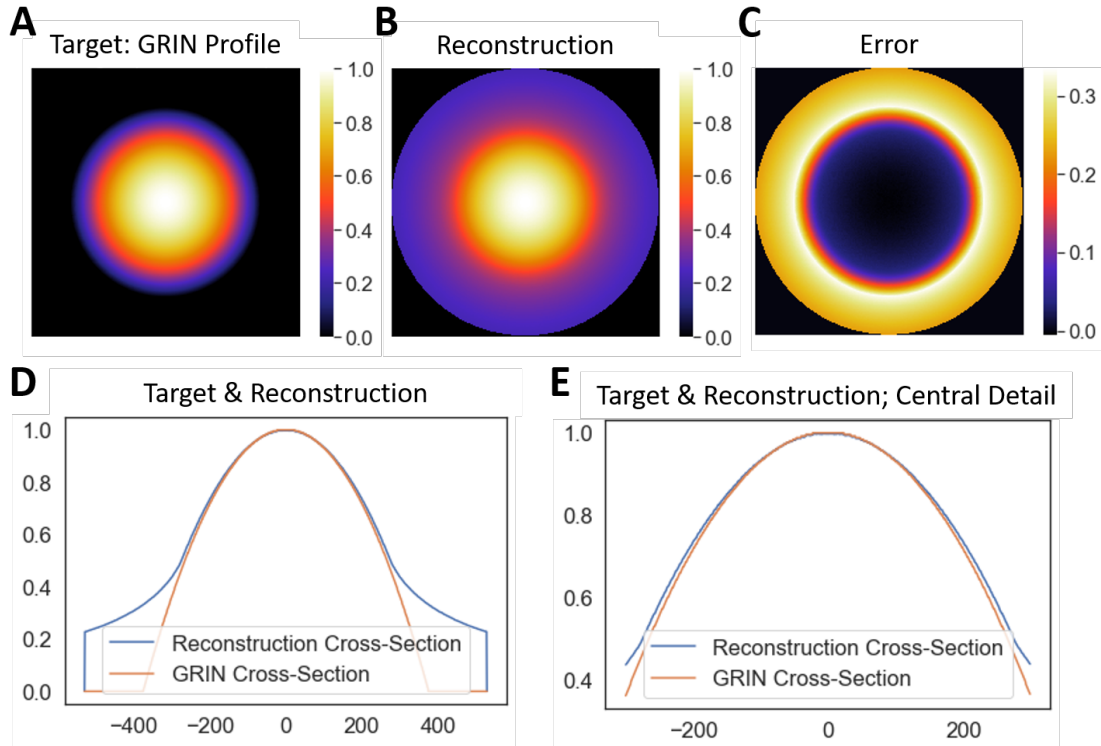


Figure 4.32: Example of reconstructing a gradient index (GRIN) lens profile via the object space method presented in this section. **A**: A target GRIN profile of the form $sech(gr)$, where g is a constant and r a radial coordinate. **B**: The reconstruction resulting from object-space method described in this section, with no constant shift added. **C**: The difference between the target and the reconstruction. **D**: Cross sections through the centers of the target and the reconstruction. **E**: The central portion of (D). Even though the approximation is poor at the outer radius of the reconstruction, it is much better in the central 50% of the print. The central region could be used to print a good approximation to a GRIN lens.

4.2.0.1 Image computation for the control of material properties

VAM is naturally suited to the problem of 3D printing structures with functionally graded properties since it has the ability to control optical dose throughout a print. However, tomographic computation addresses only the control and delivery of optical exposure. Here, we briefly discuss how the grayscale reconstruction algorithm described in Section 4.2 fits into the larger problem of material property control in VAM.

The problem of arbitrary property control by VAM requires multiple inversion steps. For material property being a function, f , of conversion, conversion being a function, q , of optical

dose, and optical dose being a function, h , of projection image set, we can say that an arbitrary 3D distribution of material property, $MP = f(q(h(imageset)))$. Note that f and q are typically intensity-dependent, requiring system-specific calibration. To compute the image set that results in a desired distribution of material properties, we must solve the inverse problem: $imageset = h^{-1}(q^{-1}(f^{-1}(MP)))$. Inverting f and q is an experimental measurement calibration step yielding monotonic functions that are empirically collected [39]. Predicting and modeling such material property dependence on conversion, and the dependence of conversion on optical exposure is an active area of research [48, 49] and beyond the scope of this discussion. However, once f and q are inverted, they provide a grayscale optical exposure target h (e.g. like the target shown in Fig. 4.31A). Inverting h can then be approximated by the grayscale reconstruction algorithm described in Section 4.2. The limitations of such a grayscale reconstruction will strongly depend on the particular target distribution, along with the particular material responses involved. Finally, the effects of inhibition can be included to increase reconstruction degrees of freedom. This is done, as in the binary-gelation printing case, by allowing negative values in the projection sets, the magnitude of which is determined by the inhibition period of the resin. While the limited fidelity of the grayscale algorithm constrains the ability to specify arbitrary material properties in 3D, the processed outlined here may provide sufficient control for many applications, and demonstrates a computational advancement for more potential control over VAM prints.

4.3 FBP Enabled by Inhibition

While the grayscale reconstruction algorithm presented in the previous section approximates grayscale control over conversion, its control is limited - it is an approximation, after all. What it approximates is a theoretically exact grayscale reconstruction algorithm such as FBP or, equivalently, BPF. The best and most arbitrary control would be achieved if we could directly apply FBP to the problem of VAM tomographic reconstruction. However, as we discussed at the beginning of this chapter, the lack of negative intensity in projected images precludes the direct application of FBP for the tomographic reconstruction of arbitrary distributions of integrated optical intensity.

This is because applying FBP’s Ram-Lak filter [40] to an image set results in negative image set values. However, it is not fundamentally optical dose that we care about - it is material conversion. This gives us an additional degree of freedom - material response - to solve the problem of reconstructing *material conversion* via optical tomography.

Because back projection is a linear operator, it is possible to add a constant to the image set, making it non-negative, and then to compensate for this shift by subtracting a constant from the resultant reconstruction. This would yield a mathematically equivalent reconstruction to the desired FBP reconstruction. The lack of negative intensity makes it impossible to subtract optical dose from a reconstruction. However, we do not need to subtract optical dose, we need to subtract material conversion.

Instead of thinking of DMD patterns as potential contributions to optical dose over time, let us consider them as potential contributions to material conversion over time, with intensity just an intermediate step. Then, the linearity of the back projection operator tells us that we can exactly compensate for a uniform increase in image-space DMD pixel values with a uniform decrease in object-space conversion values. Decreasing to zero what would have been a positive value of conversion can be achieved with the inclusion of a radical inhibitor in the print resin. This is the same concept presented in Section 4.1.9. However here, instead of adding a small period of inhibition to an OSMO reconstruction, we seek to add a sufficiently large inhibition period to enable the application of FBP. In this section, we discuss the theoretical details of this approach, how FBP could be used for both binary and grayscale VAM, and challenges inherent to the method.

4.3.1 Compensating an Image-Space Shift in Object-Space

Let us take a closer look at how a constant shift propagates through a back projection operator. Consider a 2D target function f that we wish to tomographically reconstruct. Let P be a forward projection operator, and P^* a back projection operator. The FBP algorithm gives us that $f = P^*g_R$, where g_R is the filtered image set computed by applying a Ram-Lak filter to Pf .

Suppose that the values of g_R are bounded below by $-m$, and above by M , where m and M

are positive constants. That is, $g_R \in [-m, M]$. Since our DMD cannot project the negative values of g_R , let us add the constant m to the image set: $g_{R+} \equiv g_R + m$, where $g_{R+} \in [0, m + M]$.

Let N_θ be the number of projection angles used in a back projection. Since P^* is linear, for any image sets g_1 and g_2 ,

$$P^*(g_1 + g_2) = P^*(g_1) + P^*(g_2)$$

$$P^*(g + m) = P^*(g) + P^*(m)$$

$$P^*(g_{R+}) = f + mN_\theta$$

This suggests that for the final material conversion to be equal to f , as desired, the inhibition period must cover the integrated intensity applied by the value mN_θ . Note that both the shape and magnitude of f is a function of N_θ . For the case of grayscale VAM in which a perfect reconstruction is desired - not just values above and below and threshold - the full inhibition corresponding to the dose delivered by mN_θ is required. Crucially, the fact that the reconstruction dose $f + mN_\theta$ is applied to locally deplete inhibitor means that $f + mN_\theta$ is implicitly no longer unitless and can no longer arbitrarily scale to any desired dose via increased print times; its range is capped by the dose to clear inhibitor.

An example of this is illustrated in Fig. 4.33. Here, the target is the binary structure f shown in Fig. 4.5A. Upon forward projecting f and applying a Ram-Lak filter, the image set g has a range of about $-m = -0.05$ to $M = 0.05$. Adding m to g , then back-projecting with $N_\theta = 360$ images results in an added constant of $mN_\theta = 18$ in object space. In arbitrary units, the background dose is 18, and the in-part dose is 19.25, or 7% higher than the background. Applying this dose to a resin with only a small oxygen inhibition period, and without added radical inhibitor, results in a very small conversion contrast, as illustrated in Fig. 4.33C. If instead, we apply this dose distribution to a resin with a longer inhibition period, the out-of-part dose that can be applied without causing conversion becomes higher. The absolute dose contrast then also increases, since the difference between in-part and out-of-part dose is equal to 7% of the out-of-part dose. For example, suppose the dose to clear inhibition is 18 [a.u.], as shown in Fig. 4.5D. Then the in-part dose will be 19.25.

Whether or not the resultant conversion is sufficient to cure the print and to avoid system errors depends on system precision and on the conversion/dose slope of the particular resin in question.

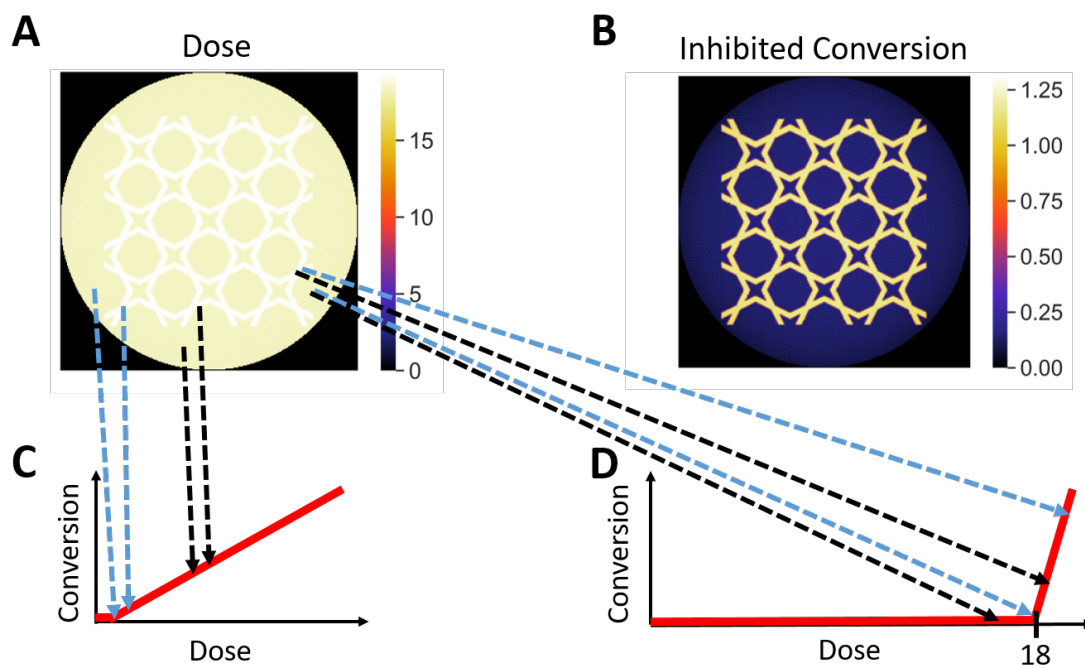


Figure 4.33: Example of a reconstruction made via FBP, expect with the image set shifted to be non-negative. **A**: The reconstruction formed by a filtered image set that was shifted to have a minimum value of zero. In arbitrary units, the background dose is 18, and the in-part dose is 19.25 (7% higher than the background). **B**: A conversion profile that would be produced if the dose in (A) were applied to a resin with a sufficiently large inhibition period, followed by a linear conversion vs dose response. Such a conversion curve is sketched in (D). **C**: Sketch of a resin response without a radical inhibitor added. Oxygen adds a very small inhibition period, so applying the dose in (A) results in a very small conversion contrast. That is, if the intensity profile in (A) is applied such that the oxygen in the out-of-part regions is exactly depleted, the total dose applied will be small since the oxygen inhibition period is small. Since the value 7% larger than this small dose is small, the in-part regions are driven only a small way up the conversion curve. This is shown with the blue arrows. The black arrows show an alternative use for the dose. **D**: An inhibition period equal to a dose of 18 results in a conversion profile (B) with a background of approximately zero, and an in-part conversion corresponding to a dose of 1.25. Since the in-part dose is 7% greater than the background dose, further increasing the inhibition period would allow for higher in-part conversion while keeping the background conversion at zero. More inhibition would allow for an further increased in-part conversion. The blue arrows show an example conversion range, likely most applicable as endpoints to grayscale reconstruction's dynamic range. The black arrows shown an example with a slightly shorter exposure time. This might be used for a binary print in which we want to ensure no out-of-part gelation.

More generally, suppose that a shifted-FBP binary reconstruction (as described in this sec-

tion) has an out-of-part to in-part dose ratio of $\alpha : \beta$. For a grayscale reconstruction, α and β are the minimum and maximum of the reconstruction, respectively. For simplicity, let us ignore material nonreciprocity, and assume that conversion C is a function of dose D ; $C = \text{conv}(D)$. Suppose that for any given amount of inhibition, the dose required to clear that inhibition and to reach the start of conversion is D_{inhib} . Suppose that we wish for the maximum printed conversion to be $C_{\text{max}} = \text{conv}(D_{\text{inhib}} + Y)$, where Y is the dose beyond D_{inhib} required to achieve the conversion C_{max} . Then since $\text{Dose}_{\text{inhib}} : Y = \alpha : \beta$, the inhibition dose D_{inhib} required to achieve a desired printed conversion of C_{max} is equal to $Y\alpha/\beta$. That is, α/β determines how much inhibitor we must add to our resin to achieve a desired degree of conversion with an FBP print. This raises two important questions: **1.** What affects α/β , given a target function? **2.** Given that, for any target function, we can compute the inhibition required for shifted FBP to theoretically print the function, can FBP practically be used to print any function?

The answer to the first of these question is unknown. While further research might produce a method to predict α/β for a given target, it is quick and easy to simply computed a shifted FBP reconstruction to evaluate the resultant contrast. Doing so includes both the fundamental nature of shifted FBP for a given target, as well as degradation occurring from numerical approximation. For example, imperfect projector interpolation or simply non-continuous spatial and angular sampling can result in particularly large image set amplitudes after the Ram-Lak filtering step. If the image set then contains an artificially small value - smaller than would occur under perfect numerics - an extra shift will be applied to the image set to make it non-negative, and thus the final dose reconstruction will be reduced in contrast. Further work is needed to understand exactly how FBP discrete approximations affect the reconstruction contrast, but it is clear that they do. An example of this is shown in Fig. 4.34.

The answer to the second question has to do with tolerances throughout a given printing system. For example, consider the target function and the shifted FBP reconstruction shown in Fig. 4.35. Here, the dose contrast is terrible - the in-part dose is only about 0.02% greater than the out-of-part dose. It is theoretically possible to add a large inhibition period (Fig. 4.35F) to print

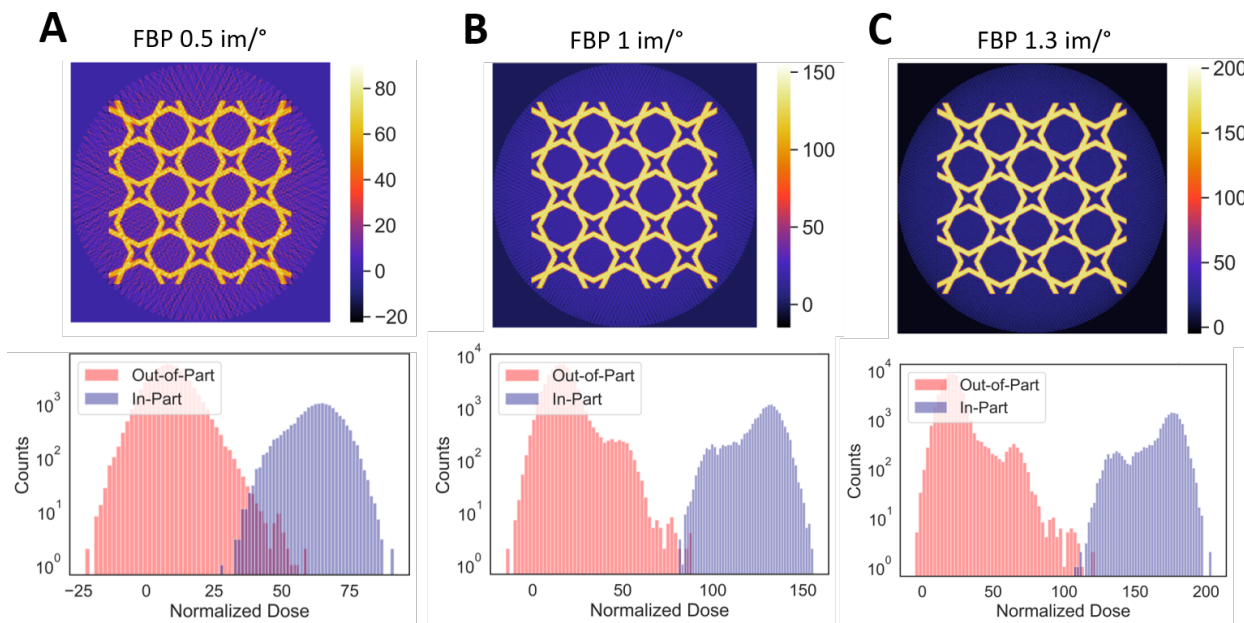


Figure 4.34: Example showing how discretely sampled FBP reconstructions vary with angular sampling density. This is relevant to VAM since spatial resolution is finite (projected beams are not infinitely small), as is angular resolution. In the continuous limit, a reconstruction would have two voxel values: zero and one (after normalization). Instead, we see broad distributions of voxel values for both in-part and out-of-part voxels. **A, B, C**: FBP reconstructions of a binary target, along with dose histograms using 0.5, 1, and 1.3 images per azimuthal degree, respectively. Further increasing angular resolution did little to improve the result. This is likely because spatial resolution - and not angular resolution - is the limiting factor to reconstruction quality. Unfortunately, while it is conceptually easy to increase angular resolution in VAM, spatial resolution is fundamentally optically limited. Any such deviations from theory could have some effect on value ranges of a shifted FBP image set.

with a desired conversion contrast. However, the optical uniformity of the printer, the sensitivity of the resin, etc. must all fall within a tolerance of about 0.02%. For existing VAM printers, this is not feasible. Additional problems exist with adding very large inhibition periods. One problem is that radical inhibitors such as TEMPO are optically absorptive - adding a high concentration could excessively increase the absorption of the resin. Another problem is that the inhibition period contributes to the total print time. A very large inhibition period could cause very long print times, potentially causing excess diffusion of various species during printing. For all these reasons, using radical inhibition to enable shifted FBP VAM reconstructions does not enable the arbitrary printing of all possible parts. However, it seems to be a promising technique for many binary structures,

and it is the only technique for exact grayscale printing.

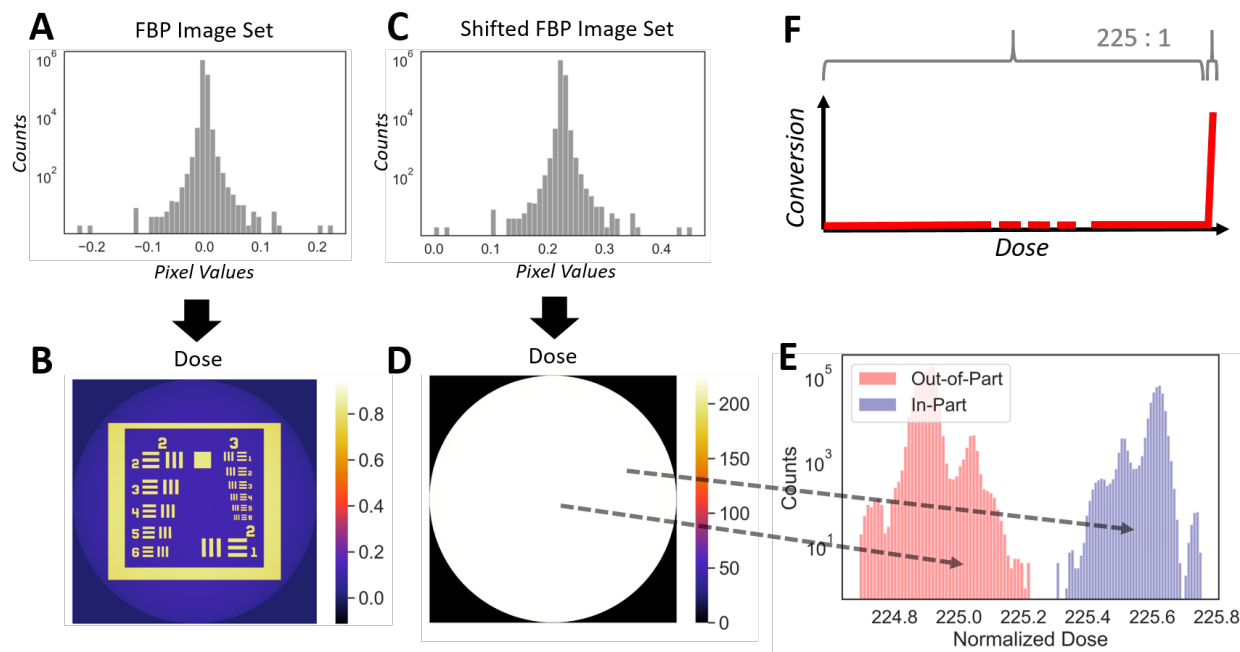


Figure 4.35: Example of the very low dose contrast that results from using shifted FBP on a particular target geometry: a resolution target. **A**: Image set histogram from FBP, producing **B**, a dose reconstruction. **C**: Histogram of the shifted FBP image set; note that non-negative value range. **D**: Shifted FBP dose reconstruction from (C), with dose histograms shown in **E**. Note the very low dose contrast. **F**: Sketch of a hypothetical inhibited resin. If we wanted an in-part conversion that would be produced by a dose value of 1 without the presence of inhibition, we would need the inhibition period to span 225 units of dose. That is, since the fractional difference between the in and out-of-part doses (E) is about 225:1, then so must be the ratio of the dose to clear inhibition and the dose to affect conversion.

4.3.2 Shifted FBP Versus OSMO for Binary VAM

Suppose this shifted FBP method is practical for a given binary structure that we wish to print. Then is it better to use inhibited shifted FBP, or is it better to use OSMO? To answer this question we must better define the goals of our binary print.

Suppose we wish to maximize the conversion process window. That is, we want to maximize the difference between the largest out-of-part conversion and the smallest in-part conversion. In this case, OSMO would be a much better choice than shifted FBP, given the same inhibition applied

in both cases. While we would expect FBP to produce binary reconstructions with perfect contrast (only 0 and 1, in normalized units), as shown in Fig. 4.34, this does not actually occur under non-ideal conditions. Here, a negative process window results, even with a target function for which OSMO produces a large positive process window. This is because FBP assumes continuous sampling and perfect interpolation, and any degradation that occurs when these conditions are not met is ignored. The OSMO algorithm, on the other hand, seeks to improve reconstructions regardless of how they were made. When OSMO improves an image set, it evaluates not only the aspects of a reconstruction that an image set would theoretically produce under ideal circumstances, but it also evaluates the effects of limited sampling resolution, projector interpolation methods, and whatever else produced the reconstruction. It is thus not surprising that OSMO would produce higher contrast reconstructions than FBP. Thus, whether we are trying to leave room for system error, attempting to drive in-part conversion as high as possible without causing out-of-part gelation, or any other application where a large dose process window is useful, OSMO is a better choice than FBP.

FBP, on the other hand, has a few potential advantages for binary printing. FBP produces a relatively instantaneous image-set calculation, and it makes no explicit deviation from the target function. If uniformity of in-part and out-of-part values is a priority, it will likely perform better than OSMO. However, OSMO allows for control over IPDR, and out-of-part uniformity is typically not a metric of concern. Where FBP (and equivalent) performs uniquely well is in its attempt to exactly reconstruct a target function; a necessary constraint for arbitrary grayscale control over print conversion.

4.4 Future Work

Here we briefly discuss future modeling work, algorithm additions, and potential research directions.

4.4.1 Relaxed Voxel Constraints

The basic goal of any binary printing algorithm is to deliver dose that is above a gelation dose-threshold for all in-part voxels, while delivering dose that is below that threshold for all out-of-part voxels. This requirement, along with the required process window, the projection geometry, the number of projection images and their resolution, etc., determines how constrained the reconstruction problem is. For some target geometries, this may result in an impossible problem, and a constraint such as desired process window size must be relaxed. It is thus desirable to find ways to relax the reconstruction requirements that do not impact print quality for a given application.

All of the preceding binary reconstructions have been, in fact, over-constrained. For binary printing, it is not actually necessary to avoid all out-of-part gelation; we are only required to avoid out-of-part gelation that interferes with gelling prints in the desired geometry, and removing them for post processing. By removing dose requirements on a subset of out-of-part voxels, we can relax reconstruction constraints, hopefully resulting in improvements to reconstruction quality metrics.

Allegorically, removing constraints on out-of-part voxels is trivial; we can simply exclude them from the list of out-of-part voxels addressed by e.g. OSMO. The interesting research question is how best to choose which out-of-part voxels to ignore. Clearly, the voxels immediately adjacent to in-part voxels should not be ignored, since if they gel, they will be connected to the print, thus changing the print-shape. Furthermore, it is important that any out-of-part print be separable from the desired in-part print. One simple approach to this would be to enforce the dose of out-of-part voxels adjacent to in-part voxels, and to ensure that any out-of-part gelled regions are smaller in their maximum dimension than the size of the smallest gaps in the desired print. For example, if we are trying to print an empty cage, we must be sure that any gelled regions inside the cage are small enough to fit between the bars of the case, and thus be separated from the print. That is, if we define allowable out-of-part gelation regions to be a collection of sufficiently small, disconnected beads, we can be confident that they will be separable from the part (assuming the part has some

gaps). However, the constraint that gelled regions be smaller in all dimensions than the gaps between the bars would be unnecessary in this case. In this example, a long, thin gelled region could also be separated from the print.

Similarly, it may be possible to relax constraints on in-part voxels. Suppose that in a desired print region, a continuous shell of voxels gel, while the interior voxels do not. That shell would encapsulate the ungelled resin, holding it in place until a flood-exposure post-processing step cures it to gelation. Thus, it may not be necessary to specify that all in-part voxels receive a gelation dose.

An interesting and likely fruitful area of research would be to explore what print metric improvements are possible via relaxed voxel constraints, and which voxels have the largest impact when relaxed. If such relaxation is helpful, then the next step would be to develop an algorithm for choosing the voxels to relax that maximize print-quality improvement while still allowing separation between the desired print and incidental out-of-part gelation.

4.4.2 Accounting for Nonreciprocity

In Section 4.2 we briefly discussed how in order to achieve grayscale control over material properties, it would be necessary to account for material nonreciprocity - where the effective dose with respect to material conversion is not simply the product of optical intensity and time. Actually doing this is not an easy task, and is an area of future work that will be important to both grayscale and binary VAM.

It would be unsurprising if material nonreciprocity was a significant contributor to the print errors we observe in binary VAM prints. This response can be approximately modeled relatively easily to understand the impact of the effect. Suppose that we assume that a fractional exponent applied to intensity approximates the effective dose upon multiplication with exposure time (as in [39]). Since projection images have units of optical intensity, we can apply such an exponent to a projection set, to simulate the resultant conversion via back projecting the images.

Fig. 4.36 shows an example simulation of the effects of nonreciprocity. (A) and (B) show

an OSMO reconstruction resulting in a positive process window of 1.6%. (C) and (D) show the effective-dose reconstruction assuming that the effective dose to which material responds linearly is defined as $t \times I^{0.86}$ [39], where t is time and I is optical intensity. Since image sets are in units of intensity, this was simulated by simply applying the 0.86 exponent to the to image set computed for (A), then back projecting and normalizing. There are a few interesting effects to note. First, the reconstruction shown in (C) has a brighter background out-of-part dose than that in (A). This is because an exponent between 0 and 1 emphasizes lower-value projection contributions. Second, we went from a positive process window to a negative one, with a huge VER of 17%! This is a large error that dominates the choice of reconstruction algorithm, slight timing errors, etc. It is unknown if the effective dose model used for this simulation is at all accurate for our materials. The effect is likely reduced when much of the print occurs in the material inhibition period where the exponent on intensity is one. However, a nonlinear intensity response could be significantly affecting part of the printing process, and it highlights the potential importance of accounting for reciprocity failure.

There are a few computational approaches to compensate for this effect. The solution is computationally simple if we assume that the nonreciprocal of dose on material conversion follows superposition. For example, suppose that one exposure imparts an effective dose of $\Delta t \times I_{\text{effective}} = \Delta t \times 0.1$, and a subsequent exposure imparts an effective dose of $\Delta t \times 0.4$. Under the assumption of superposition, these exposures will affect a material conversion equivalent to a single exposure with effective dose $\Delta t \times (0.1 + 0.4)$. If this, indeed, holds, then we simply need to invert a computed image set as per an effective dose exponent, a more complex model, or even simply empirically collected data. This was the approach outlined in Section 4.2, and the assumption of dose superposition is likely a good approximation.

Suppose, however, that dose superposition does not hold. Returning to our example, that would mean that an effective dose of $\Delta t \times 0.1$ + an effective dose of $\Delta t \times 0.4$ does not yield the same conversion effect of a single effective dose of $\Delta t \times (0.1 + 0.4)$. This could occur if polymerization kinetics, and the polymer chains formed, are significantly different at different intensities. In

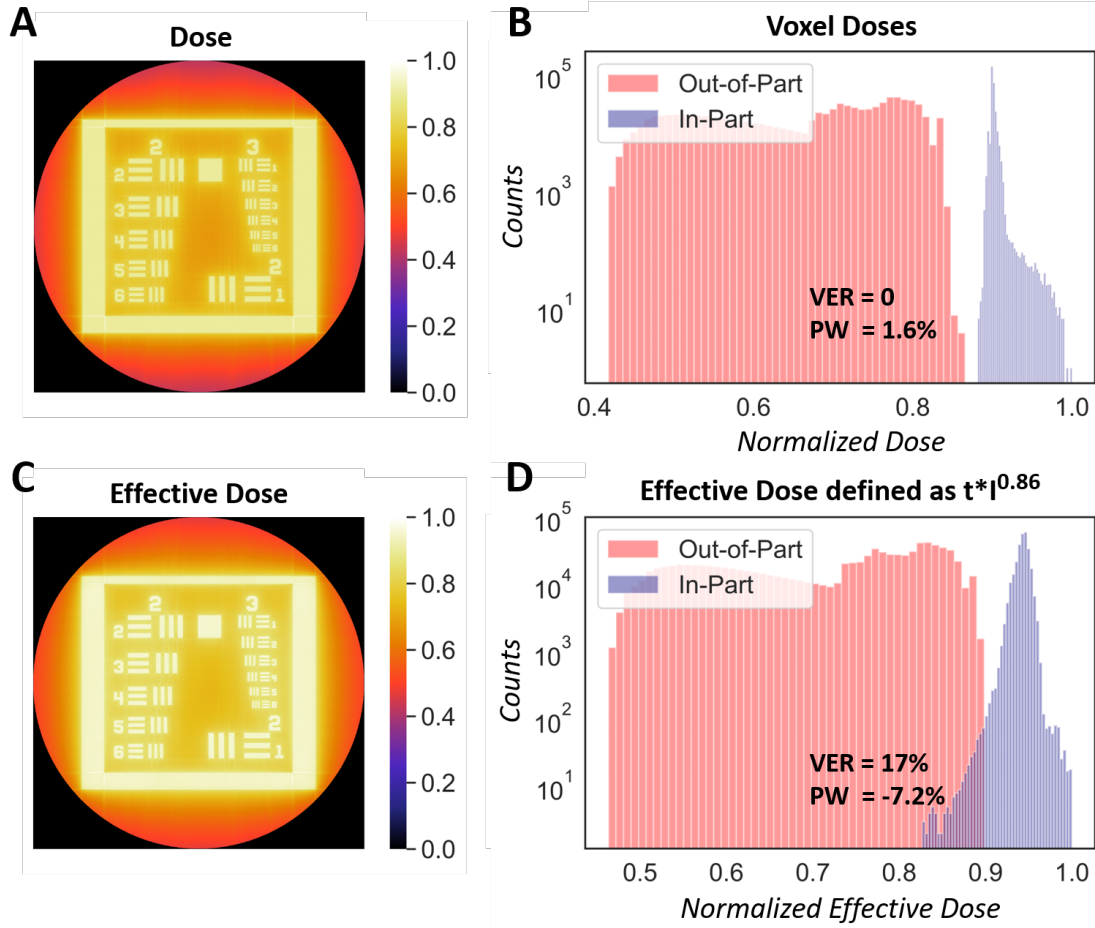


Figure 4.36: Simulated reconstruction showing the impact of nonreciprocity, assuming a simple model for effective dose: $t \times I^{0.86}$. **A-B**: OSMO dose reconstruction with a process window of 1.6%. Suppose we assume that the resin responds linearly to an effective dose of $t \times I^{0.86}$, where t is time and I is optical intensity. Applying this scaling exponent to the project images, then back projecting, results in a reconstruction of effective dose, which will be proportional to resin response. Applying this to the image set that generated (A-B), we get **C-D**, the simulated effect of material nonreciprocity. A negative process window occurs, along with a large VER.

this case, we have an additional layer of tomographic optimization. Depending on how different effective doses add, with respect to material conversion, this problem could be easier or harder than a problem assuming the linearity of back projections. An approach to this problem that is likely tractable, is to include a model of material response upon each OSMO iteration. That is, for each OSMO dose reconstruction, we compute the expected material response as a function of superimposed optical intensities. If needed, we could even simulate these intensities sequentially,

to simulate the effect of subsequent exposures over the the print time. This could include an exponent on intensity that changes during the print, as different material-response regimes are reached, such as the inhibition, pre-gelation, and post-gelation periods. The resultant computed conversion distribution would then serve as the output of a given OSMO model, and would then be used to update the subsequent OSMO model. In this scheme, OSMO would optimize for material conversion instead of only optical dose.

Likely the most difficult part of accounting for material nonreciprocity lies in the experimental process of measuring material response. Once, e.g., FTIR is used to answer the question of whether effective dose obeys superposition, many careful exposures and conversion measurements will be necessary, with fewer measurements necessary if this superposition is, in fact, obeyed. A simpler version of this would involve measuring only the dose to gelation as a function of intensity. This could even be done in a single print: for each of some set of positions along the axis of rotation, a particular intensity (DMD gray level) could be continuously applied. The times to gelation for each gray level could then be recorded.

4.4.3 Accounting for Diffusion

An interesting area of future work in VAM computation is to predict and compensate for diffusion. Multiple species can diffuse during VAM printing: photoinitiator, radicals, inhibitor molecules, oligomers prior to gelation, etc. We have observed that long print times (e.g. greater than 5 min for our 3:1 BPA-GDA:PEGDA250 resin) tend to result in print-shape degradation, manifesting as rounded corners, edges, etc. For low print powers and low initiator concentrations, we have observed prints that we expected would take 10 min fail to gel even after 30 min of print-time. This suggests an out-diffusion of oligomer at a rate high enough that gelation is never able to occur. Computational modifications will never make this situation good for accurate VAM printing, but we may be able to improve print quality for less extreme cases of diffusion.

Suppose that we have a model for diffusion that we can apply to a simulated distribution of conversion - e.g. first order Fickian diffusion. That is, suppose we can model diffusion given some

initial distribution of molecules. This leaves us with two problems to solve. **1.** How to apply this to model a VAM printing process in a way that is sufficiently accurate, without being prohibitively computationally expensive. **2.** Given such a simulation of diffusion in VAM printing, we are left with the inverse problem of choosing the optical exposures to deliver that, after diffusion, will leave us with a desired conversion distribution. These two problems are coupled by our general requirement to limit computational cost.

Let us first consider the second problem, as it informs complexity limits on the first. One approach would be to include diffusion in each OSMO iteration. That is, after back projecting intensities, we compute the final conversion distribution that would result from diffusion, then that becomes the output upon which OSMO then iterates. In this case a diffusion calculation would be performed with each OSMO iteration, suggesting that computationally inexpensive diffusion model be applied. An alternate approximation would be to perform an optical dose OSMO optimization as usual, then apply diffusion to the final output. This could then stand alone as its own inverse problem: computing the dose distribution that compensates for this diffusion. Once complete, with a new optical dose target established, OSMO could be re-run to achieve that new dose target. This approximation might not be sufficiently accurate, as the new OSMO dose distribution might be different enough from the prior one to significantly change diffusion behavior. However, it would be significantly less computationally expensive, and it might be adequate.

The first problem also has multiple potential approximation approaches. For example, given an image set, one approach would be to simulate diffusion one projected image at a time. E.g. the effects of a single exposure could be added to the existing distribution of species, then the diffusion model could be run forward in time by one exposure time (e.g. 12.5 ms for prints at 80 °/s with 1 image per degree). This would then be repeated until print completion. Clearly, this would be horribly expensive if nested within OSMO iterations. But perhaps it would be tractable if a dose update to compensate for diffusion were only performed once, as described above. Total computation could be reduced by taking larger time steps when calculating diffusion. For example, it could be assumed that the dose from from each vial rotation is delivered simultaneously; then

diffusion time would be simulated for the time of one rotation. In the least accurate limit, we could assume that dose for the entire print is delivered simultaneously, and then the simulated diffusion time would equal the print time.

Future work would involve determining which compromises are sufficiently accurate, while minimizing computational cost. One way to approach the problem could be to perform a simulation and dose-compensation as accurately as possible for a test structure. Then various approximations to both of the above problems could be made, and the results compared with the accurate result. Hopefully, it will be found that only very simple, inexpensive approximations are necessary for accurate results. Developing fast and effective methods that yield good results with understood bounds on error would be an impactful step forward in control over VAM printing.

4.4.4 OSMO Extrapolation

It may be possible to speed OSMO convergence by observing how voxels in the model are changing during iteration, then extrapolating voxels in the model to improve the model without costly OSMO iterations. This concept is illustrated in Fig. 4.37. In this example, (B) shows an initial OSMO model resulting from 2D filtering. (C) shows the final model that was generated after many iterations. Note that they are significantly different in magnitude. Suppose that, upon iterating OSMO, we observe that a given voxel is trending in a particular direction. E.g., represented by the leftmost red points in (A). Instead of iterating OSMO to slowly push that voxel towards its final value, suppose that we perform an extrapolation in an attempt to replace those costly iterations in a quick step. This, or other series acceleration methods, could potentially dramatically speed OSMO convergence.

Unfortunately, OSMO models are sensitive, and their final value distribution is highly dependent upon the target structure. We found that for a given target geometry, it was possible to tune extrapolation parameters and achieve 50% reductions in OSMO iterations to convergence. Such parameters included: how far to extrapolate, how many prior data points to use for computing the extrapolation slope, whether the extrapolation was linear or included higher order terms, etc.

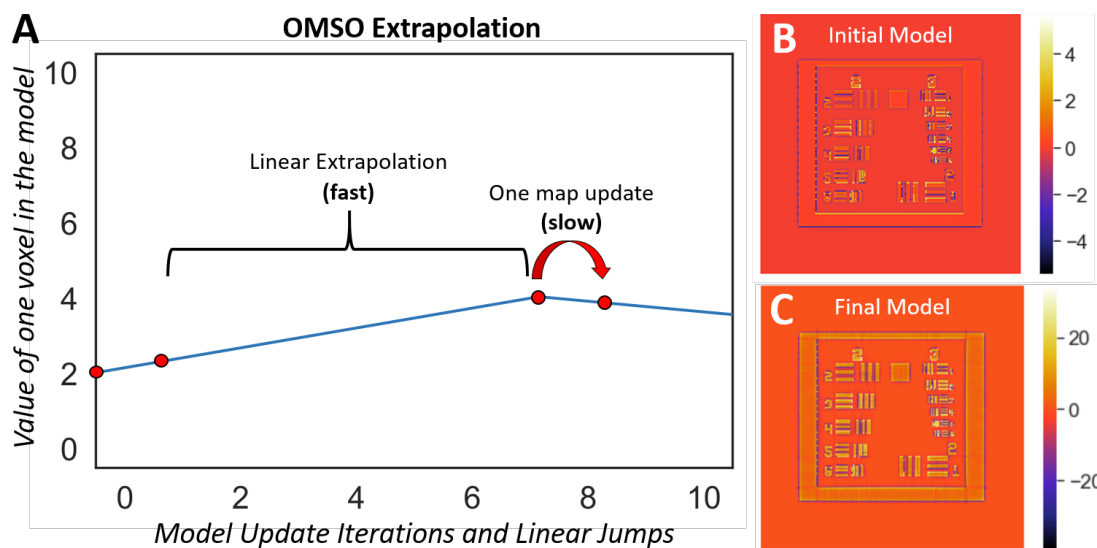


Figure 4.37: OSMO extrapolation strategy. **A**: Illustration of the value of a single voxel in an OSMO model. A single iteration is performed by an OSMO step. Extrapolating model values can speed OSMO convergence. However, we have not yet found an extrapolation speed that, for all targets, always speeds convergence while never slowing it. **B**: An initial OSMO model. **C**: A final OSMO model.

Unfortunately, when chosen incorrectly, including model extrapolation slowed OSMO convergence. However, some of the initial results were promising. One interesting area of future work is to analyze model convergence for a variety of structures, and to devise an extrapolation method that reliably speeds convergence.

4.4.5 More Projection Directions

Thus far, tomographic VAM has used a single latitudinal ring of projection angles with on the order of one image per azimuthal angle. The earliest method of volumetric printing [41] used only three projections, yet produced remarkably accurate prints. This was achieved by using orthogonal projections along the three Cartesian directions. While this method was limited to only a handful of target geometries, it highlights the importance of projection direction when attempting to achieve a given dose distribution, and suggests that a latitudinal ring of projections at a fixed tomosynthesis angle (with 0° being standard tomography) only scratches the surface of potential

control over delivered optical dose.




Figure 4.38: Increased projection direction concepts for VAM. **A:** The addition of projections parallel to the axis of rotation would significantly increase options for dose delivery. **B:** Tilting the axis of resin rotation during printing would allow access to a full range of tomosynthesis angles in a single print. This would likely require index matching, although compensation for image distortion is possible. This approach could also be applied to flat samples.

Fig. 4.38 illustrates just a few of the potential approaches to increasing angular access to a VAM print region. Mathematically, projection access from all hemispherical angles would be ideal. This would be practically very difficult, and it may be the case that most of the benefits are possible with a significantly reduced subset of all the possible angles. Future work to advance control via this method would first involve simulation to understand the benefit yielded by increasing angular access, and to understand which subset of those angles yield the biggest improvement.

4.5 Conclusions

In this chapter we discussed one of the problems fundamental to VAM: how to compute the light patterns to project into resin to get some desired print. We first discussed how reconstructing arbitrary distributions of optical dose is generally impossible with projections of optical intensity, as the only methods of doing this exactly - FBP or equivalent - demand negative projection intensity values.

Instead, we defined the problem of *Binary VAM* in which dose need only be above or below a material gelation threshold. We presented an object-space approach to this problem, and used

it to define the OSMO Algorithm. We defined reconstruction quality metrics for binary VAM reconstructions, and showed OSMO to be superior to prior approaches. We introduced the concept of allowing material inhibition to relax reconstruction constraints, and we showed improved OSMO reconstructions under such conditions. We demonstrated the flexibility of OSMO by performing reconstructions using projection arrangements other than than the standard full range, azimuthal geometry.

We demonstrated an object-space approach to approximating grayscale reconstructions for the purposes of VAM printing with functionally graded materials. We discussed the additional inverse problems - beyond optical reconstruction - that would be required to achieve material control. We extended this concept further by discussing how material inhibition could allow for the application of FBP to printing arbitrary conversion distributions by taking advantage of the linearity of the back projection operator with respect to material conversion. We discussed how to calculate the inhibition period required to print arbitrary distributions of material conversion. We examined that while sufficient inhibition theoretically allows for the printing of any function, not all prints are practically possible, as dose contrast may be prohibitively low. Lastly, we discussed a variety of future research directions.

Chapter 5

Tomosynthesis

A conventional VAM printing geometry has limited printing applications for a few reasons. First, printing must be done into a roughly cylindrical package whose diameter is limited by printer depth of focus. This requires access to the sample from a large azimuthal range of projection directions, and this does not naturally lend itself to printing into non-cylindrical packages. While it might be possible to, e.g., print into a flat sample packaged between glass slides that is then immersed in index matched fluid, not all packages would necessarily be amenable to immersion in index-match fluid, and multiple technical problems would need to be overcome. Second, if we wish to print larger structures by e.g. using a larger projection pattern, by stitching together multiple adjacent prints, or by any any other method of increasing the number of patterned voxels, we can only do so in one dimension - parallel to the axis of resin rotation. This is due to the printer's limited depth of focus which constrains the size of two of the three dimensions of a print.

In this chapter we discuss an alternative printing geometry - tomosynthesis - inspired by laminography. We first introduced this geometry in section 4.1.11.1. Here, images are projected from a circular trajectory, but at a fixed angle to the resin axis of rotation (Fig. 5.1). This has multiple potential advantages over a conventional VAM geometry. First, it may be possible to print directly into flat sample packages, avoiding a curved cylindrical surface and need for index matching. Second, since the depth of focus limits only sample thickness, print size is constrained in only one dimension, not in two as in conventional VAM. This could enable printing large thin samples, either in other step or by stitching together prints in two dimensions. Such example

application concepts are illustrated in Fig. 5.1.

We begin by examining some fundamental limitations of tomosynthesis. We then discuss OSMO reconstructions in detail for a variety of target geometries, we examine proof of concept prints, we detail possible optical system design concepts for tomosynthesis, and we discuss future work.

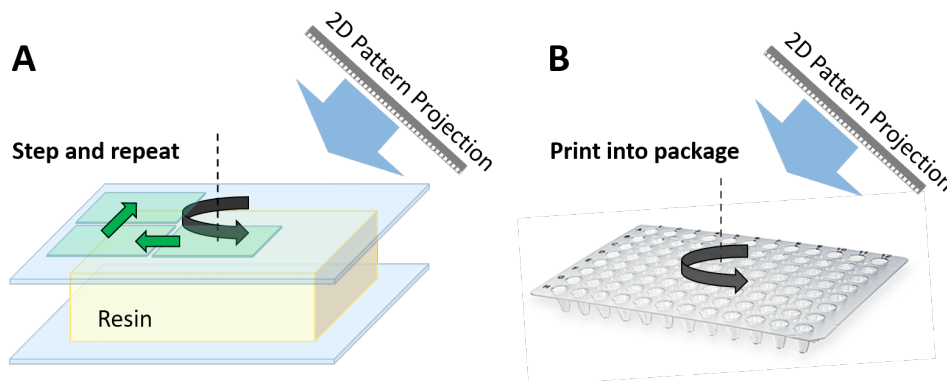


Figure 5.1: Potential applications of tomosynthetic printing geometry. **A**: 2D pattern stitching scheme. Note that if the glass slides are part of the final package, we should be careful to only include resin voxels in the optimization process. That is, we do not care what dose is delivered to glass, air, etc., so we should take care to not over-constrain the optimization problem by including surrounding out-of-part voxels that physically will lie outside of the resin. **B**: Example of printing into a well plate - something that conventional VAM would be unable to do.

5.1 Tomosynthetic Tomography Fundamentals

There are some fundamental questions we should ask before attempting tomosynthetic VAM. Are arbitrary reconstructions possible with this geometry? Is this a slight perturbation from our standard tomographic geometry, or do we expect fundamentally different behavior? Are there fundamental limitations or degradation to print quality? Before blindly exploring the space of trade-offs with this geometry, it is helpful to at least understand the theoretical limits of that space.

Reassuringly, a tomosynthetic geometry is not novel; it is used in tomographic metrology. One of its most ubiquitous applications is diagnostic medical imaging of breast tissue [45]. In that

application, it is used instead of a standard tomographic geometry because of the practical lack of access to full equatorial range of projections. It is for exactly the same reason that tomosynthesis for VAM is appealing - it allows for a different range of access angles. As has been discussed in the literature on tomosynthetic VAM imaging, nearly arbitrary reconstructions are indeed possible. However, there are limitations relating to accessible Fourier content, resulting in an increased minimum voxel size.

It is not surprising that voxel size would increase under tomosynthesis. Imagine we wish to print a single voxel. If for each angle, a beam is projected to that voxel, it is obvious that the intersection of the beams will no longer form a small cylinder as shown in Fig. 2.1, but will instead form a shape closer to a bicone. We can understand this more formally via the Projection Slice Theorem [40]. From this theorem we know that a single 2D tomosynthesis projection will sample frequency data along a plane going through the origin in frequency space, perpendicular to the projection direction. This is similar to what we described when discussing FBP, but instead we consider a 2D projection and 3D Fourier space instead of a 1D projection and 2D Fourier space. Considering the collection of all possible planes formed by a tomosynthetic projection set at some tomosynthesis angle, we see that we can access much of the same region of Fourier space that we could with the standard tomography geometry, except for missing cone-regions above and below the Fourier space origin [45, 46]. Intuitively, this makes sense; if we look at an object from above at an angle instead from from the sides, we will not be able to see the same detailed resolution in the up and down directions (features parallel to the rotation axis). Once we have filled the accessible regions in frequency space, we can take a Fourier transform to compute the resultant points spread function (PSF) of the projection geometry. An example of this is shown in Fig. 5.2. Here, the regions formed by sweeping a plane azimuthally were computed via simple geometry. 2D slices of these regions are shown in yellow in the left column of Fig. 5.2 for the tomosynthesis angles of 10, 45, and 80 degrees, where 0° is standard tomography. Fourier transforming these gives us a real-space PSF, shown in the center and right columns of Fig. 5.2.

Aside from these resolution penalties, FBP works for tomosynthesis, just as it does for to-

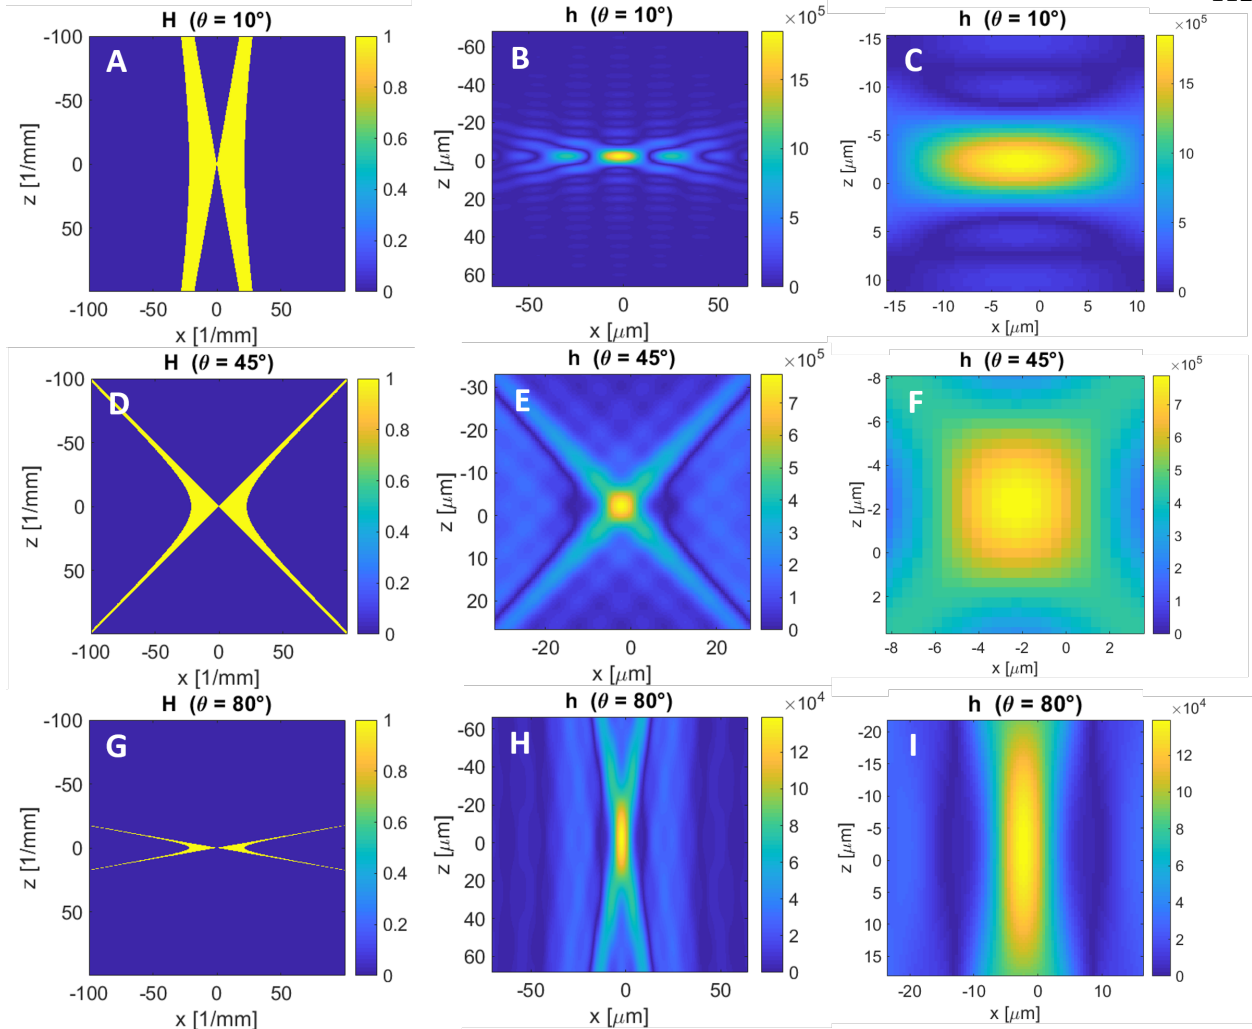


Figure 5.2: Tomosynthesis PSF computed via the Projection Slice Theorem. **Left Column:** 2D slices of frequency space filled by sweeping a plane through the origin about all azimuthal angles. The plane angles are perpendicular to the real space projection directions, where $\theta = 0^\circ$ is standard tomography. **Middle Column:** The PSF of the projection geometry, computed by taking the Fourier transform of the filled Fourier space. **Right Column:** Enlarged view of the PSFs. (A) is not far from standard tomography at 10° , so we see the PSF (C) similar in shape to a cylindrical voxel. At a very steep tomosynthesis angle (G) of 80° , the PSF (I) is elongated vertically; we lose resolution in the vertical direction.

mography, except with a different frequency filter to account for the differing sampling density in frequency space [cite]. This suggests that tomosynthetic VAM is indeed possible. However, without negative intensity (or sufficient material inhibition) we cannot apply FBP, and must instead rely on the material gelation threshold. It is thus not obvious that arbitrary-shape binary reconstructions

are possible with tomosynthetic VAM in the same way they are with standard tomographic VAM. Unlike with the standard geometry, where beams were at least confined to a 2D slice region, with tomosynthesis, each 2D projection addresses all voxels in the 3D volume. It might be the case that without negative projection values, this higher degree of coupling makes some shapes impossible to reconstruct without error voxel for binary VAM. In the next section we explore this by applying the OSMO algorithm to various geometries with various tomosynthesis angles.

5.2 Applying OSMO to Tomosynthesis

Let us consider a few example tomosynthesis reconstructions made by OSMO. We have already shown one example in Fig. 4.25D. Here, a thinker geometry was reconstructed with a 50° tomosynthesis angle was reconstructed with a VER near zero ($1.7\text{E-}4$). It is significant that the VER was positive. Unlike with the standard tomographic projection geometry where it seems that a VER of zero is possible for any structure, we are more constrained with tomosynthesis. Future work could likely develop theory to predict printability under various conditions, but let us start to get a feel for the nature of tomosynthetic VAM with some examples.

Fig. 5.3 shows an example of a reconstruction achieving a positive process window of 0.5% (and therefore a VER of zero). This required just over 1000 OSMO iterations - many more than required by the conventional geometry. This is unsurprising given the more coupled nature of the optimization problem. Note that the areas between the pillars in Fig. 5.3 B & C are high in dose, and the tops and bottoms of the features in Fig. 5.3 C & D have a rounded dose profile - likely a consequence of PSF described in the previous section. Nevertheless, a material threshold nonlinearity would theoretically allow this dose reconstruction to print a perfectly-shaped part. In this case, the tomosynthesis angle was relatively small and the target geometry was sparse - both attributes that intuitively suggest an easy reconstruction.

Fig. 5.4 shows an example of a reconstruction with a moderate tomosynthesis angle of 35° (where 0° is standard tomography), that converged to a relatively large VER of 0.2%. Even though the tomosynthesis angle was 15° shallower than than the 50° reconstruction shown in Fig.

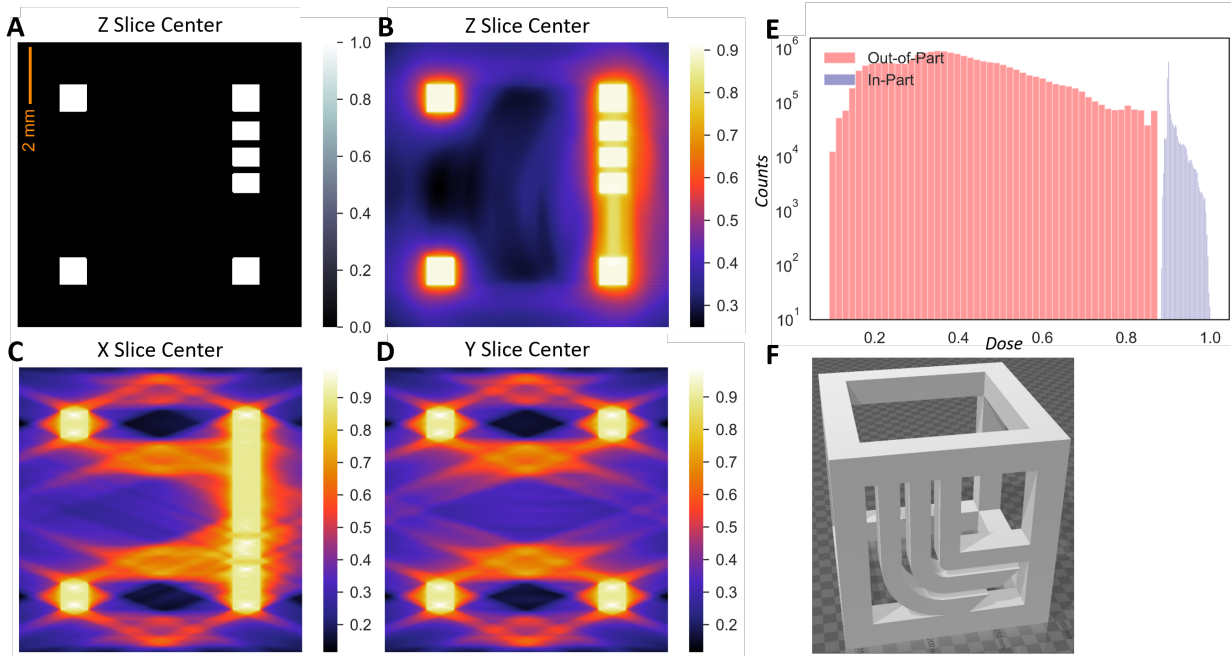


Figure 5.3: Example of a tomosynthetic OSMO reconstruction with a positive (0.5%) process window. D_l was 0.86, D_h was 0.9, just over 1k iterations were used, and the tomosynthesis angle was 20° . **A**: Slice of the target geometry (shown in **F**), made through the center of the shape and normal to the axis of rotation. **B**: The same slice, but through the 3D dose reconstruction. Note that the gaps between the pillars are high in dose - this drives the PW to be small. **C & D**: Slices through the center of the part along different coordinate axes. In both, the axis of rotation lies vertically along the center of the reconstruction slice. The tomosynthesis angle can be seen in the streaks of dose. **D**: Dose reconstruction histogram. **F**: The target geometry.

4.25D, the resultant VER was higher. This confirms the obvious suspicion that the target geometry has a large influence on reconstruction quality metrics, perhaps even more so than tomosynthesis angle. So what is different about this structure? It might be that because struts meet at an angle, features come arbitrarily close to one another. E.g. in Fig. 5.4G, we can see high dose values between the struts. These are regions that, especially given the elongated tomosynthesis PSF, are difficult to keep at a low dose.

If we suspect that tomosynthesis struggles with features that come near each other, especially with respect to position along an axis parallel to the rotation axis, then we might expect a good reconstruction to result from an object with large gaps between features. Let us keep the tomosynthesis angle the same - at 35° - but reconstruct an object without closely spaced features.

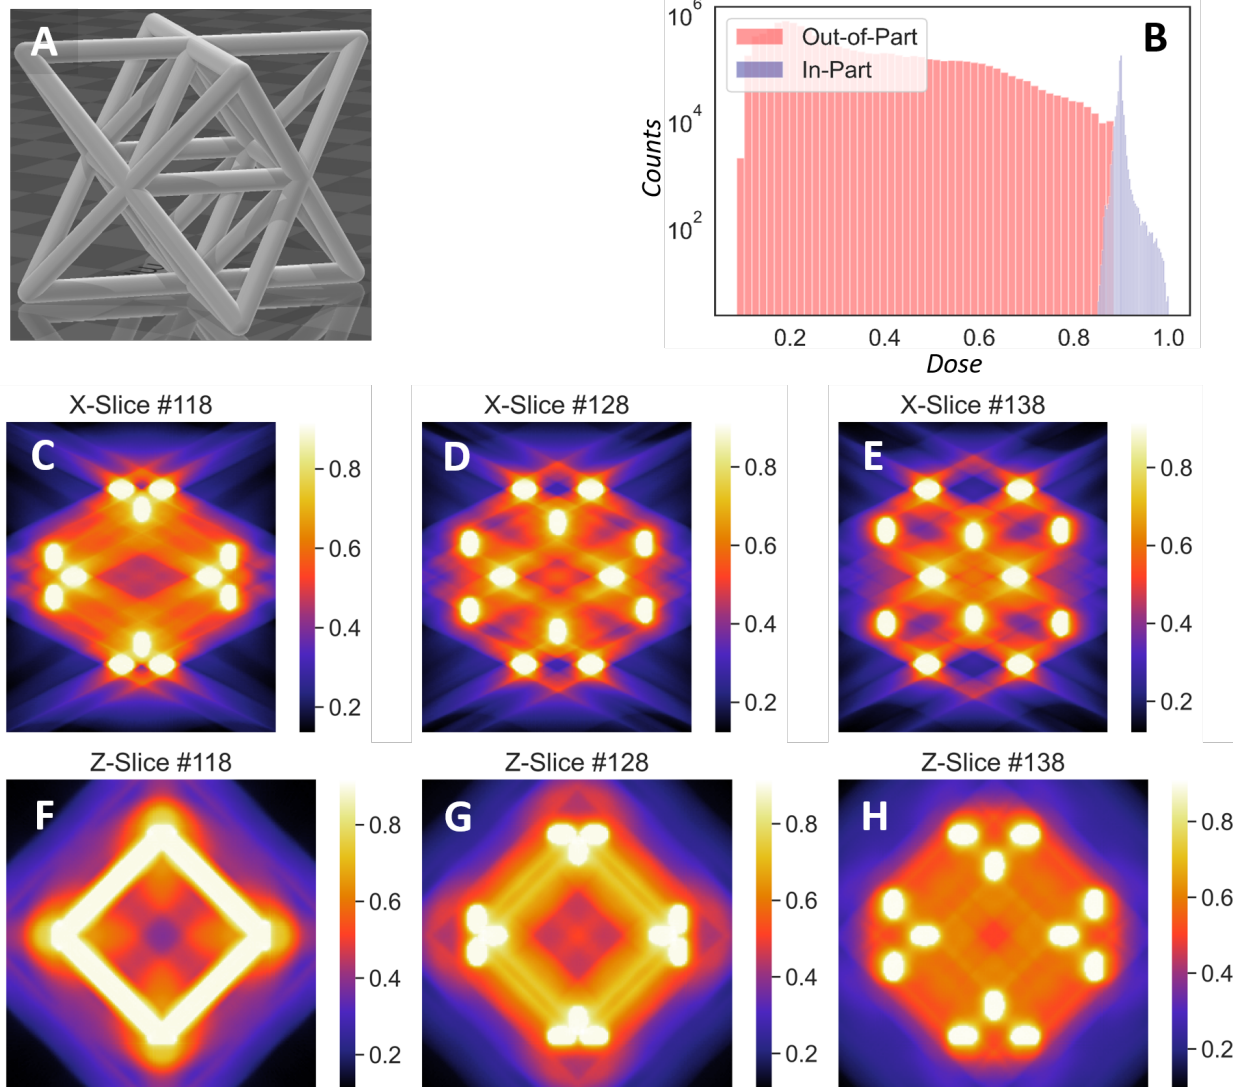


Figure 5.4: Example of a tomosynthetic OSMO reconstruction with a VER of 0.2%. D_l was 0.87, D_h was 0.9, 5k iterations were run, and the tomosynthesis angle was 35° . **A**: An octet-strut target geometry. **B**: Reconstruction dose histograms. $VER > 0$ implies histogram overlap. **C-E**: Slices through the 3D dose Reconstruction, parallel to the axis of rotation. (C) is a central slice. **F-H**: Slices perpendicular to the axis of rotation. (F) is a central slice.

An example of this is shown in Fig. 5.5. Here, under the same print geometry as the previous example, a relatively large process window of 3% is achieved.

Further theory or extensive simulation would be necessary to accurately predict the printability of a given geometry under tomosynthesis at a given angle. However, initial observations suggest

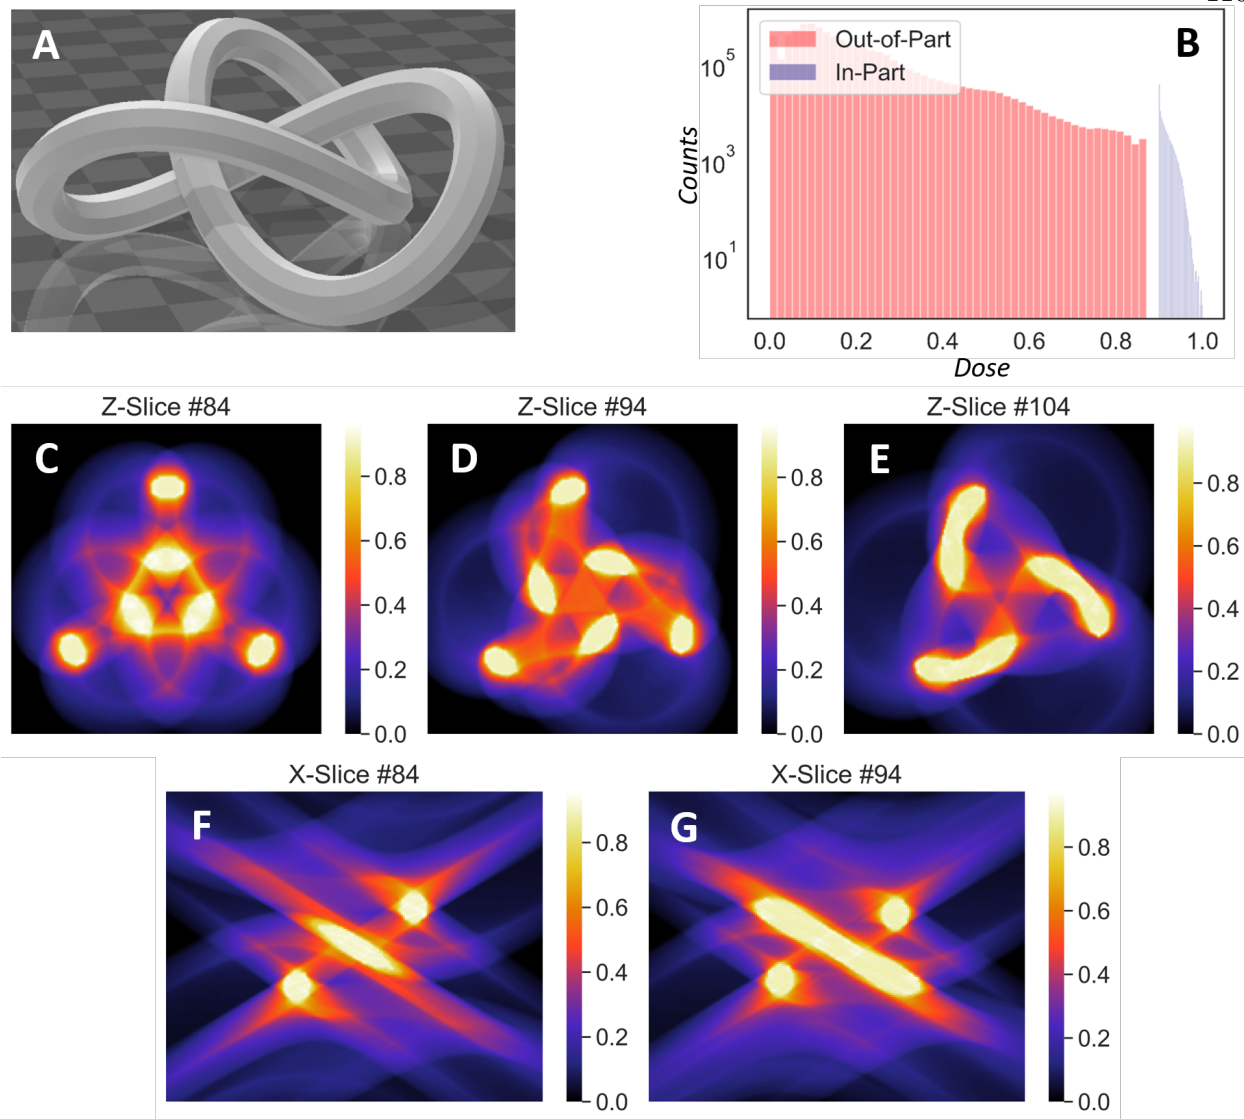


Figure 5.5: Example of a tomographic OSMO reconstruction with PW of 3%. D_l was 0.87, D_h was 0.9, 2k iterations were run, and the tomography angle was 35° . **A**: A knot geometry. **B**: Reconstruction dose histograms. **C-E**: Slices through the 3D dose Reconstruction, perpendicular to the axis of rotation. (C) is a central slice. **F-G**: Slices parallel to the axis of rotation. (F) is a central slice.

that smaller tomography angles (closer to standard tomography) perform better, and that structures without small gaps - especially vertical gaps - are easier to reconstruct well. Given this intuition, let us choose a target that we expect will be easily printable, simulate a reconstruction, then try a real print of the structure.

5.3 Tomosynthesis Experimental Demonstration

Here we present a proof of concept experimental tomosynthesis print. First, as per the previous section, let us choose a target geometry without small gaps in the vertical (parallel to the axis of rotation) direction. For this experiment, let us consider a small section of a dental aligner. Its simulated reconstruction is shown in Fig. 5.6. Note that while this geometry is relatively simple, it still has concavities, convexities, and the ends wrap down (following the contours of teeth). With a process window over 2%, we expect this to be printable.

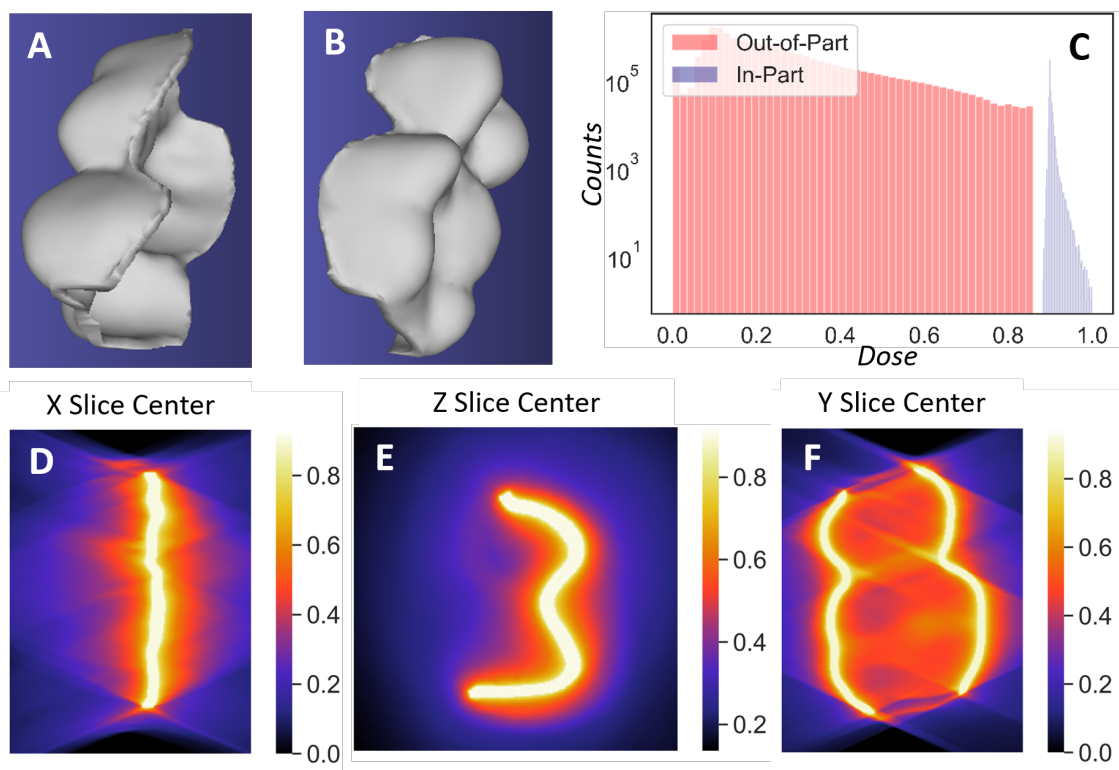


Figure 5.6: A tomosynthetic OSOMO reconstruction with PW of 2.1%. D_l was 0.85, D_h was 0.9, 1.5k iterations were run, and the tomosynthesis angle was 25° . **A & B**: A dental aligner target geometry. **C**: Reconstruction dose histograms. **D**: Central slice through the 3D dose Reconstruction, parallel to the axis of rotation. **(E)** Central slice perpendicular to the rotation axis. **F**: Central slice, normal to the prior slices.

In the lab, the simplest method of applying tomosynthesis projections to resin is to simply tilt the axis of rotation with respect to the projection beam direction. This is illustrated in Fig. 5.7.

The axis of the rotation stage was tilted, and its angle measured to within 0.5° of 25° with a digital angle meter. To accommodate the tilted vial, an index-match bath with large side lengths was used. Before installing the bath, the center of the vial was aligned to the center of the rotation axis via the shadowgram process described in Chapter 2. Rotational alignment of this bath was likely the largest source of beam pointing error. If the light was not normal to the input window of the bath, the resultant refraction would steer the beam to the side - away from the axis of rotation. An attempt was made to align the back-reflection of a red laser to achieve good rotational alignment of the bath, but the method was ad-hoc and its precision is unknown.




Figure 5.7: Experimental setup used to print with the 25° tomosynthesis image set discussed in the previous figure. **A:** Side view. The rotation axis was tilted by 25° , achieving the desired tomosynthesis angle with the writing light. **B:** View of vial with resin (3:1 BPA-GDA:PEGDA250, 20 mM I907) immersed in index-match fluid (the same monomer blend, but without photoinitiator). **C:** Image of the printed part immediately after print completion. The maximum print intensity was 230 mW/cm^2 , the print time was 109 s, and the part was about 1 cm long.

The part was extracted from the remaining resin, it was washed in ethanol, and it was post-cured under a 405 nm flood exposure. Photos of the part are shown in Fig. 5.8. Notably, striation effects are significantly different from standard tomographic prints - an issue we will explore at depth in the next chapter. No quantitative metrology of the part was performed, but qualitatively it appears to have printed correctly. Better, more predictable washing methods are likely necessary for any dimensional measurements of a print to be meaningful beyond very obvious flaws. Regardless of the lack of metrology, this print experimentally demonstrates that tomosynthetic VAM is indeed possible. As far as we are aware, this was the first demonstration of tomosynthetic VAM.

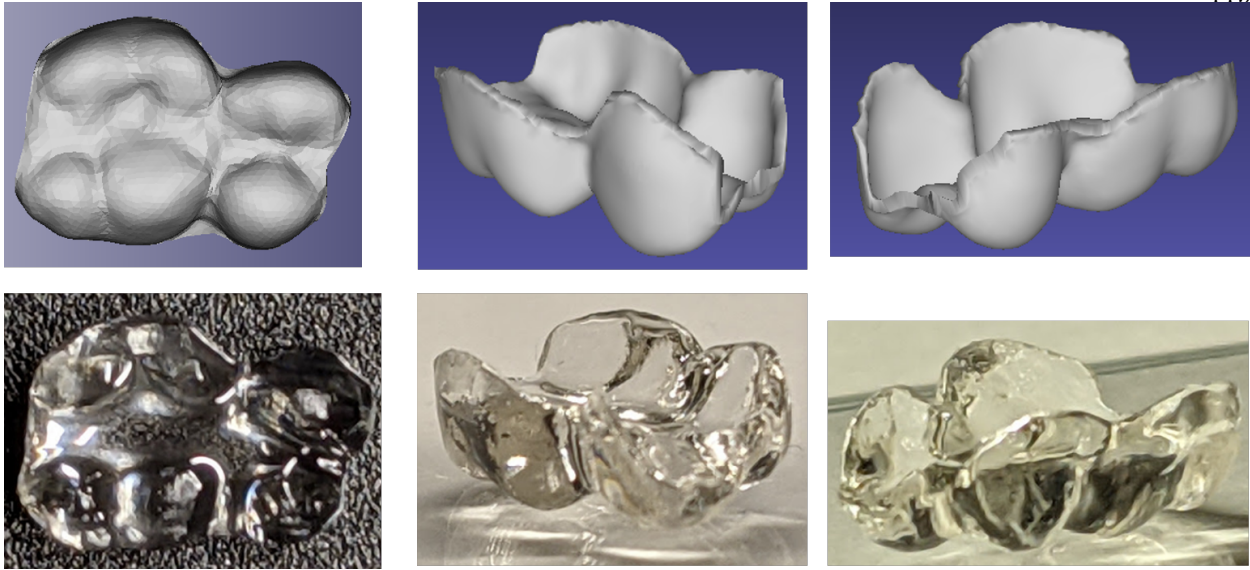


Figure 5.8: Photos of the part printed in Fig. 5.7 for a qualitative comparison to the target geometry. The perspectives are different (camera-part distance, and simulated distance to the STL file), making any but the most crude qualitative comparison difficult. However, the shape looks generally correct, we believe this to be the first experimental demonstration of tomosynthetic printing, and it proves that the method is, indeed, possible.

5.4 Future Work

While the print shown in the previous section demonstrates the feasibility of tomosynthetic VAM, it is still only a print into a vial - something that we could already achieve with standard tomographic VAM. A primary motivation for pursuing tomosynthesis is to print into flat surfaces. In this section, we discuss some of the challenges that will need to be overcome to achieve this, as well as some smaller, intermediate steps that could facilitate the development process. We conclude by discussing the computational work needed to apply tomosynthetic VAM to the problem of grayscale reconstructions for functionally graded material control.

5.4.1 Light Coupling

A basic prerequisite to printing into a flat surface is a method of coupling writing light into it. A few potential approaches to this are outlined in Fig. 5.9. The option shown in (A) is the simplest modification of the previous vial setup. However, a large bath is required to accommodate the

rotation axis, and the bath must scale up with the sample size. This approach also excludes samples that, for whatever reason, can not tolerate immersion in fluid. Refracting in directly (B) Seems simple, but is impractical for a few reasons. First, the image is stretched in one dimension. While we could compensate for this digitally, matching the compressed image with the tomosynthesis angle would present yet another source of potential error. More importantly, however, the index jump from air to resin requires a very high input angle. The 3:1 BPA-GDA:PEGDA250 resin we use was measured on an Abbe refractometer to be 1.533. To achieve a tomosynthesis angle of e.g. 40° in the resin would require a very steep angle incident on the sample. This would make the internal angle very sensitive to the external input angle, likely resulting in error, and certainly being difficult to align. A third approach is to substitute the index-match bath in (A) with a prism, coupled to the sample with a thin layer of index-match oil (C). This approach has the optical advantages of the setup in (A), but most of the sample is in air.




Figure 5.9: A few conceptual approaches to coupling light into flat samples. **A:** Immersing the flat sample in an index-match bath. **B:** Refracting light directly into the sample. Avoiding a back-reflection at the exiting air-glass interface is important for avoiding stray light in the sample. An optical absorber, such as black polydimethylsiloxane, laminated to the back surface, prevents this. **C:** A prism with index-match oil between the prism and the sample.

A first attempt to implement the coupling method shown in Fig. 5.9C is shown in Fig. 5.10. Unfortunately, the resultant print failed. There are a few difficulties. First, the flat sample bottom must be fairly precisely perpendicular to the axis of rotation. Otherwise, if the tip/tilt of the axis is adjusted to be parallel to the prism surface, it will not remain so upon sample rotation. This

causes a non-uniform distances between any given point on the sample and the prism face, resulting in air bubbles being introduced into the oil film. Any air, of course, either does not allow light to enter the sample, or it deflects the light at the wrong angle in the sample. Another difficulty is that gravity tends to drain the oil down the angled surface of the prism. Dripping oil at the top of the sample-prism interface during printing can help replenish the oil, but it does not guarantee a continuous oil film, and it makes a mess. Rotating the entire setup with respect to gravity such that the flat sample is perpendicular to gravity would help with this problem. In general, we expect this method to work, but further developments must first be made in sample packaging, sample clamping, and alignment methods.




Figure 5.10: Experimental attempt at coupling light into a flat sample surface with a prism and index-match oil. **A**: Side view. The sample holder was a glass slide glued to a tube. In terms of the coupling interface, this is equivalent to a sandwich-style sample. **B**: View through the prism of the sample. Where oil does not fill the gap between sample and prism, a reflection is seen instead. The glass slide was not perfectly normal to the tube's central axis, causing the introduction of air into the oil film during sample rotation.

5.4.2 Tomosynthesis Optics

An accurate tomosynthetic VAM printer will also require a modified optical projection system. With a typical imaging system, the plane of best focus will be normal to the projection direction. However, a tilted sample is not normal to the projection beams. One approach is to extend the

DOF. This is illustrated in Fig. 5.11A. However, as we discussed in the chapter on printer design, increasing DOF decreases resolution. (B) shows one approach to tilting the focus plane to be aligned with the sample. If we tilt the object (the DMD), the image will also tilt. However, this introduces aberrations. While it is possible to compensate for such aberrations - as is done in complex designs for tilt-shift photography lenses - it is a complicated way of addressing the problem. A method that is perhaps simpler, is to use a diffractive optical element (DOE) such as a transmissive blazed grating (C). There are two advantages to this approach. First, the writing beams are deflected at an angle without the inherent addition of aberrations, and secondly, the DOF is decoupled from the lateral sample size. A DOE approach is thus one of the most promising options, however future work will be needed to implement it.

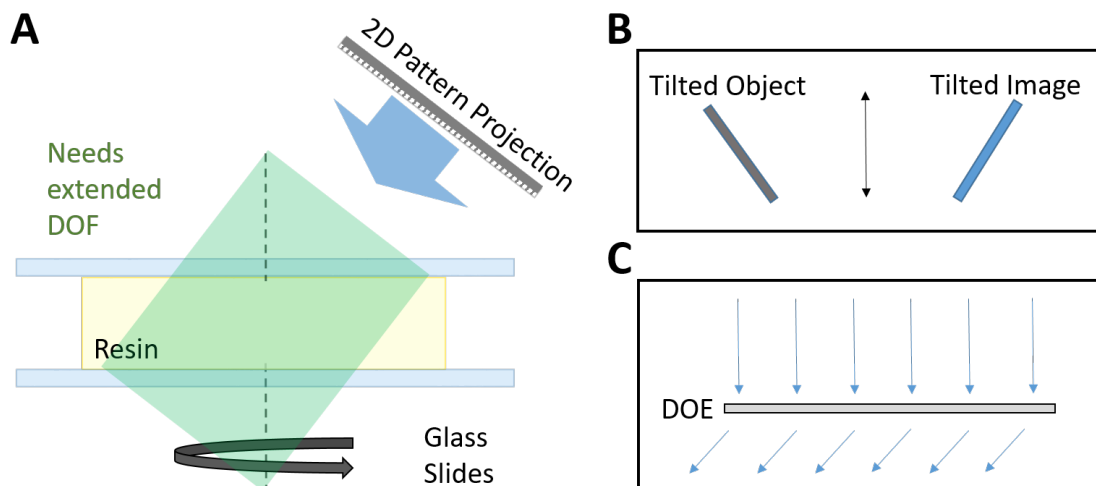


Figure 5.11: Illustration of tomosynthesis optics challenge, ignoring the problem of coupling light into a sample that we discussed in the previous section. **A:** To print the same sized part as with a standard tomography setup, tomosynthetic VAM requires an extended DOF. This requirement becomes even more noticeable at low tomosynthesis angles. Unfortunately, it is these lower angles that we would prefer to use, as they result in improved resolution (see Fig. 5.2). **B:** One approach could be to tilt the DMD. This will result in a tilted image whose degree of tilt goes as the imaging magnification. This, however, creates imaging aberrations. **C:** A diffractive optical element (DOE), such as a blazed transmission grating, could be used to angle the writing light, while keeping the plane of focus parallel to the DOE. This is a particularly attractive solution as it decoupling DOF from the length and width of the sample.

5.4.3 Computation for Stitched VAM

One of the promising advantages of tomosynthetic VAM is the ability to pattern a large sheet of material via the step and repeat approach sketched in Fig. 5.1A. Theoretically, this is a straightforward problem. A tomosynthesis optimization can be performed under the assumption that the entire sheet lies within the object-space reconstruction area. Then, as a printer is scanned over a surface during a step and repeat stitched print, the computed beam intensities can be delivered from various positions of the projector. That is, a projector moving over a 2D print area can essentially synthesize, to mathematical equivalence, an enormous projector capable of printing the entire sheet in a single tomosynthetic reconstruction.

While conceptually simple, this approach requires performing an arbitrarily large reconstruction over an arbitrarily large number of voxels - something that is not generally possible without infinite compute resources. Thus, it will be practically necessary to formulate and solve a new reconstruction problem for the printing of arbitrarily large continuous structures.

One simple approach would be to first perform a reconstruction over one region of resin addressable by a single tomosynthesis position (that is, a print before any stepping). The outermost beams of such a reconstruction will inadvertently deliver some dose to adjacent regions of resin. After stepping to such an adjacent region with some dose already in it, a new reconstruction could be performed, given that existing dose. E.g., if OSMO were being used, any back projected dose would be added to the existing dose, per iteration. For simple structures, this would likely work. For more complex ones, this could easily eat into available process window, as only some subset of beams are optimized at once. One interesting research area would be to devise schemes of optimization that do better than a purely sequential approach, while remaining computationally tractable. This could involve: propagating reconstruction constrains across print regions, computing a few regions at a time, weighting the importance of some beams over others, etc. Finding good approaches to this problem would be helpful not only to step and repeat tomosynthetic VAM, but also to any VAM printing approach that involves scanning or stitching together print regions.

5.4.4 Grayscale Tomosynthetic VAM

One area of future work is to extend tomosynthetic VAM to the problem of grayscale conversion control for the printing of functionally graded materials. As with the standard tomographic geometry, this could be approached by approximating reconstructions of grayscale optical dose, it could be approached by tuning material response, or some combination of both.

Just as it is possible to apply FBP to standard-geometry, positive intensity VAM printing via the inclusion of material inhibition, it should be equally possible to do so with a tomosynthetic printing geometry. When a forward projection is performed on a target in a tomosynthesis geometry, the frequency space of the object is non-uniformly sampled, as discussed in Section 5.1. As with the standard tomographic geometry, this can be compensated by performing a filtering step before back projection. This filter is two dimensional, and is well-known from metrology applications [40]. Note that back projection is still linear. A constant added to a set of images adds a constant to the reconstruction. Thus, we can expect material inhibition to replace negative projection intensities, just as with the standard projection geometry.

Alternatively, it may be possible to tomosynthetically produce approximations to grayscale dose distributions that are sufficiently accurate for some applications. In Section 4.2, we described a method of approximating grayscale dose reconstructions for 2D tomographic reconstructions. This involved applying a 2D frequency filter to the target, perhaps adding a constant shift, forward projecting, then negative-clipping. This filter was one dimension larger than the FBP image-space filter. Similarly, a 3D tomosynthesis filter could be applied to a target function. As described in Section 5.1, this frequency space filter would have a shape of a cylinder with a bicone removed. Forward projection, negative clipping, and back projection would then produce an approximation to a grayscale reconstruction. This approximation could be tuned by adding a shift to the filtered target before performing forward projection. The performance of this method is unknown for tomosynthesis. It is possible that would work well for some applications, and it is also possible that the highly coupled nature of tomosynthesis reconstructions would always result in poor grayscale

reconstructions. However, checking this is straightforward - it is a matter of coding up the 3D frequency filter [40], and simulating some sample reconstructions.

5.5 Conclusions

In this chapter we discussed using a tomosynthetic geometry for VAM. In addition to overcoming print-size limitation in two dimensions, tomosynthesis holds promise as a method of rapid printing into flat surfaces. This could have multiple advantages. It could enable convenient printing into flat sample assemblies that make up part of the finished product; this is particularly relevant to the manufacturing of microfluidic devices. It also could be relevant to any application that involves a limited range of accessible print angles - e.g. printing into a well plate. We discussed the fundamental tomographic limitations associated with tomosynthesis, we presented the first experimental implementation of this printing geometry, we examined algorithm methods and extensions, we discussed the future optical and algorithmic research necessary for an accurate implementation of this approach, and we discussed how tomosynthesis could be used to achieve grayscale optical dose control for VAM printing of functionally graded properties.

Chapter 6

Striations and Latent Image VAM

Multi-step additive manufacturing (AM) is being pursued for a wide range of applications [7] such as: printing optics [8, 9, 10], microfluidics [11, 12, 13], composites with overprinting [14, 15], regenerative medicine [16, 17, 19, 20], etc. However, inherent to multi-step AM are layering effects [50, 51, 52] which limit the performance and applications of printed parts via mechanical inhomogeneity and anisotropy, non-uniform refractive index of printed optics, and cosmetic imperfection.

VAM prints into a volume of photosensitive resin in a single lithographic step, and is thus free of inherent layering effects. Fundamental to this process is a material nonlinearity – a gelation threshold – which enables delivery of optical dose for selective gelation of the desired print volume without unwanted gelation of the surrounding resin. This approach has multiple advantages over multi-step AM. Non-contiguous prints become possible; the lack of material movement during printing removes resin viscosity constraints, thus increasing the range of accessible material properties; print times are dramatically reduced and layering is fundamentally absent from the printing process. However, VAM prints [26, 27, 28, 30, 53] suffer from large striations – similar in appearance to layering, and on the order of print feature size – impacting the homogeneity and shape-accuracy of printed parts, and presenting a significant cosmetic defect.

In this chapter, we examine the source of striations in VAM, and discuss the hypothesis that the material nonlinearity upon which VAM relies also drives striations via a *self-writing waveguide* (SWW) effect. We present a simple and effective method of dramatically reducing striations (Fig. 6.1F) via the addition of a uniform optical exposure applied to the entire volume of resin at the

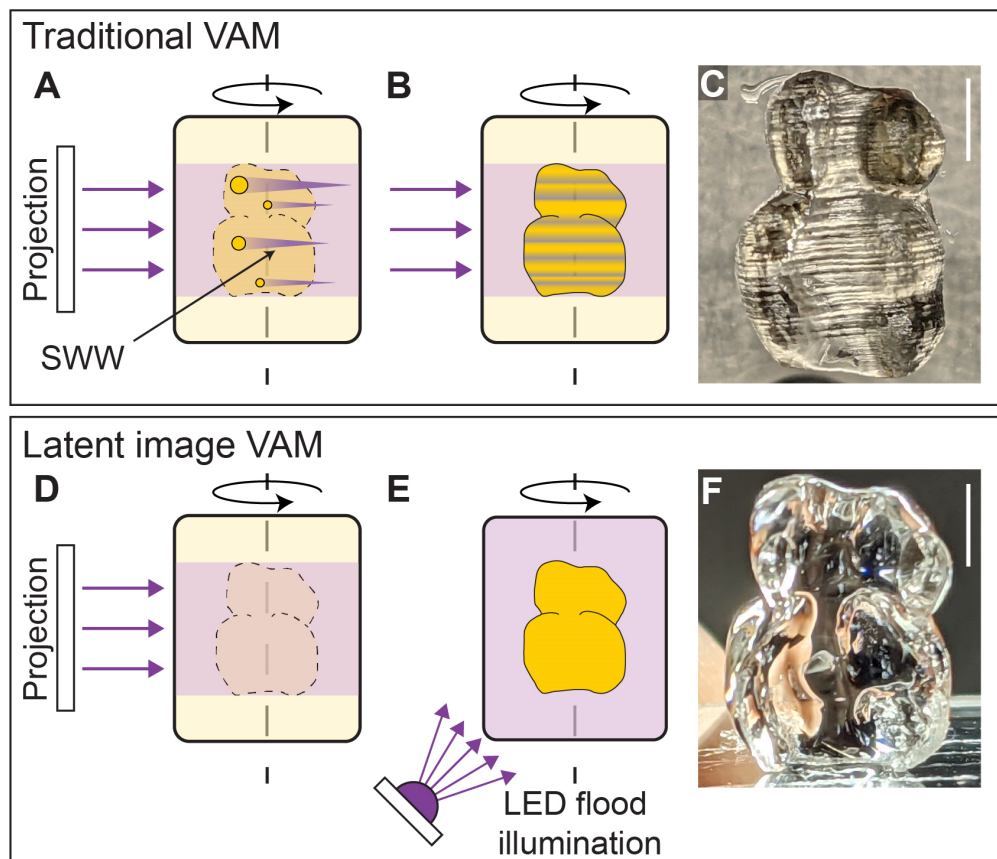


Figure 6.1: **A–B**: A conventional VAM print process. **C**: Resultant dental aligner print exhibiting typical striations. As detailed in Section 6.1, initial local gelation causes focusing, leading to a self-written-waveguide (SWW) effect, manifesting as striations (scale bars: 3 mm). **D–E**: The latent-image process proposed here. **D**: Patterning is stopped just before gelation occurs, leaving a 3D latent image of polymer conversion that is higher in the desired print region than in the surrounding resin. **E**: The latent image is developed across the gelation threshold via diffuse, uniform LED illumination, driving only the desired region to gelation. **F**: The resultant print exhibits dramatically reduced striations, with a smooth surface and improved refractive index homogeneity. Image from [6].

end of a print. Here, VAM patterning starts as usual, steadily increasing polymer conversion in the desired print-regions, but not past the gelation threshold as is typical. Instead, the patterning is stopped just before this material nonlinearity is reached (see Section 6.2.3 for further timing details), then a flood exposure develops the resultant latent image to gelation (Fig. 6.1D-E). We discuss how the technique affects the system tolerances required to maintain theoretically perfect print fidelity. We show that the method has the additional advantage of drastically reducing the

gelation period (GP) of a print, defined here as the time between the first appearance of any gelation and the print completion. This mitigates the problem of partial-print sinking in low power or low viscosity prints, thus further expanding resin options and increasing the efficacy of low-cost VAM printers. We conclude by discussing future work and potential improvements to the method.

6.1 Striations in VAM

Striations are a ubiquitous problem in the field of VAM. A variety of volumetric printers applied to a wide range of materials suffer from large striations on the order of feature size [26, 27, 28, 53, 30]. Such striations not only degrade print-shape accuracy, but any non-uniformities in polymer conversion can manifest as inhomogeneity or anisotropy of modulus, or as unwanted refractive index variability. Thus, although VAM is free from layering, it is not free from layer-like effects. Examples of typical VAM striations are shown in Fig. 6.2 and 6.3. Here, we discuss the hypothesis that striations are caused by a SWW effect, driven by the material nonlinearity upon which VAM fundamentally relies, and that striations are not caused by the direct recording of non-uniform patterning beams.

The strongest evidence that VAM striations are not a direct recording of writing beams is that the latent cure method detailed in this chapter effectively suppresses striations. With this method, typically about 90% of the printing dose is delivered by patterned writing beams before gelation, and then the final 10% of dose (approximately equal to the IPDR - see Section 6.2.1) is delivered by a uniform exposure. If striations were simply a direct recording of non-uniform beams, we would expect to see them develop along with the rest of the part during latent cure. Striation visibility would likely be reduced by the uniform exposure, but only by a small amount. Instead, we observe a nearly complete elimination of striations, suggesting a fundamental difference between using highly directional patterned illumination and an angularly-diffuse flood exposure to cross the gelation threshold.

We hypothesize that striations in VAM are due to a SWW effect. This effect is well-known to occur in many photopolymer exposure systems [54], and has been employed to manufacture




Figure 6.2: Scanning electron microscopy (SEM) scans of typical striations in conventional VAM prints. **A:** A cone printed with the axis of rotation aligned with the central axis of the cone. A roughly $60\ \mu\text{m}$ striation pitch was measured both near the tip of cone, where the focused beam size was measured to be $11\ \mu\text{m}$, and near the base of the cone, where the beam size was $23\ \mu\text{m}$. A single projector pixel was measured to have a focused size of 6 and $18\ \mu\text{m}$ at the center and at the edge of the print, respectively. **B:** A slab print – the same as in Fig. 6.4A, B, F, G – exhibiting irregular striations. Scale bars: $1\ \text{mm}$. Shadowgram images of the cone print formation are shown in Fig. 6.3. See Section 6.2.3 for further print details. Image from [6].

waveguides [55, 56, 57, 58, 59]. Here, as light gels a small region of material, the refractive index increase of this phase change acts as a lens, concentrating the light to the resin just beyond it. The increased intensity causes this next region to also gel more quickly than the surrounding resin. This continues, with each new region of gelation building upon the growing waveguide, until a long waveguide – a striation of increased index – extends through the print region. The properties of SWWs are complex and have been extensively studied in the literature. Notably, depending on material index-change dynamics and on the intensity profile of writing beams, SWWs can form with variable widths, not matching the size of the writing beams. SWWs can be unstable in




Figure 6.3: Shadowgram showing the development of the cone print from Fig. 6.2. The formation of striations is visible upon the initial gelation. Image from [6].

their propagation direction, and filamentation can also occur, with stochastic splitting and merging behavior [59, 60].

The SWW effect is most pronounced when the refractive index of the material changes most quickly with conversion; that is, during the gelation phase change. The much smaller index change of pre-gelation conversion also produces SWWs, as perhaps seen in the liquid resin surrounding the prints in Fig. 6.4A, F, D. However, their impact appears to be negligible, as evidenced by the striation-free prints presented in this paper. Thus, prominent print striations would be expected only when highly directional illumination, such as VAM patterning beams, are incident on a material

near the start of gelation. As detailed in the next section, avoiding directional illumination as the material gels dramatically reduces striations in VAM prints.

6.2 Latent Cure Method

Here we describe a simple, effective, and inexpensive method of avoiding striations in any VAM printing process by using diffuse light instead of highly directly patterning beams to cross the material nonlinearity. In this method, a print starts as usual with patterned illumination locally increasing polymer conversion in the resin. Just before gelation is reached, the writing is stopped, leaving a latent image of the desired print in the form of higher polymer conversion and oxygen depletion for the desired in-part regions than the surrounding out-of-part regions. Then, an LED with a diffuser applies a low-spatial-coherence, uniform exposure to the entire volume of resin such that the light is incident from a wide range of angles (Fig. 6.1E). This develops the latent image into a gelled part, and the flood illumination is ceased before unwanted gelation occurs in the surrounding out-of-part regions, thus preserving selective gelation. With this method, patterning illumination is ideally stopped just before any gelation occurs. In practice, stopping the print slightly earlier or slightly later – at the first sign of gelation on a shadowgram - still produces good results. See Section 6.2.3 for further exposure and timing details. The resultant print exhibits dramatically reduced striations, as shown in Figures 6.1 and 6.4 which compare the current and proposed VAM processes. The shadowgrams in Fig. 6.4A-H, which highlight even small refractive index changes [5, 32] suggest improved uniformity, although measurements without the surrounding resin in place would be more conclusive. Lastly, surface striation effects are eliminated, as shown via SEM scans in Fig. 6.4E and 6.4K. The uniform exposure decreases the dose-contrast between in-part and out-of-part regions, reducing system error tolerances, but only by a small fraction equal to the in-part dose range [2] of the tomographic reconstruction - about 10% for typical prints. Thus, a theoretically perfect VAM print – one free of missing or unwanted gelled voxels – will remain so upon latent cure (see Section 6.2.1 for further details). The latent cure method appears to only require an LED and diffuser for good results; however, the non-directionality of the latent cure

could be improved via multiple light sources, an integrating sphere around the print, or any other modification to increase the range of input angles over which the flood exposure is delivered.




Figure 6.4: Striation mitigation by the latent cure method proposed in this paper. A–D: Shadowgrams and E: SEM scan of conventional VAM prints of a slab, mesh, and tilted mesh showing striations on the order of feature size upon completion. F–J: The same geometries printed using the latent cure method. Striations are largely eliminated, both in surface roughness (J), and in refractive index uniformity shown via shadowgrams of prints still immersed in resin (F–J). A–B and F–G illustrate the improved shape fidelity yielded by the method. Scale bars: 1 mm for A–J, with $500\ \mu\text{m}$ for zooms. The shadowgrams were captured with the basic collimated LED shadowgraph setup described in [5]. Image from [6].

6.2.1 Dose Contrast Reduction Upon Latent Cure

Here we detail the reduction in dose contrast introduced by the latent cure step, and we discuss when such a reduction could impact print fidelity.

The uniform latent cure flood exposure adds optical dose to both the in-part and out-of-part resin regions. This reduces the contrast between these regions, but only by a percentage equal to the in-part dose range [5]. Figure 6.5 shows an example of a tomographic dose distribution with and without a latent cure. Thus, for the simple case of a positive process window with ideal materials (e.g. with a well-defined nonlinearity, no dark-polymerization, etc.), we can expect no degradation

of print-accuracy upon latent cure.




Figure 6.5: In-part and out-of-part dose histograms, in normalized units, of VAM reconstructions. The top image shows a process window [2] of 9.1%. Adding a latent flood exposure to the entire volume shrinks that to 8.2%, slightly reducing the error tolerance of the printing process (e.g. timing, illumination uniformity, etc.). When a VAM print is stopped just prior to the gelation of any voxels, the in-part histograms highest values are just under the gelation threshold. Thus, to push these in-part voxels to gelation, a dose equal to the width of the in-part histogram (the In Part Dose Range [2]), must be added. Adding this value of dose to all voxels, both in and out of part, reduces the contrast of the reconstruction. However, as long as the histograms are separated by a positive value, the gelation threshold can be chosen to lie between them, thus gelling all in-part voxels without inadvertently gelling any out-of-part voxels. The latent VAM technique would likely work well in combination with the image-set re-binning technique described in Section 4.1.7. A re-binned image set would require a shorter flood exposure period, and the images in the final rotation would better match a uniform exposure than would a standard image set. Both these factors would minimize the extent to which the process window size was reduced. Image from [6].

In the case of a tomographic reconstruction unable to achieve a voxel error rate of zero [2] - one that would produce error voxels - the addition of a latent cure would further degrade the print. That is, in the case of a reconstruction where the in-part and out-of-part histograms overlap, adding a latent cure exposure will further that overlap, thus increasing the number of error voxels regardless of where the gelation threshold is set.

It can be the case that a small positive process window is insufficient for avoiding gelation errors. This can be due to the process window being taken up by timing or other experimental errors. However, even under otherwise ideal conditions, the resin curing dynamics may require not only a positive process window, but one of at least a certain size. For example, consider a material that is so mechanically weak upon initial gelation that it does not survive removal and washing

without breaking or deforming. To print such a resin, it is necessary to drive the in-part regions of a print beyond gelation, to a degree of conversion yielding sufficient strength. The additional optical exposure required to do this also increases optical dose for the out-of-part regions. If the process window is positive but very small, the dose required to sufficiently cure in-part voxels beyond gelation can also cause gelation in out-of-part voxels. Thus, some materials require a minimum process window size to be printable without gelation errors. An example of such a resin response is illustrated in Fig. 6.6.




Figure 6.6: Sketch of a resin response requiring a large process window to avoid gelation errors. Consider a resin that requires a degree of conversion beyond gelation in order to survive post processing. The additional dose required to attain this conversion requires a process window of at least that dose in order to avoid unwanted out-of-part gelation. Image from [6].

In summary, printing without gelation errors can require a positive process window of a minimum size. This can allow for system errors in timing, optical non-uniformity, etc. This can also allow for curing in-part regions beyond gelation without any unwanted gelation in out-of-part regions. Thus, for the latent cure step to introduce no print-shape errors, the final reduced process window size must still be sufficient given the constraints of the entire printing process.

6.2.2 Timing the Stopping of VAM Prints for Latent Cure

When using the latent cure method, patterning illumination is ideally stopped just before any gelation occurs. In practice, stopping the print slightly earlier or later still produces good results. If the print is stopped well before gelation instead of at the last possible moment before gelation, the contrast of the final dose-reconstruction will be reduced. This is because a larger latent-cure dose will then be required to bring the part to gelation. Adding a larger uniform dose to a smaller reconstruction patterned dose distribution will reduce the final contrast of the reconstruction. However, as long as the reconstruction process window is positive (as shown in Figure 6.5), the final reconstruction will still have a positive process window, and thus will produce a theoretically perfectly-gelled part. A smaller process window, however, will leave less room for tolerating system errors while maintaining correct gelation, as discussed in Section 6.2.1.

A print can also be stopped late, after some gelation within the print region has occurred. In this case, we would expect that any such gelation would cause striations via the SWW effect. However, in practice we have found that if only a small percentage of the part gels before the print is stopped, a latent cure still produces a print free of noticeable striations. Thus, the first signs of gelation during a print may be used as a signal to stop the print, avoiding the need for careful control over print timing.

6.2.3 Details of Illumination and of Example Prints

In this section we briefly further describe the illumination setup used in the latent cure process. We also present details of the prints shown in Figs. 6.1-6.4.

The prints shown in Fig. 6.1 were made on a laser-based printer with a peak print-plane intensity of 550 mW/cm^2 at a wavelength of 405 nm. The resin used was a 3:1 mixture of Bisphenol A glycerolate (1 glycerol/phenol) diacrylate and Poly(ethylene glycol) diacrylate (Mn 250). The photo-initiator used was Irgacure 907 at 80 mM and 20 mM concentrations for the conventional and latent VAM prints, respectively. For the conventional VAM print (Fig. 6.1C), the print time

was 48 seconds. For the latent cure print (Fig. 6.1F), patterned illumination was incident for 77 seconds at which point, a few small portions of the part had started to gel. Then 405 nm uniform illumination was applied for 10 seconds with an LED with a diffuser, completing the print. Note that the latent print was slower than the conventional print, since it only had a quarter of the initiator concentration. The LED intensity was approximately 1.5 W/cm^2 , however, it was incident on the corner of an index-matching bath (illustrated in Fig. 6.7, refracting towards the print from multiple directions. Thus, the flood-cure intensity received by the print could have been as much as double the LED intensity. Future work could carefully measure the flood dose delivered. However, the intensity was arbitrarily set; the flood source was applied until print completion, as observed on a shadowgram.




Figure 6.7: **A**: Sketch of latent-cure illumination setup. An LED with diffuser is aimed at the corner of the VAM index match bath; it is also aimed upwards so that it is not directly primarily orthogonally to the rotation axis. **B**: Image of a VAM print in progress. As soon as the first hint of gelation appeared on a shadowgram, the writing illumination was turned off. **C**: Image of the same print during the latent-cure phase. Similar setups have been used on other VAM printers to good effect. The particulars of the flood-exposure setup seem to not be important; so far they have all worked.

Prints shown in Fig. 6.2 and 6.4 were made on a laser-based printer with peak print-plane intensity of 204 mW/cm^2 at a wavelength of 442 nm. The resin used was pentaerythritol tetraacrylate with 10 mM camphorquinone photoinitiator. An equal mass of ethyl 4-(demethylamino)benzoate (to camphorquinone) was added to the resin. For the cone made by conventional VAM in Fig. Fig. 6.2, patterned illumination was incident for 55 seconds. For the slab and lattice made by conven-

tional VAM in Fig. 6.4, patterned illumination was incident for 55 and 60 seconds, respectively. For the slab and lattice made by latent image VAM in Fig. 6.4, patterned illumination was incident for 23 and 30 seconds, respectively. Then, 455 nm uniform LED illumination was applied for 10 seconds for each, completing the print. LED intensity at the print plane was approximately 40 mW/cm^2 .

6.3 Latent Cure to Avoid Partial-Print Sinking

The latent cure step also serves as an opportunity to apply high intensity optical exposure during the GP of a print. This can dramatically shorten the GP, ameliorating the distortion of a print due to partial part sinking during printing [27].

This effect occurs when the polymerization rate is slow relative to the settling velocity of the gelled material, as is the case for systems with low print power, low viscosity resins, or large change in the density upon gelation [61, 62]. An inexpensive latent cure source, free of the etendue requirements that practically limit patterned print power [27], can be added to any VAM system to expand the range of materials that it can accurately print. This could allow low-power VAM printers to use resins previously only accessible by costly, high-power systems, just as it could allow state of the art, high-power printers access to a new range of resins. Figure 6.8 shows an example of how latent cure enables printing in a low viscosity ($< 1 \text{ Pa s}$) resin that would otherwise be untenable, given printer-power.

Prints for Fig. 6.8 were made on a 405 nm LED-based CAL printer (DLI 3DLP9000 UV retrofitted with a telecentric lens) with a peak print-plane intensity of 65 mW/cm^2 . Low viscosity resin ($< 1 \text{ Pa s}$, $\alpha = 0.2852 \text{ cm}^{-1}$) was prepared with a 1:1 molar ratio of tris[2-(acryloyloxy)ethyl] isocyanurate (TAE-ICN, Sigma) to poly(ethylene glycol) diacrylate (PEGDA, $M_n = 575$, Sigma) and 40 mM of Irgacure 907 (Sigma) as photoinitiator. The mixture was stirred at $60 \text{ }^\circ\text{C}$ until all the TAE-ICN crystallites were dissolved. In conventional VAM printing of the hollow cube in Fig. 6.8, the print was deemed completed after 330 s. In the latent VAM case, after 215 s of patterned irradiation a 12 s latent exposure is applied to complete the print. Latent curing was done with




Figure 6.8: Latent cure as a method of speeding the GP in a low power VAM print. **A–D**: Conventional VAM process, suffering from low print power and low viscosity resin. The partially-gelled regions sink before the part is complete. **D**: Severely distorted final part. **E–H**: VAM print with latent cure used to quickly develop the entire print through its GP before significant sinking occurs. Even though the patterning process and viscosity were unchanged from **A–D**, the rapid gelation via latent cure results in, **H**: a significantly improved print. The GP could be easily further shortened by increasing the power of the latent exposure. In this case, a GP of only a few seconds was difficult to control, so the latent-cure power was reduced so that gelation could be easily observed, and the latent exposure stopped before unwanted gelation. Image from [6].

an OmniCure S2000 Spot UV Curing system with a 405 nm filter with a custom diffuser on the output resulting in an approximate latent curing intensity of 150 mW/cm^2 .

6.4 Conclusions

Like the layering effects that limit the applications of typical 3D printing methods, striations in VAM can cause non-uniformity in modulus and refractive index, and degrade print-surface accuracy. We discussed the hypothesis that striations are not simply recorded writing beams, and are instead caused by a SWW effect [54, 55, 56, 57, 59], driven by the material nonlinearity upon which VAM fundamentally relies. This effect can be avoided by adding a diffuse latent-

cure step. We showed the improved uniformity yielded by this method, we discussed how the resultant reduction in system error tolerances is small, and how theoretically-perfect [2] printing is maintained by the method. We demonstrated the latent cure step as method of quickly delivering optical power during the GP of a print, mitigating the problem of partial-print sinking that plagues low-power VAM printers, or resins with a particularly low viscosity or high degree of densification upon gelation. Future work could improve timing control of the latent cure, allowing for sub-second GPs. Lastly, the non-directionality of the latent cure could be improved by an array of lamps, by a surrounding integrating sphere, etc., improving the uniformity of the method, and minimizing the SWW effect.

Chapter 7

Conclusions & Summary of Future Work

VAM holds promise as the future of additive manufacturing for a wide range of applications. Compared to competing techniques, VAM can print a broader range of geometries, it promises dramatic improvements to strength and other properties of printed parts, it offers improved uniformity of such properties, and it does so with orders of magnitude of improvement to print speed over a wide range of size scales. However, prior to the work presented in this thesis, VAM suffered from multiple limitations that both restricted its scope of application and prevented its practical use as a manufacturing technique.

The **first** limitation was a lack of tomographic computation methods to provide sufficient optical dose control. From printing 3D parts without shape error, to controlling grayscale optical dose for e.g. printing optics or functionally graded materials, control over dose is critical in VAM. We presented a simple and effective approach to VAM image computation that significantly improved on prior methods. We demonstrate the flexibility of the approach by extending it to include material response models, optical models including beam-shape and occlusions for overprinting, as well as the problem of controlling grayscale dose throughout a print.

The **second** limitation was the requirement of cylindrical sample packages. We presented a tomosynthetic VAM printing geometry that allows for printing into flat surfaces. This lends itself to applications such as the manufacturing of flat microfluidic chips, printing into well plates, or any other printing application requiring beam access through a surface. It further allows for the patterning of arbitrarily large sheets. We discussed theory and image computation for this printing

geometry, and we presented experimental proof of concept prints.

The **third** limitation was that VAM ubiquitously suffered from large striations – similar in appearance to layering, and on the order of print feature size – impacting the homogeneity and shape-accuracy of printed parts, and presenting a significant cosmetic defect. We presented evidence that the material nonlinearity upon which VAM relies also drives striations via a self-writing waveguide (SWW) effect. We demonstrated a simple and effective method of dramatically reducing striations via a latent-cure, flood-exposure step. We additionally showed this to drastically shorten the period from initial gelation to print completion. This mitigates the problem of partial-print sinking in low power or low viscosity prints, thus further expanding resin options and increasing the efficacy of low-cost VAM printers.

These advancements have enabled significant improvements to VAM print resolution and uniformity (e.g. [5, 6]), and as of this writing, VAM commercialization efforts are already underway. However, there remains much potential for further advancement. The last three chapters end with in-depth discussions of future work. But we conclude here with a high-level summary of the most important work to come.

The three pillars of VAM - optics, materials, and computation - all have potential for significant improvement. These are highly coupled and will likely be gradually improved together; solving a pressing problem within one pillar will likely leave another pillar as the next bottleneck to VAM performance.

An impactful area for improvement in VAM lies in process control. Improved control over optical dose delivery could come from more precise, better calibrated projection systems. Measuring the nonreciprocal effects of intensity and time, then applying calibrations via tomographic calculation, would improve the ability to predict and control print conversion. Similarly, models and measurement of diffusion could allow for compensation by projection images. Improvements to print post processing as well as in-situ metrology techniques will likely be necessary to fully realize the potential improvements from better process control. Furthermore, a good understanding of how a print changes during post-processing could allow writing images to compensate for detrimental

effects. The aggregate effect of all these improvements would likely dramatically advance VAM print control and fidelity.

The development of new printing geometries and sample packaging schemes is an exciting area for growth in VAM. This work presented the fundamentals of tomosynthetic VAM, along with the first proof of concept prints. However, tomosynthesis is just one of many possible VAM printing geometries. Multiple projectors, various sample package shapes, new material movement schemes, etc. will all enable new VAM modalities. Conceptually, VAM printing is possible whenever light can be delivered to photosensitive material from multiple directions, and there are many ways to package material and to deliver light. It seems likely that the VAM modalities presented in this work will be only the first of many, with much exciting work ahead to bring the field of VAM to maturity.

Bibliography

- [1] NIST Chemistry WebBook, SRD 69. <https://webbook.nist.gov>.
- [2] Charles M. Rackson, Kyle M. Champley, Joseph T. Toombs, Erika J. Fong, Vishal Bansal, Hayden K. Taylor, Maxim Shusteff, and Robert R. McLeod. Object-space optimization of tomographic reconstructions for additive manufacturing. *Additive Manufacturing*, 48:102367, December 2021.
- [3] Caitlyn C. Cook, Erika J. Fong, Johanna J. Schwartz, Dominique H. Porcincula, Allison C. Kaczmarek, James S. Oakdale, Bryan D. Moran, Kyle M. Champley, Charles M. Rackson, Archish Muralidharan, Robert R. McLeod, and Maxim Shusteff. Highly Tunable Thiol-Ene Photoresins for Volumetric Additive Manufacturing. *Advanced Materials*, 32(47):2003376, 2020. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/adma.202003376>.
- [4] Bryan Moran, Erika Fong, Caitlyn Cook, and Maxim Shusteff. Volumetric additive manufacturing system optics. In Benjamin L. Lee and John Ehmke, editors, *Emerging Digital Micromirror Device Based Systems and Applications XIII*, page 2, Online Only, United States, March 2021. SPIE.
- [5] Joseph Toombs, Manuel Luitz, Caitlyn Cook, Sophie Jenne, Chi Chung Li, Bastian Rapp, Frederik Kotz-Helmer, and Hayden Taylor. Volumetric Additive Manufacturing of Silica Glass with Microscale Computed Axial Lithography. [arXiv:2110.01651 \[physics\]](https://arxiv.org/abs/2110.01651), October 2021. arXiv: 2110.01651.
- [6] Charles M. Rackson, Joseph T. Toombs, Martin P. De Beer, Caitlyn C. Cook, Maxim Shusteff, Hayden K. Taylor, Robert R. McLeod, and Robert R. McLeod. Latent image volumetric additive manufacturing. *Optics Letters*, 47(5):1279–1282, March 2022. Publisher: Optica Publishing Group.
- [7] Carl Schubert, Mark C. van Langeveld, and Larry A. Donoso. Innovations in 3D printing: a 3D overview from optics to organs. *British Journal of Ophthalmology*, 98(2):159–161, February 2014. Publisher: BMJ Publishing Group Ltd Section: Innovations.
- [8] Du T. Nguyen, Cameron Meyers, Timothy D. Yee, Nikola A. Dudukovic, Joel F. Destino, Cheng Zhu, Eric B. Duoss, Theodore F. Baumann, Tayyab Suratwala, James E. Smay, and Rebecca Dylla-Spears. 3D-Printed Transparent Glass. *Advanced Materials*, 29(26):1701181, 2017. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/adma.201701181>.

- [9] Timo Gissibl, Simon Thiele, Alois Herkommer, and Harald Giessen. Two-photon direct laser writing of ultracompact multi-lens objectives. Nature Photonics, 10(8):554–560, August 2016. Number: 8 Publisher: Nature Publishing Group.
- [10] Simon Thiele, Kathrin Arzenbacher, Timo Gissibl, Harald Giessen, and Alois M. Herkommer. 3D-printed eagle eye: Compound microlens system for foveated imaging. Science Advances, 3(2):e1602655, February 2017. Publisher: American Association for the Advancement of Science Section: Research Article.
- [11] Hua Gong, Bryce P. Bickham, Adam T. Woolley, and Gregory P. Nordin. Custom 3D printer and resin for 18 m × 20 m microfluidic flow channels. Lab on a Chip, 17(17):2899–2909, August 2017. Publisher: The Royal Society of Chemistry.
- [12] Sidra Waheed, Joan M. Cabot, Niall P. Macdonald, Trevor Lewis, Rosanne M. Guijt, Brett Paull, and Michael C. Breadmore. 3D printed microfluidic devices: enablers and barriers. Lab on a Chip, 16(11):1993–2013, 2016. Publisher: Royal Society of Chemistry.
- [13] Reza Amin, Stephanie Knowlton, Alexander Hart, Bekir Yenilmez, Fariba Ghaderinezhad, Sara Katebifar, Michael Messina, Ali Khademhosseini, and Savas Tasoglu. 3D-printed microfluidic devices. Biofabrication, 8(2):022001, June 2016. Publisher: IOP Publishing.
- [14] Catalin Jianu, Giuseppe Lamanna, and Constantin Gheorghe Opran. Research Regarding Embedded Systems of Robotic Technology for Manufacturing of Hybrid Polymeric Composite Products. Materials Science Forum; Pfaffikon, 957:267–276, June 2019. Num Pages: 267-276 Place: Pfaffikon, Switzerland, Pfaffikon Publisher: Trans Tech Publications Ltd.
- [15] Md Naim Jahangir, Jeremy Cleeman, Hyun-Jun Hwang, and Rajiv Malhotra. Towards out-of-chamber damage-free fabrication of highly conductive nanoparticle-based circuits inside 3D printed thermally sensitive polymers. Additive Manufacturing, 30:100886, December 2019.
- [16] Elliot S. Bishop, Sami Mostafa, Mikhail Pakvasa, Hue H. Luu, Michael J. Lee, Jennifer Moriatis Wolf, Guillermo A. Ameer, Tong-Chuan He, and Russell R. Reid. 3-D bioprinting technologies in tissue engineering and regenerative medicine: Current and future trends. Genes & Diseases, 4(4):185–195, December 2017.
- [17] Nadav Noor, Assaf Shapira, Reuven Edri, Idan Gal, Lior Wertheim, and Tal Dvir. 3D Printing of Personalized Thick and Perfusable Cardiac Patches and Hearts. Advanced Science, 6(11):1900344, 2019. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/advs.201900344>.
- [18] Bagrat Grigoryan, Samantha J. Paulsen, Daniel C. Corbett, Daniel W. Sazer, Chelsea L. Fortin, Alexander J. Zaita, Paul T. Greenfield, Nicholas J. Calafat, John P. Gounley, Anderson H. Ta, Fredrik Johansson, Amanda Randles, Jessica E. Rosenkrantz, Jesse D. Louis-Rosenberg, Peter A. Galie, Kelly R. Stevens, and Jordan S. Miller. Multivascular networks and functional intravascular topologies within biocompatible hydrogels. Science, 364(6439):458–464, May 2019. Publisher: American Association for the Advancement of Science Section: Report.
- [19] Michael Randazzo, Jared M. Pisapia, Nickpreet Singh, and Jayesh P. Thawani. 3D printing in neurosurgery: A systematic review. Surgical Neurology International, 7(Suppl 33):S801–S809, November 2016.

- [20] L. A. Hockaday, K. H. Kang, N. W. Colangelo, P. Y. C. Cheung, B. Duan, E. Malone, J. Wu, L. N. Girardi, L. J. Bonassar, H. Lipson, C. C. Chu, and J. T. Butcher. Rapid 3D printing of anatomically accurate and mechanically heterogeneous aortic valve hydrogel scaffolds. *Biofabrication*, 4(3):035005, August 2012. Publisher: IOP Publishing.
- [21] Paulo Jorge Bártolo, editor. *Stereolithography: Materials, Processes and Applications*. Springer US, 2011.
- [22] Haoyuan Quan, Ting Zhang, Hang Xu, Shen Luo, Jun Nie, and Xiaoqun Zhu. Photo-curing 3D printing technique and its challenges. *Bioactive Materials*, 5(1):110–115, March 2020.
- [23] Jingchao Jiang, Xun Xu, and Jonathan Stringer. Support Structures for Additive Manufacturing: A Review. *Journal of Manufacturing and Materials Processing*, 2(4):64, December 2018. Number: 4 Publisher: Multidisciplinary Digital Publishing Institute.
- [24] Xin Wang, Man Jiang, Zuowan Zhou, Jihua Gou, and David Hui. 3D printing of polymer matrix composites: A review and prospective. *Composites Part B: Engineering*, 110:442–458, February 2017.
- [25] Halil L. Tekinalp, Xiangtao Meng, Yuan Lu, Vlastimil Kunc, Lonnie J. Love, William H. Peter, and Soydan Ozcan. High modulus biocomposites via additive manufacturing: Cellulose nanofibril networks as “microsponges”. *Composites Part B: Engineering*, 173:106817, September 2019.
- [26] Brett E. Kelly, Indrasen Bhattacharya, Hossein Heidari, Maxim Shusteff, Christopher M. Spadaccini, and Hayden K. Taylor. Volumetric additive manufacturing via tomographic reconstruction. *Science*, 363(6431):1075–1079, March 2019. Publisher: American Association for the Advancement of Science Section: Report.
- [27] Damien Loterie, Paul Delrot, and Christophe Moser. High-resolution tomographic volumetric additive manufacturing. *Nature Communications*, 11(1):852, December 2020.
- [28] Caitlyn C. Cook, Erika J. Fong, Johanna J. Schwartz, Dominique H. Porcincula, Allison C. Kaczmarek, James S. Oakdale, Bryan D. Moran, Kyle M. Champley, Charles M. Rackson, Archish Muralidharan, Robert R. McLeod, and Maxim Shusteff. Highly Tunable Thiol-Ene Photoresins for Volumetric Additive Manufacturing. *Advanced Materials*, 32(47):2003376, 2020. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/adma.202003376>.
- [29] Martin Regehly, Yves Garmshausen, Marcus Reuter, Niklas F. König, Eric Israel, Damien P. Kelly, Chun-Yu Chou, Klaas Koch, Baraa Asfari, and Stefan Hecht. Xolography for linear volumetric 3D printing. *Nature*, 588(7839):620–624, December 2020. Number: 7839 Publisher: Nature Publishing Group.
- [30] Antony Orth, Kathleen L. Sampson, Kayley Ting, Kayley Ting, Jonathan Boisvert, and Chantal Paquet. Correcting ray distortion in tomographic additive manufacturing. *Optics Express*, 29(7):11037–11054, March 2021. Publisher: Optical Society of America.
- [31] Brett Kelly, Indrasen Bhattacharya, Maxim Shusteff, Robert M. Panas, Hayden K. Taylor, and Christopher M. Spadaccini. Computed Axial Lithography (CAL): Toward Single Step 3D Printing of Arbitrary Geometries. *arXiv:1705.05893 [cs]*, May 2017. arXiv: 1705.05893.

- [32] Chi Chung Li, Joseph Toombs, and Hayden Taylor. Tomographic color Schlieren refractive index mapping for computed axial lithography. In Symposium on Computational Fabrication, SCF '20, pages 1–7, New York, NY, USA, November 2020. Association for Computing Machinery.
- [33] Damien Loterie and Paul Delrot. VOLUMETRIC 3D PRINTING OF ELASTOMERS BY TOMOGRAPHIC BACK-PROJECTION. page 11, 2018.
- [34] Pantazis Mouroulis and John Macdonald. Geometrical Optics and Optical Design. Oxford University Press, 1997. Google-Books-ID: Qx5vsmVyGK0C.
- [35] James W. Wydra, Neil B. Cramer, Jeffrey W. Stansbury, and Christopher N. Bowman. The reciprocity law concerning light dose relationships applied to BisGMA/TEGDMA photopolymers: Theoretical analysis and experimental characterization. Dental Materials, 30(6):605–612, June 2014.
- [36] Nicole Stephenson Kenning, Beth A Ficek, Cindy C Hoppe, and Alec B Scranton. Spatial and temporal evolution of the photoinitiation rate for thick polymer systems illuminated by polychromatic light: selection of efficient photoinitiators for LED or mercury lamps. Polymer International, 57(10):1134–1140, 2008. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/pi.2455>.
- [37] Darren L. Forman, Michael C. Cole, and Robert R. McLeod. Radical diffusion limits to photoinhibited superresolution lithography. Physical Chemistry Chemical Physics, 15(36):14862–14867, August 2013. Publisher: The Royal Society of Chemistry.
- [38] T. R. Waite. Theoretical Treatment of the Kinetics of Diffusion-Limited Reactions. Physical Review, 107(2):463–470, July 1957. Publisher: American Physical Society.
- [39] Asais Camila Uzcategui, Archish Muralidharan, Virginia L. Ferguson, Stephanie J. Bryant, and Robert R. McLeod. Understanding and Improving Mechanical Properties in 3D printed Parts Using a Dual-Cure Acrylate-Based Resin for Stereolithography. Advanced Engineering Materials, 20(12):1800876, 2018. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/adem.201800876>.
- [40] Frank Natterer. The Mathematics of Computerized Tomography. SIAM, June 2001. Google-Books-ID: pZMdcI0WZxcC.
- [41] Maxim Shusteff, Allison E. M. Browar, Brett E. Kelly, Johannes Henriksson, Todd H. Weisgraber, Robert M. Panas, Nicholas X. Fang, and Christopher M. Spadaccini. One-step volumetric additive manufacturing of complex polymer structures. Science Advances, 3(12):eaao5496, December 2017. Publisher: American Association for the Advancement of Science Section: Research Article.
- [42] computed-axial-lithography/CAL-software, August 2020. original-date: 2019-10-02T06:52:14Z.
- [43] Ribin Varghese Pazhamannil, P. Govindan, and Abhilash Edacherian. Optimized projections and dose slices for the volumetric additive manufacturing of three dimensional objects. Materials Today: Proceedings, December 2020.

- [44] Wim van Aarle, Willem Jan Palenstijn, Jeroen Cant, Eline Janssens, Folkert Bleichrodt, Andrei Dabrovolski, Jan De Beenhouwer, K. Joost Batenburg, and Jan Sijbers. Fast and flexible X-ray tomography using the ASTRA toolbox. Optics Express, 24(22):25129, October 2016.
- [45] Max Viergever and Andrew Todd-Pokropek, editors. Mathematics and Computer Science in Medical Imaging. Nato ASI Subseries F:. Springer-Verlag, Berlin Heidelberg, 1988.
- [46] Feng Xu, Lukas Helfen, Tilo Baumbach, and Heikki Suhonen. Comparison of image quality in computed laminography and tomography. Optics Express, 20(2):794–806, January 2012. Publisher: Optical Society of America.
- [47] David J. Glugla, Madeline B. Chosy, Marvin D. Alim, Amy C. Sullivan, and Robert R. McLeod. Transport-of-intensity-based phase imaging to quantify the refractive index response of 3D direct-write lithography. Optics Express, 26(2):1851–1869, January 2018. Publisher: Optica Publishing Group.
- [48] Gregory I. Peterson, Johanna J. Schwartz, Di Zhang, Benjamin M. Weiss, Mark A. Ganter, Duane W. Storti, and Andrew J. Boydston. Production of Materials with Spatially-Controlled Cross-Link Density via Vat Photopolymerization. ACS applied materials & interfaces, 8(42):29037–29043, October 2016.
- [49] Xiao Kuang, Jiangtao Wu, Kaijuan Chen, Zeang Zhao, Zhen Ding, Fengjingyang Hu, Daining Fang, and H. Jerry Qi. Grayscale digital light processing 3D printing for highly functionally graded materials. Science Advances, 5(5):eaav5790, May 2019.
- [50] Nawal Alharbi, Reham Osman, and Daniel Wismeijer. Effects of build direction on the mechanical properties of 3D-printed complete coverage interim dental restorations. The Journal of Prosthetic Dentistry, 115(6):760–767, June 2016.
- [51] Kavin Kowsari, Biao Zhang, Sahil Panjwani, Zaichun Chen, Hardik Hingorani, Saeed Akbari, Nicholas X. Fang, and Qi Ge. Photopolymer formulation to minimize feature size, surface roughness, and stair-stepping in digital light processing-based three-dimensional printing. Additive Manufacturing, 24:627–638, December 2018.
- [52] Aniruddha V. Shembekar, Yeo Jung Yoon, Alec Kanyuck, and Satyandra K. Gupta. Trajectory Planning for Conformal 3D Printing Using Non-Planar Layers. American Society of Mechanical Engineers Digital Collection, November 2018.
- [53] Indrasen Bhattacharya, Joseph Toombs, and Hayden Taylor. High fidelity volumetric additive manufacturing. Additive Manufacturing, 47:102299, November 2021.
- [54] Ra’ed Malallah, Haoyu Li, Damien P. Kelly, John J. Healy, and John T. Sheridan. A Review of Hologram Storage and Self-Written Waveguides Formation in Photopolymer Media. Polymers, 9(8):337, August 2017. Number: 8 Publisher: Multidisciplinary Digital Publishing Institute.
- [55] Manabu Kagami, Tatsuya Yamashita, and Hiroshi Ito. Light-induced self-written three-dimensional optical waveguide. Applied Physics Letters, 79(8):1079–1081, August 2001. Publisher: American Institute of Physics.
- [56] Anthony S. Kewitsch and Amnon Yariv. Self-focusing and self-trapping of optical beams upon photopolymerization. Optics Letters, 21(1):24–26, January 1996. Publisher: Optical Society of America.

- [57] Ra'ed Malallah, Derek Cassidy, Inbarasan Muniraj, James P. Ryle, John J. Healy, and John T. Sheridan. Self-written waveguides in photopolymer. Applied Optics, 57(22):E80, August 2018.
- [58] A. J. Jacobsen, W. Barvosa-Carter, and S. Nutt. Micro-scale Truss Structures formed from Self-Propagating Photopolymer Waveguides. Advanced Materials, 19(22):3892–3896, 2007. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/adma.200700797>.
- [59] Ra'ed Malallah, Haoyu Li, Inbarasan Muniraj, Derek Cassidy, Nebras Al-Attar, John J. Healy, and John T. Sheridan. Controlling the trajectories of self-written waveguides in photopolymer. Journal of the Optical Society of America B, 35(8):2046, August 2018.
- [60] Nicolas Hô, Bruno Bourliaguet, Jacques M. Laniel, Réal Vallée, and Alain Villeneuve. Observation of filamentation in As₂S₃. In Nonlinear Guided Waves and Their Applications, page WC8, Clearwater, Florida, 2001. OSA.
- [61] Yu Jian, Yong He, Liyan Zhao, Agnieszka Kowalczyk, Wantai Yang, and Jun Nie. Effect of Monomer Structure on Real-Time UV-Curing Shrinkage Studied by a Laser Scanning Approach. Advances in Polymer Technology, 32(1), 2013. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/adv.21331>.
- [62] Devdatt L. Kurdikar and Nikolaos A. Peppas. The volume shrinkage, thermal and sorption behaviour of polydiacrylates. Polymer, 36(11):2249–2255, January 1995.

Appendix A

OSMO Sample Code

```

%----- SETUP -----
% Get in-part and out-of-part indices
gelInds    = find(ft==1);          % Where ft is a 2D binary target geometry
voidInds   = find(ft==0);

N          = size(ft,1);          % Parameters for Matlab's radon()
theta     = 0:179;              N_theta = length(theta);

% Make initial model:
% Apply 2D Ramp Filter to target to generate initial OSMO model.
% Target is square with an even side-length.
[Nx, Ny]   = size(ft);
fx_fx      = (-Nx/2):1:(Nx/2-1);
fy_fy      = (-Ny/2):1:(Ny/2-1);
[fx, fy]   = meshgrid(fx_fx, fy_fy);
R2D        = sqrt(fx.^2 + fy.^2); %The 2D ramp filter

f_ft       = fftshift(fft2(ft));
fm         = real(ifft2(ifftshift(f_ft.*R2D)));
fm         = fm./max(fm(:));     %Initial model

% Get an initial reconstruction, f
g          = radon(fm,theta);    g(g<0) = 0;          % Forward project & clip negatives
f         = iradon(g, theta, 'none', 1.0, N);      % Back project
f         = f./max(f(:));       f(f<0)=0;         % Normalize and clip negatives

%----- OSMO ALGORITHM -----
for ii = 1:Nitr
    %Step 1: Change map to reduce void dose.
    voidDif    = f(voidInds)-Dcl;          %Extra out-of-part dose
    voidDif(voidDif<0) = 0;
    fm(voidInds) = fm(voidInds) - voidDif;
    %Make New Reconstruction
    g          = radon(fm,theta);    g(g<0) = 0;          % Forward project & clip negatives
    f         = iradon(g, theta, 'none', 1.0, N);
    f         = f./max(f(:));       f(f<0)=0;         % Normalize and clip negatives

    %Step 2: Change map to increase in-part (gel) dose.
    gelDif     = Dch - f(gelInds);        %Missing in-part dose
    gelDif(gelDif<0) = 0;
    fm(gelInds) = fm(gelInds) + gelDif;
    %Make New Reconstruction
    g          = radon(fm,theta);    g(g<0) = 0;          % Forward project & clip negatives
    f         = iradon(g, theta, 'none', 1.0, N);
    f         = f./max(f(:));       f(f<0)=0;         % Normalize and clip negatives
end
% g is the final writing 1D image-set, producing the dose-reconstruction, f.

```

Figure A.1: Sample OSMO code implemented in Matlab