

Lecture Notes on Symmetry Optics

Lecture 7:

The Symmetry Optics Validation Study

To accompany <https://youtu.be/5wzwEnF-nFI>

v1, April 13, 2022

Paul Mirsky

paulmirsky633@gmail.com, symmetryoptics.com

1 Laboratory setup


1.1 Introduction

Welcome to Lectures on Symmetry Optics. I'm Paul Mirsky. This is Lecture 7 of the series, and the topic is: The Symmetry Optics Validation Study

Over the course of the past 6 lectures, we've introduced a very novel way of thinking about light. It's distinct from wave optics, and it operates according to its own strange logic. We've claimed that symmetry optics works to describe diffraction and interference. But what is the evidence for such a claim? Why should anyone believe it? Why not just dismiss it as nonsense?


This lecture will lay out the evidence from the *Symmetry Optics Validation Study*. This is a research project to compare experimental data against the theoretical model which we've developed over the course of this series. If the data and the model agree, it's evidence in support of the theory.

1.2 Grating pattern at the flat



We'll start by discussing the experimental setup. This picture shows a grating film. A wide laser beam is illuminating the film from behind, so we are looking into the beam. The wavefronts are all flat. This is the starting point for Many-Slit interference.


1.3 Camera images, one plane



In this picture, the grating is over on the right side and the light is propagating from right to left. The device on the left is a camera. It captures an image of one single plane. Right now, it is imaging the plane of the grating.

What happens if we move the camera further away from the grating, say by 1mm? In one sense, the camera would see a *blurry* image of the *grating*. But actually it sees a *sharp* image of a *different plane*. Instead of imaging the grating plane, it instead images a theoretical plane which lies 1mm away from the grating.


1.4 Scanning over many planes



We can image any plane like this, by moving the camera. This camera is mounted on a mechanical stage, which moves the camera away from the grating in small steps. At each position, it captures an image of one plane.

We're going to assemble these individual images to create a video which shows the screen pattern evolving as a function of Z . This is what symmetry optics calculates, so we are getting data which we can compare against the model.

1.5 Lens-limited configuration



Next we need to introduce one additional level of complexity. Up to now, we have mostly studied the *free configuration*. That means, the light emerges from the grating and propagates out into free space. This is what we know how to calculate. But, it's very difficult to verify experimentally, because the light may propagate hundreds of meters, and the pattern there may be tens of meters wide. The camera only has an aperture of around 20mm, and the stage has only a couple hundred mm of travel, so it's too small to capture those distant planes.

To overcome these practical difficulties, we study the *lens-limited configuration* instead. To make this configuration, we place a lens one focal length away from the flat; in our experiment, this is 100mm away.

The lens does not affect any of the planes that come before it, so they are exactly the same as in the free configuration. But any plane *after the lens* gets imaged onto a different plane. The pattern that would have formed without the lens acts as a *virtual object*, and it forms a *real image*. You can calculate that using the usual thin-lens equation. The image is closer to the lens, and it's also demagnified.

The rear flat, which is one focal length beyond the lens, is an image of *infinity* in the free configuration. The lens effectively collapses an infinite amount of space into this small region. This configuration makes it practical to study interference with a small camera.

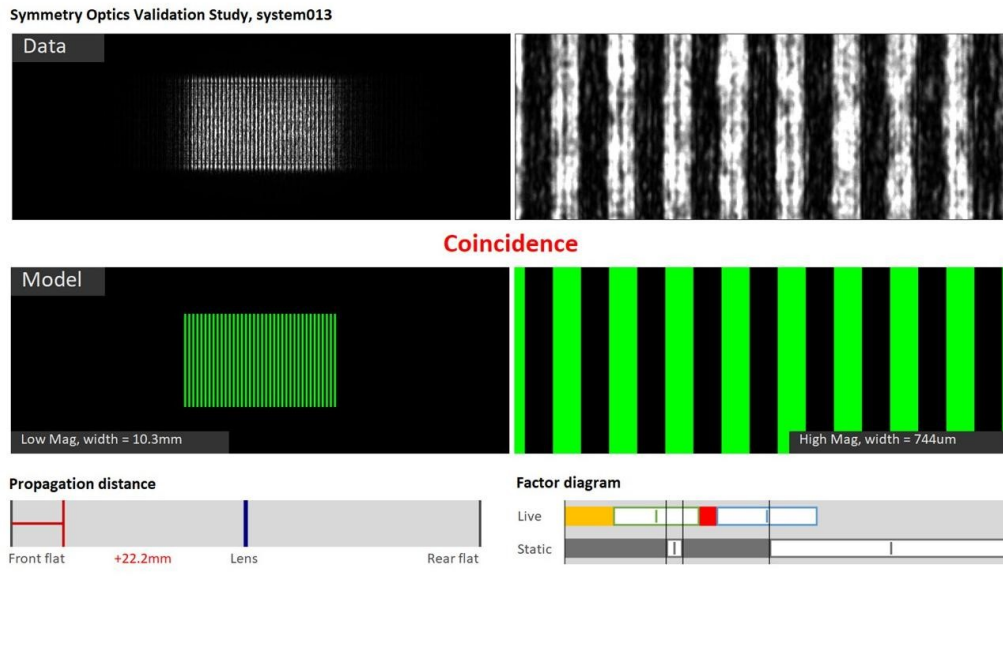
1.6 Experimental bench



We're not going to discuss all these components in detail. But, this is the experimental bench that gathered the study data. This gives you a general sense of what it looks like.

2 Experimental results

2.1 One frame




The result data is presented as a set of videos. Each video is made up of frames. Here is the format of a typical frame:

First, at the top are two windows showing actual images collected by the camera. The different windows show the same pattern in the same plane, but the left image is at low magnification, and the right image is at high magnification. This is very useful because the pattern has both larger features, and smaller features, and we need to be able to see both.

Next come two additional windows. These contain *simulated images*, which were calculated using the theoretical model. Each simulated image corresponds to the experimental image above it.

Next there is a note which describes the plane. There is also a continuous factor diagram, which shows the static and live at the current plane. And finally, each frame contains this element showing where the plane is in physical space. This the front flat, the lens, and the rear flat, and right now we are seeing the plane indicated by this vertical red line, which is 22.2mm after the front flat.

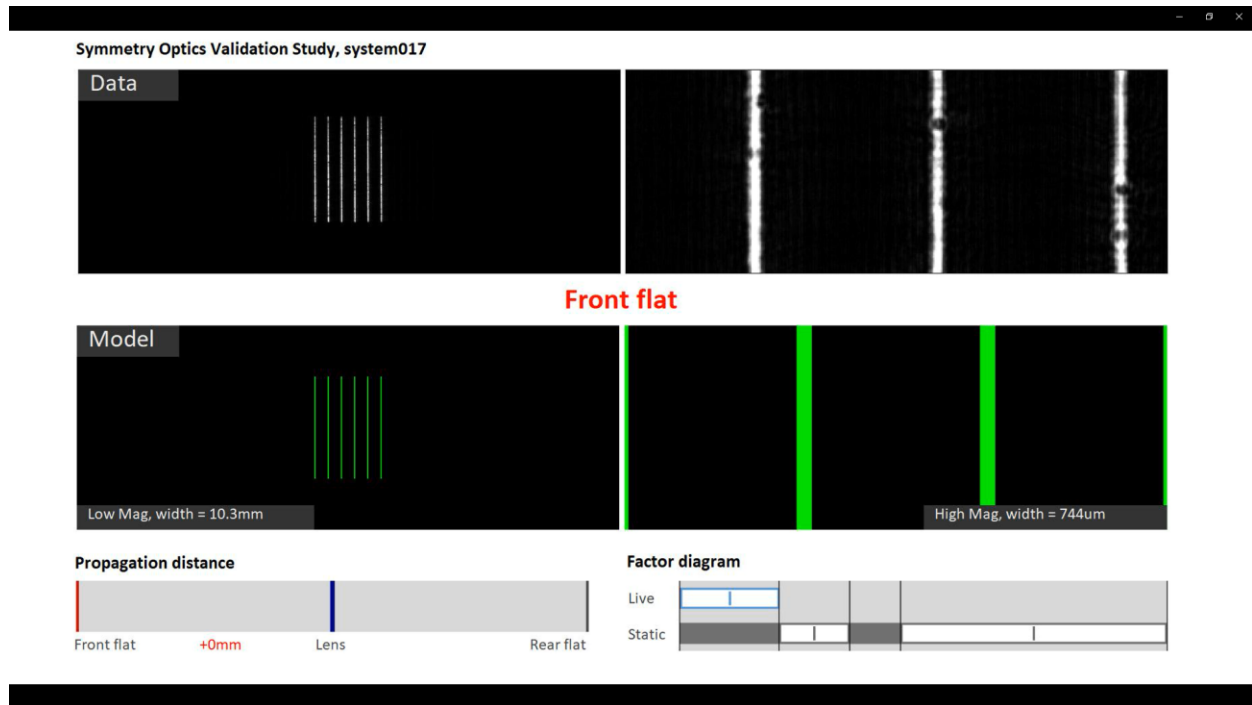
2.2 SOVaS system008



Let's watch one of the result videos. Here are the parameters.

- This is the starting pattern, directly at the grating.
- No change up to the early elbow.
- The stripe gets wider up to the core start.
- The core start is solid.
- Then the period gets wider approaching the interval elbow.
- And there are some coincidences.
- Interval elbow
- More coincidences
- Eventually, we get to breakout
- Then stripe and array both get wider, so now there are changes at both magnifications
- We pass the lens
- Then reach the core end
- Then the stripe gets narrower up to the late elbow
- And then no change up to the far limit, which is also the rear flat.


2.3 SOVaS system017



Here's another video; this time, $B > C$, so this one is early-breaking.

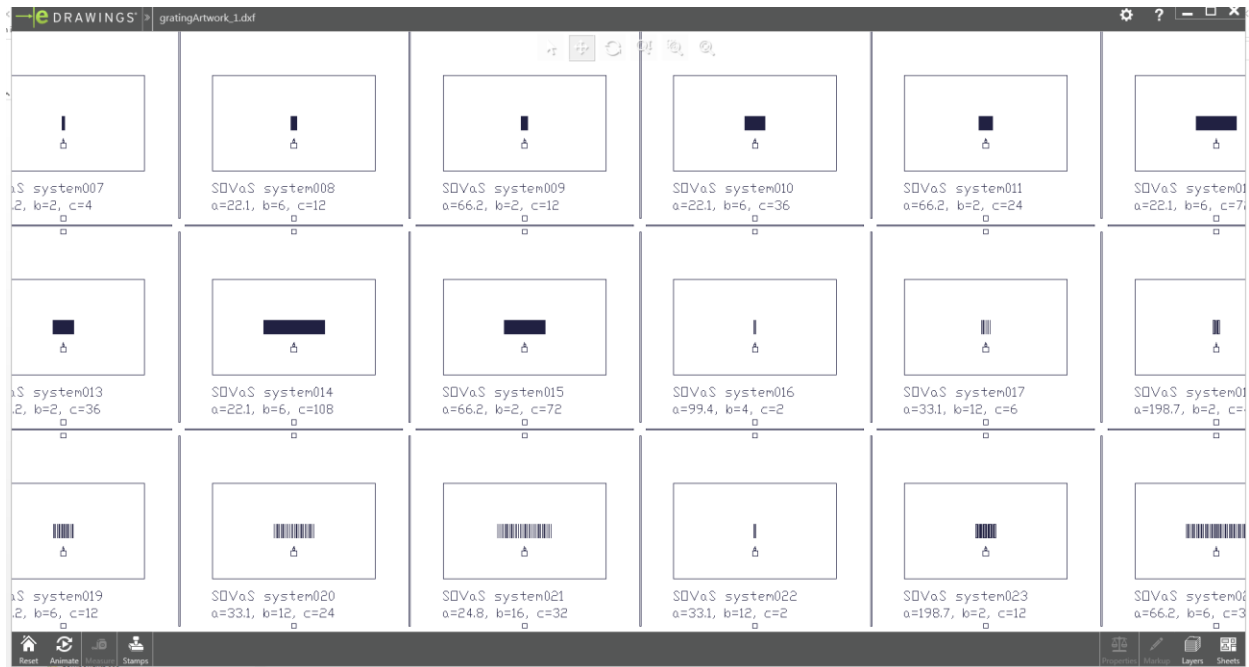
- Front flat
- Early elbow
- Core start
- Breakout
- Interval elbow, and passing the lens
- core end
- Late elbow
- And, the rear flat.

2.4 One grating



Each video begins with a different *grating*. This slide shows just one of them. It's printed on clear film using a very high-resolution printer. The smallest features can be about 15 microns wide.

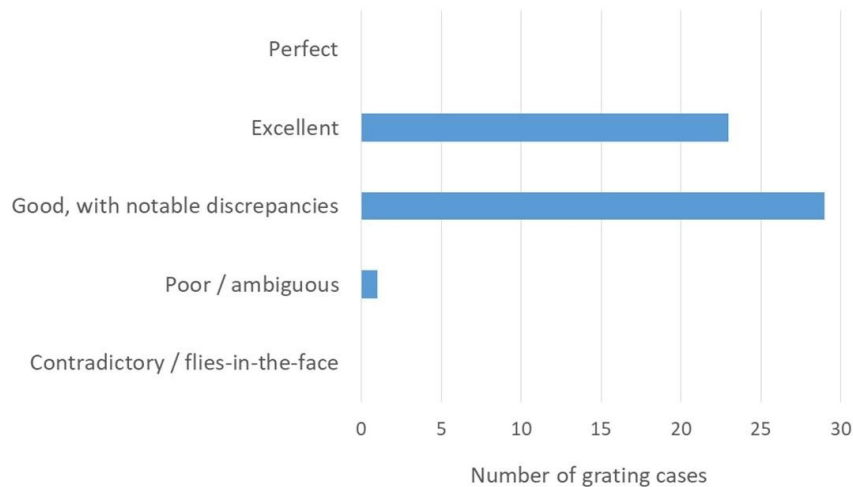
2.5 Set of gratings



The films were printed from .dxf files which you can also download for yourself. You can see that there are a range of different patterns, but they all follow the basic format of a stripe, a period, and an array. Some have quite large features that are visible with the naked eye; others have much smaller features that can only be seen under a microscope.

It's impossible to test every conceivable combination, because there are an infinite number. But we will validate a diverse collection of 64 different gratings. These are spread out all over the parameter space, in order to validate the model under the widest possible variety of different conditions, given the practical constraints.

2.6 Data-model agreement

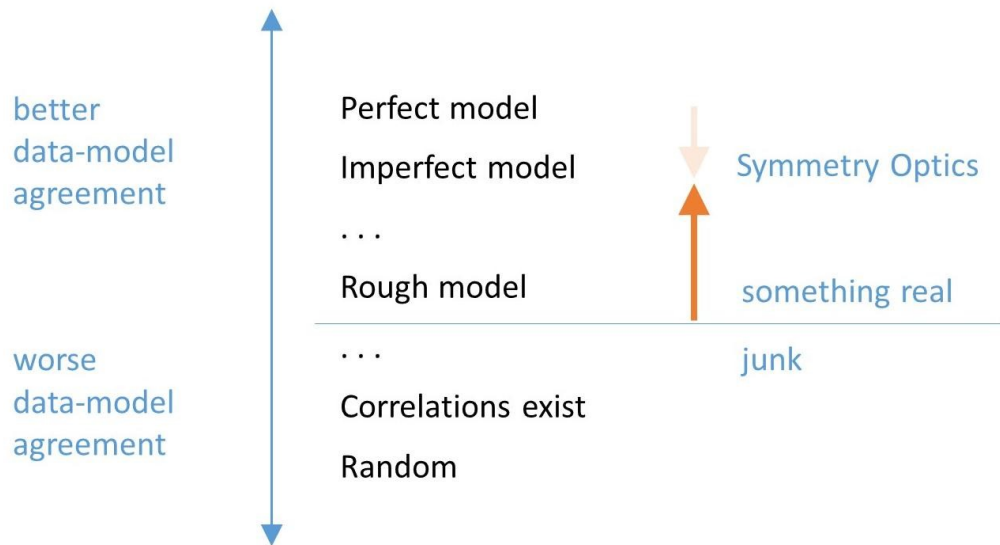


It's difficult to quantify how well the data matches the theory, so instead I subjectively rated all the samples on a scale with 5 levels: Perfect, Excellent, Good, Poor, and Contradictory.

None were ever truly perfect. But I classified 23 as excellent, 29 as good, and 1 as poor. The two I showed a moment ago were considered excellent. Also very importantly, not a single example was contradictory. That would have been if it showed clear and definitive trends that flew in the face of the model. Additionally, 11 systems were duplicates and I didn't score those separately.

These ratings are admittedly somewhat subjective, but you can look at the videos yourself and reach your own conclusions. The first few are the worst, so don't look only at those.

2.7 Degrees of data-model agreement



So, what is the verdict of the study? In the simplest understanding, an experiment can either prove or disprove a hypothesis. But in fact, it's quite a bit more complicated than that. We can think of data-model agreement as existing along a spectrum.

At the lowest end is randomness. Here, the data and the model are not correlated at all. This means that you haven't proven anything.

However, if you find correlations between the model and data, then you have something real. This is what you often find, for example, in medicine, where a certain pill may or may not have a statistically significant effect. This divide – between junk and something real – is the bare minimum threshold for saying, 'the evidence supports my claim'.

Much, much further up the scale, we have models with a mathematical form. Some of them work in a very rough way, such as models in economics or in epidemiology. Weather and climate models are also like this. These models are far from perfect, but they can guide important decisions because they are far, far better than randomness and even far better than a simple correlation like in medicine. For a real quantitative field like optics, the threshold between junk and something real, is really somewhere around here. It's a much stricter standard.

At the very top of the scale is the perfect model, which gets every detail correct with infinite accuracy. This is the ideal. Some models get very close to this.

But in fact, there is almost always some error, resulting in an imperfect. It may be a case like non-relativistic physics, which is almost perfect to within a miniscule fraction. Or, it may be a case where the agreement is really quite good, but nevertheless there exist noticeable discrepancies.

Symmetry optics falls into this range. We do notice discrepancies between the model and the data. And, we take careful note of these because they define the boundary within which the model works. But they're not so critical at this point, because we're concerned with proving a weaker claim, which is that Symmetry Optics is *something real*. And the evidence absolutely does support this claim. Symmetry optics is working well – actually, very well. It's something real.

When I say 'it's something real', what I mean is that it's worthy of investigation by other scientists.

2.8 Why study symmetry optics?

These scientists, these researchers, are really the audience that I'm addressing all of this to. I'm hoping to persuade you to start investigating symmetry optics yourself, and to make your own contributions to this new field.

You can think of Symmetry Optics like an unexplored desert island. It's unknown precisely what's there, but it's very clear that there *is* territory there to explore. In physical sciences, if a model succeeds for anything, then it can probably be extended. Symmetry Optics succeeds in giving an elegant, clean description of interference. There has to be a powerful principle underneath it, even if we don't fully understand it yet.

I believe that there is much more to symmetry optics. Some of it, I've already found. The rest is waiting for pioneers like you to discover. Don't wait for consensus, because at that point, there won't be anything left for you to discover.

3 Technical notes

3.1 Overview

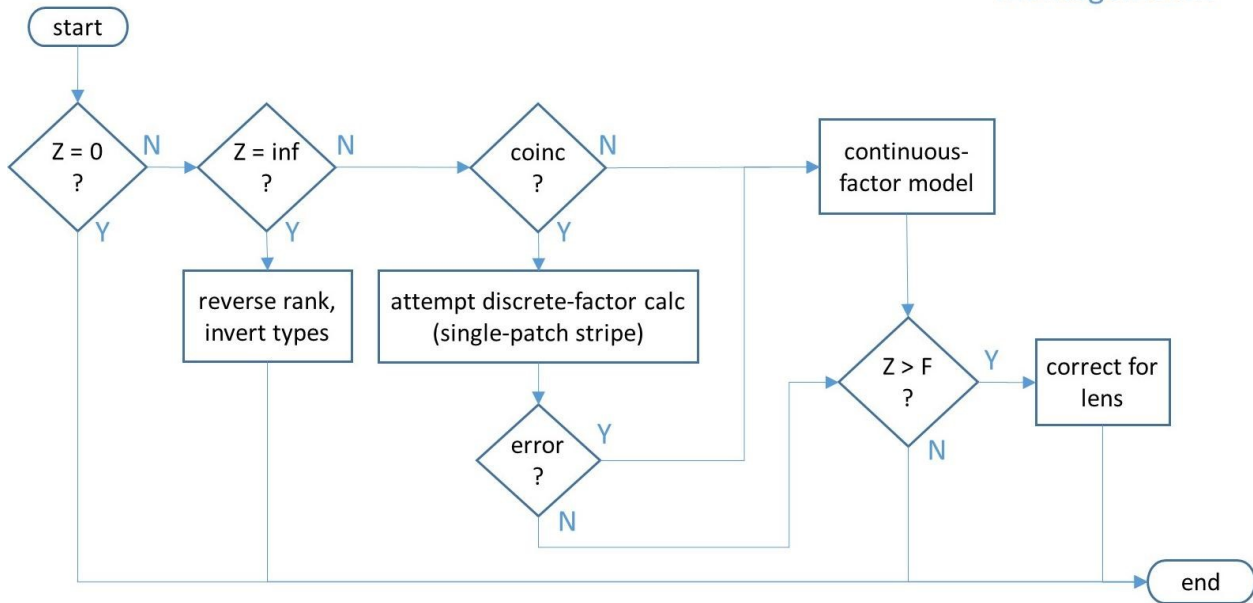
For the final section of this lecture, we'll go over some technical notes.

1. Algorithm montage
2. No solid core end for $C = 2$
3. Pinch before early elbow is only a display artifact

Two of these explain discrepancies between the data and the model. They're not serious issues, but they are points for potential criticism and so they need to be addressed proactively.

3.2 Algorithm montage

montageMSI.m



First, let's discuss the algorithm montage. The theoretical 'model' that generates the simulated images is actually a montage of several different algorithms. To see the precise details, I recommend downloading the companion code. This flowchart shows the general logic.

Each plane has a distance Z , which applies to the free configuration, *not to the actual lensed configuration that we are using. If Z is 0, then we don't need to do any calculating, because the pattern at the flat is already given as the input.

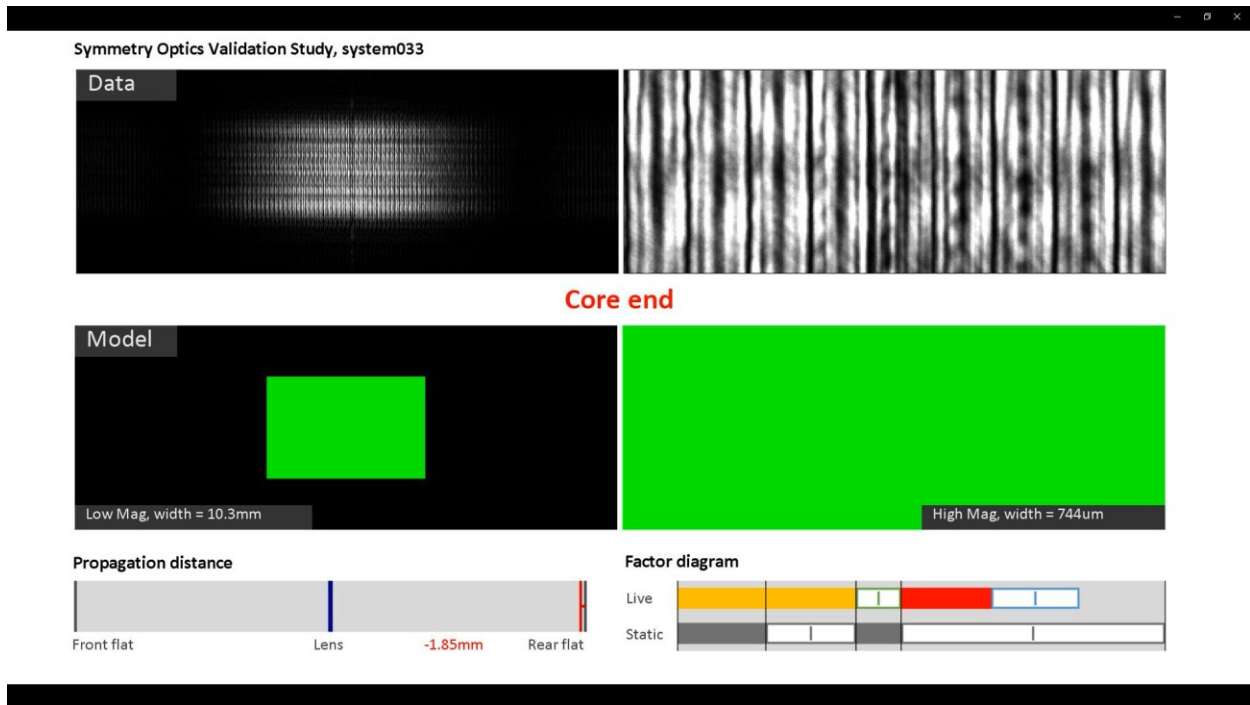
Next, if Z is at infinity, or over a certain threshold, then it corresponds to the rear flat in the lensed configuration. For this special case, the calculation is particularly easy – it's the rule we learned in Lecture 2. We reverse the rank, and invert the types, in the factor chain, and that tells us the pattern.

If Z is not 0 or infinity, the algorithm branches depending on whether or not it's a coincidence. If it's not a coincidence, we calculate using the continuous-factor model. Then, we check whether Z is greater than the focal length, which is to say, whether it's before or after the lens. If it's before the lens, we're done. If it's after the lens, we need to correct for the lens by finding the plane that it gets imaged into, and scaling by the magnification, and then we have an answer.

Finally, there is the case of coincidences. When you call this function, you flag whether or not the plane is a coincidence. This was set up to flag a coincidence for any integer multiple of the core start. For practical reasons, we only flag coincidences for late-breaking gratings, between the core start and breakout. If the plane is flagged as a coincidence, we attempt to calculate it

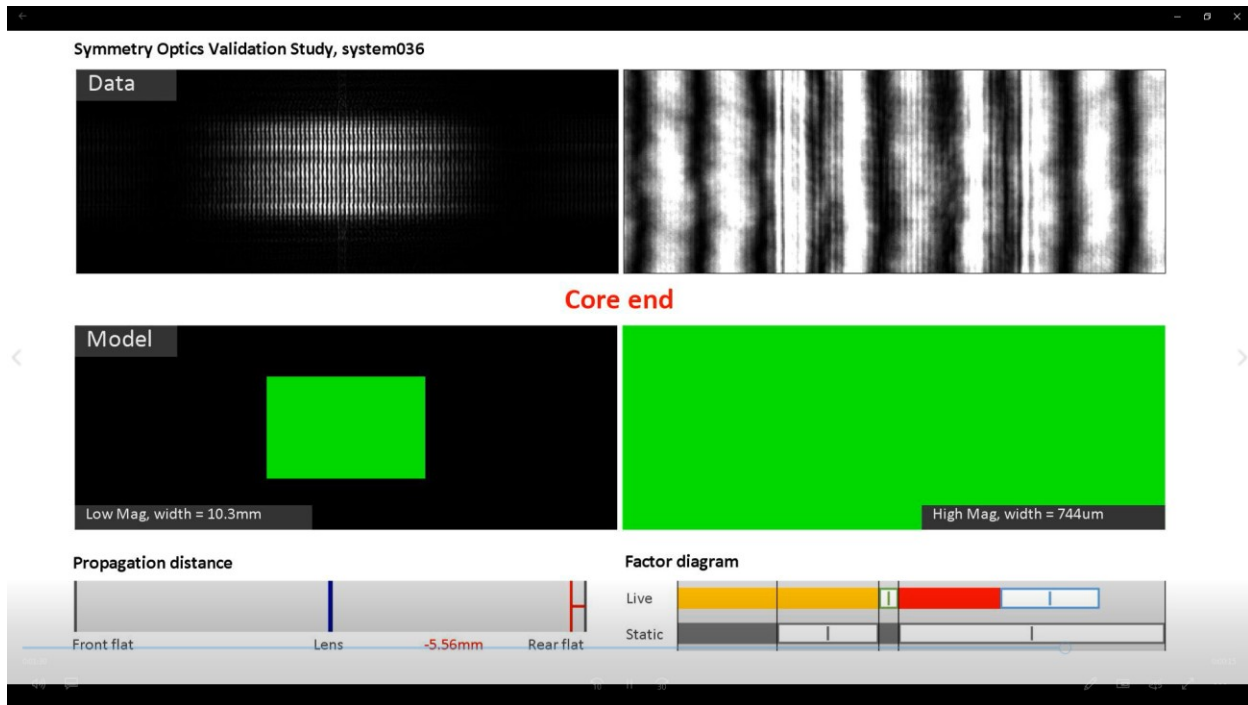
using the discrete-factor model. To make it tractable, we actually use a scaling trick where we scale the input pattern so that the stripe is just one patch wide; the details of this trick are the code. If the discrete-factor calculation throws an error, then we fall back on the continuous model instead. If there is no error, then we proceed to correct for the lens, and then we are done.

3.3 Solid core end, C = 6



The next technical note concerns a problem that occurs at the core end. First, let's see an ordinary case where the problem does *not* occur. Note that because of the lens, it's a little different from the free configuration that you're familiar with. Both in the model and in the data, the period decreases until finally, at the core end, the pattern is a single solid bright region.

3.4 No solid core end when $C = 2$

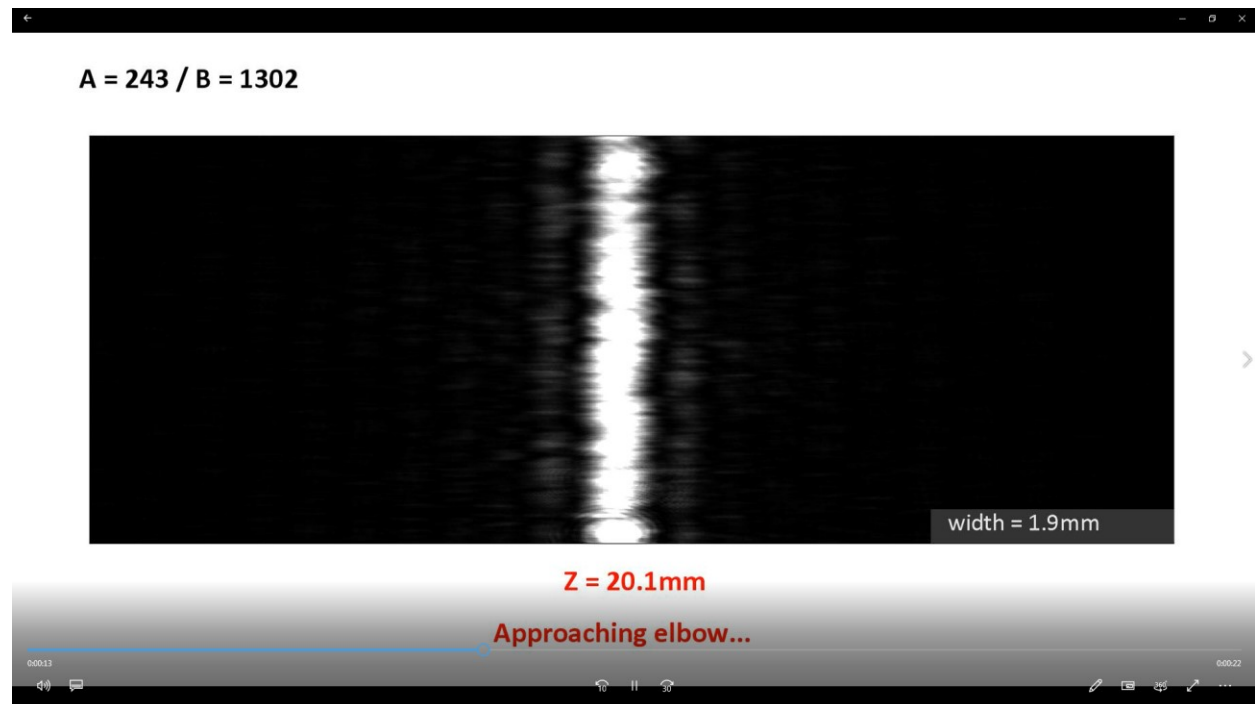


But in another grating when $C = 2$, when there are just two starting slits, the data and the model *don't* agree. In this case, the model still predicts a solid core end, but the data does not match the model. To explain this, we need to go beyond what we've learned so far and broach the topic of *symmetry*. Surprisingly, we've come quite far in this subject while barely mentioning symmetry, even though it's so fundamental as to give the field its name. The symmetry of a pattern is described mathematically by a *transformation* that leaves the pattern looking the same. In this case, the transformation is to slide each stripe by a distance of one period, so that each stripe takes the place of its neighbor.

For an infinite series of stripes, this is a perfect symmetry. For a large but finite set of stripes, the symmetry is slightly imperfect, because it fails to apply for the first and last stripes, which are each missing a neighbor. For fewer stripes, this minor imperfection becomes much more significant, and when there are only two stripes, it's a substantial problem, and so the model doesn't work as well.

This may be counterintuitive, because in conventional optics the two-slit experiment is often studied as a paradigm example. The fact that there are only two is an advantage, because it keeps the problem simple. But in symmetry optics, two slits is actually an edge case, where the model is on the verge of breaking, and this manifests in the failure to form a solid core end. In the future, we'll learn more about symmetry, and it will become integral to how we understand symmetry optics.

3.5 Pinching before early elbow, as image

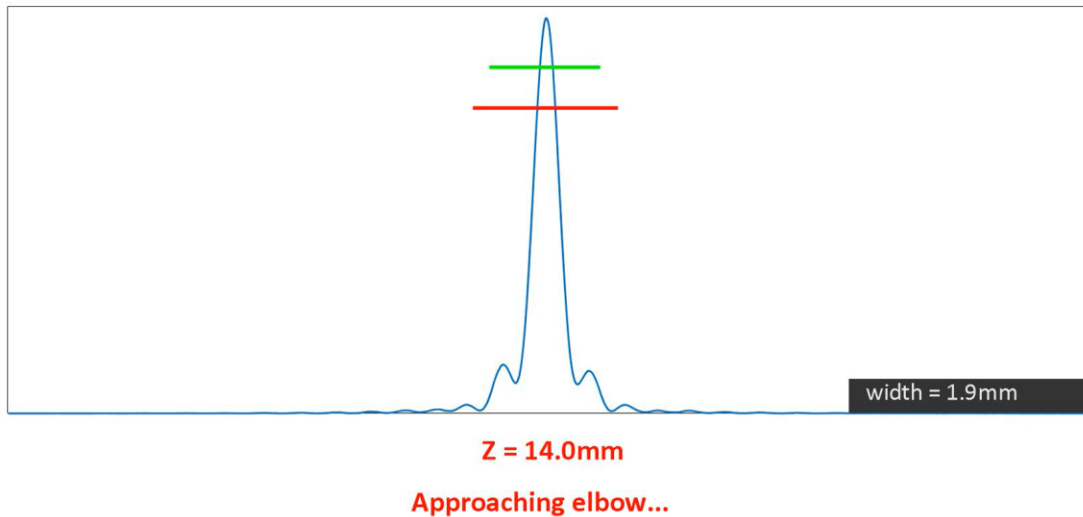


The final note concerns *pinching*. This effect occurs between the flat and the early elbow. The stripe starts at one width, then it appears to grow narrower, roughly by half. Then, it grows larger again and it gets back to its original width at the early elbow. This disagrees with the symmetry-optical model, which says that the stripe width should stay essentially constant during this stage. But we also know this isn't right even according to conventional optics, because when light passes through a slit it's supposed to diffract and get slightly wider.

Actually, the pinching effect is just a display artifact. It occurs because the dimmer pixels are hard to see, and it's difficult to determine visually where the edge is, especially when the edge is fuzzy.

3.6 Pinching before early elbow, as histogram

A = 243 / B = 1302



To see this quantitatively, we can plot it as a 1-d function of power vs position. The red bar represents the width as determined by four standard deviations of the distribution.

When we advance, we see that a narrow spike forms right in the center. This is what appears as pinching. But actually, the calculated width is almost the same or even slightly wider than the green bar, which is the width at the flat. This occurs partly because even small amounts of light can have a large effect on standard deviation when they are far out from the center. So these side-lobes, and even these are having a substantial effect.

In fact, there are multiple ways to define the width of a shape like this. They are all reasonably close. The important point is simply that the pinching is not a serious challenge to the symmetry-optical model, or to the conventional idea that light spreads out after passing through a slit.

4 Conclusions

4.1 More resources

Wherever you found this video, you'll also find links to more resources to help you understand the Symmetry Optics Validation study, and to evaluate it for yourself.

There are lecture notes available, as always.

You can also download the companion code, which will let you follow every step of the model in detail.

The model-and-data videos are all posted up on YouTube, so you can watch them yourself and form your conclusions.

I'm also providing some of the raw data used to make the videos. There are many gigabytes of raw data, so I'm providing a subset that includes all of the critical planes and all of the coincidences, for all of the videos. This is enough to let you confirm the results yourself, if you want to.

Finally, you can download the engineering data needed to print the gratings, if you'd like to try these experiments yourself.

4.2 The future of symmetry optics, outro

This marks the end of the first ‘season’ of Lectures on Symmetry Optics, if you want to think of it that way. The lectures up to now give you enough to understand the basics, calculate an answer, and verify for yourself that it works. You’ve also seen the power and the elegance of looking at light this way.

But this is the beginning of the subject, not the end. In the future, I’ll post more videos, or possibly I’ll post papers instead.

But more importantly, I hope that *you* will be inspired to do your own research and add to the field. Symmetry optics is still mostly uncharted territory, and a lot of fruit is hanging very low. It’s a great opportunity for you to make original discoveries.

Here are just a few ideas of what you might investigate: try relating Symmetry Optics to conventional wave optics. Why is there no concept of phase? What does that tell us about destructive interference? Or, relate Symmetry Optics to the Fast Fourier Transform. They share something very deep, and seem to be made of the same kind of logic. Surely each one could give us some insight about the other. Most importantly of all, how can we generalize Symmetry Optics so that it works for arbitrary patterns, rather than just repeating stripe patterns? And while it’s probably too soon for this, are there any technological or engineering applications? These are just a few ideas.

Undoubtedly, many more questions have occurred to you. My advice is to follow your curiosity and investigate. I hope you discover something great. I’ll be rooting for you!

I’m Paul Mirsky, thanks for listening.